FICO® Xpress Optimizer

45.01

REFERENCE MANUAL

FICO® Xpress Optimization

FICO®

Xpress Optimizer 45.01 (FICO® Xpress 9.7)

Deliverable Version: A

Last Revised: 29 July, 2025

Contents

Chapter 1

Introduction

The FICO Xpress Optimization Suite is a powerful mathematical optimization framework well-suited to a broad range of optimization problems. The core solver of this suite is the FICO Xpress Optimizer, which combines ease of use with speed and flexibility. It can be interfaced via the command line Console Optimizer, via the graphical interface application IVE and through a callable library that is accessible from all the major programming platforms. It combines flexible data access functionality and optimization algorithms, using state-of-the-art methods, which enable the user to handle the increasingly complex problems arising in industry and academia.

The Console Optimizer provides a suite of 'Console Mode' Optimizer functionality. Using the Console Optimizer the user can load problems from widely used problem file formats such as the MPS and LP formats and solve them using any of the algorithms supported by the Optimizer. The results may then be processed and viewed in a variety of ways. The Console Mode provides full access to the Optimizer control variables allowing the user to customize the optimization algorithms to tune the solving performance on the most demanding problems.

The FICO Xpress Optimizer library provides full, high performance access to the internal data structures of the Optimizer and full flexibility to manipulate the problem and customize the optimization process. For example, the Cut Manager framework allows the user to exploit their detailed knowledge of the problem to generate specialized cutting planes during branch and bound that may improve solving performance of Mixed Integer Programs (MIPs).

Of most interest to the library users will be the embedding of the Optimizer functionality within their own applications. The available programming interfaces of the library include interfaces for C/C++, .NET and Java. Note that the interface specifications in the following documentation are given exclusively in terms of the C/C++ language. Short examples of the interface usage using other languages may be found in the FICO Xpress Getting Started manual.

1.1 The FICO Xpress Optimizer

The FICO Xpress Optimizer is a mathematical programming framework designed to provide high performance solving capabilities. Problems can be loaded into the Optimizer via matrix files such as MPS and LP files and/or constructed in memory and loaded using a variety of approaches via the library interface routines. Note that in most cases it is more convenient for the user to construct their problems using FICO Xpress Mosel or FICO Xpress BCL and then solve the problem using the interfaces provided by these packages to the Optimizer.

The solving algorithms provided with the Optimizer include the primal simplex, the dual simplex, the hybrid gradient and the Newton barrier algorithms. For solving Mixed Integer Programs (MIPs) the Optimizer provides a powerful branch and bound framework. The various types of problems the Optimizer can solve are outlined in Chapter 3.

Solution information can be exported to file using a variety of ASCII and binary formats or accessed via memory using the library interface. Advanced solution information, such as solution bases, can be both exported and imported either via files or via memory, using the library interface. Note that bases can be useful for so called 'warm-starting' the solving of Linear Programming (LP) problems.

1.2 Starting the First Time

We recommend that new FICO Xpress Optimizer users first try running the Console Optimizer 'optimizer' executable from the command prompt before using the other interfaces of Optimizer. This is because (i) it is the easiest way to confirm the license status of the installation and (ii) it is an introduction to a powerful tool with many uses during the development cycle of optimization applications. For this reason we focus mainly on discussing the Console Optimizer in this section as an introduction to various basic functions of the Optimizer.

1.2.1 Licensing

To run the Optimizer from any interface it is necessary to have a valid licence file, `xpauth.xpr`. The FICO Xpress licensing system is highly flexible and is easily configurable to cater for the user's requirements. The system can allow the Optimizer to be run on a specific machine, on a machine with a specific ethernet address or on a machine connected to an authorized hardware dongle.

If the Optimizer fails to verify a valid license then a message can be obtained that describes the reasons for the failure and the Optimizer will be unusable. When using the Console Optimizer the licensing failure message will be displayed on the console. Library users can call the function `XPRSgetlicerrmsg` to get the licensing failure message.

On Windows operating systems the Optimizer searches for the license file in the directory containing the Xpress libraries, which are installed by default into the `C:\xpressmp\bin` folder. To avoid unnecessary licensing problems, it is recommended that the `XPAUTH_PATH` environment variable is not set on Windows.

On Unix systems it is necessary to set the `XPAUTH_PATH` environment variable to the full path to the license file. For ease of support it is recommended that the license file is placed in the `bin` directory within your Xpress installation and the `XPAUTH_PATH` environment variable is set accordingly, e.g., `/opt/xpressmp/bin/xpauth.xpr`.

1.2.2 Starting the Console Optimizer

The Console Optimizer is an interactive command line interface to the Optimizer. The Console Optimizer is started from the command line using the following syntax:

```
C:\> optimizer [problem_name] [@filename]
```

From the command line an initial problem name can be optionally specified together with an optional second argument specifying a text "script" file from which the console input will be read as if it had been typed interactively.

Note that the syntax example above shows the command as if it were input from the Windows Command Prompt (i.e., it is prefixed with the command prompt string `C:\>`). For Windows users, the Console Optimizer can also be started by typing `optimizer` into the "Run ..." dialog box in the Start menu.

The Console Optimizer provides a quick and convenient interface for operating on a single problem loaded into the Optimizer. Compare this with the more powerful library interface that allows one or more problems to co-exist in a process. The Console Optimizer problem contains the problem data as well as (i) control variables for handling and solving the problem and (ii) attributes of the problem and its solution information.

Useful features of the Console Optimizer include a help system, auto-completion of command names and integration of system commands.

Typing "`help`" will list the various options for getting help. Typing "`help commands`" will list available commands. Typing "`help attributes`" and "`help controls`" will list the available attributes and controls, respectively. Typing "`help`" followed by a command name or control/attribute name will list the help for this item. For example, typing "`help lpoptimize`" will get help for the `LPOPTIMIZE` command.

The Console Optimizer auto-completion feature is a useful way of reducing key strokes when issuing commands. To use the auto-completion feature, type the first part of an optimizer command name followed by the Tab key. For example, by typing "`lpopt`" followed by the Tab key, the Console Optimizer will complete to the `LPOPTIMIZE` command. Note that once you have finished inputting the command name portion of your command line, the Console Optimizer can also auto-complete on file names. For example, if you have a matrix file named `hpw15.mps` in the current working directory then by typing "`readprob hpw`" followed by the Tab key the command should auto-complete to the string "`readprob hpw15.mps`". Note that the auto-completion of file names is case-sensitive.

The Console Optimizer also features integration with the operating system's shell commands. For example, by typing "`dir`" (or "`ls`" under Unix) you will directly run the operating system's directory listing command. Using the "`cd`" command will change the working directory, which will be indicated in the prompt string:

```
[xpress bin] cd \
[xpress C:\]
```

Finally, note that when the Console Optimizer is first started it will attempt to read in an initialization file named `optimizer.ini` from the current working directory. This is an ASCII file that may contain any lines that can normally be entered at the console prompt, such as commands or control parameter settings. The lines of the `optimizer.ini` file are run with at start up, and can be useful for setting up a customized default Console Optimizer environment for the user (e.g., defining custom controls settings on the Console Optimizer problem).

1.2.3 Scripting with the Console Optimizer

The Console Optimizer interactive command line hosts a TCL script parser (http://www.tcl.tk). With TCL scripting the user can program flow control into their Console Optimizer scripts. Also TCL scripting provides the user with programmatic access to a powerful suite of functionality in the TCL library. With scripting support the Console Optimizer provides a high level of control and flexibility well beyond that which can be achieved by combining operating system batch files with simple piped script files. Indeed, with scripting support, the Console Optimizer is ideal for (i) early application development, (ii) tuning of model formulations and solving performance and (iii) analyzing difficulties and bugs in models.

Note that the TCL parser has been customized and simplified to handle intuitive access to the controls and attributes of the Optimizer. The following example shows how to proceed with write and read access to the `MIPLOG` Optimizer control:

```
[xpress C:\] miplog=3
[xpress C:\] miplog
3
```

The following shows how this would usually be achieved using TCL syntax:

```
[xpress C:\] set miplog 3
3
[xpress C:\] $miplog
3
```

The following set of examples demonstrate how with some simple TCL commands and some basic flow control constructs the user can quickly and easily create powerful programs.

The first example demonstrates a loop through a list of matrix files where a simple regular expression on the matrix file name and a simple condition on the number of rows in the problem decide whether or not the problem is solved using `lpoptimize`. Note the use of:

- the creation of a list of file names using the TCL `glob` command
- the use of the TCL square bracket notation (`[]`) for evaluating commands to their string result value
- the TCL `foreach` loop construct iterating over the list of file names
- dereferencing the string value of a variable using '`$`'
- the use of the TCL `regexp` regular expression command
- the two TCL `if` statements and their condition statements
- the use of the two Optimizer commands `READPROB` and `MINIM`
- the TCL `continue` command used to skip to the next loop iteration

```
set filelist [glob *.mps]
foreach filename $filelist {
  if { [regexp -all {large_problem} $filename] } continue
  readprob $filename
  if { $rows > 200 } continue
  lpoptimize
}
```

The second example demonstrates a loop though some control settings and the tracking of the control setting that gave the best performance. Note the use of:

- the TCL `for` loop construct iterating over the values of variable `i` from `-1` to `3`
- console output formatting with the TCL `puts` command
- setting the values of Optimizer controls `CUTSTRATEGY` and `MAXNODE`
- multiple commands per line separated using a semicolon
- use of the `MIPSTATUS` problem attribute in the TCL `if` statement
- comment lines using the hash character '`#`'

```
set bestnodes 10000000
set p hpw15
for { set i -1 } { $i <= 3 } { incr i } {
  puts "Solving with cutstrategy : $i"
  cutstrategy=$i; maxnode=$bestnodes
  readprob $p
  mipoptimize
  if { $mipstatus == 6 } {
    # Problem was solved within $bestnodes
    set bestnodes $nodes; set beststrat $i
  }
}
puts "Best cutstrategy : $beststrat : $bestnodes"
```

1.2.4 *Interrupting Console Optimizer*

Console Optimizer users may interrupt the running of the commands (e.g., `lpoptimize`) by typing Ctrl-C. Once interrupted, the Console Optimizer will return to its command prompt. If an optimization algorithm has been interrupted in this way, any solution process will stop at the first 'safe' place before returning to the prompt. Optimization iterations may be resumed by re-entering the interrupted command. Note that by using this interrupt-resume functionality the user has a convenient way of dynamically changing controls during an optimization run.

Lastly, note that "typing ahead" while the console is writing output to screen can cause Ctrl-C input to fail on some operating systems.

1.3 Manual Layout

So far the manual has given a basic introduction to the FICO Xpress Optimization Suite. The reader should be able to start the Console Optimizer command line tool and have the license verified correctly. They should also be able to enter some common commands used in the Console Optimizer (e.g., `READPROB` and `LPOPTIMIZE`) and get help on command usage using the Console Optimizer help functionality.

The remainder of this manual is divided into two parts. These are the first chapters up to but not including Chapter 8 and the remaining chapters from Chapter 8.

The first part of the manual, beginning with Chapter 2, provides a brief overview of common Optimizer usage, introducing the various routines available and setting them in the typical context they are used. This is followed in Chapter 3 by a brief overview of the types of problems that the FICO Xpress Optimizer can be used to solve. Chapter 4 provides a description of the solution methods and some of the more frequently used parameters for controlling these methods along with some ideas of how they may be used to enhance the solution process. Finally, Chapter 5 details some more advanced topics in Optimizer usage.

The second half of the manual is the main reference section. Chapter 8 details all functions in both the Console and Advanced Modes alphabetically. Chapters 9 and 10 then provide a reference for the various controls and attributes, followed by a list of Optimizer error and return codes in Chapter 11. A description of several of the file formats is provided in Appendix A.

CHAPTER 2

Basic Usage

The FICO Xpress Optimization Suite is a powerful and flexible framework catering for the development of a wide range of optimization applications. From the script-based Console Optimizer providing rapid development access to a subset of Optimizer functionality (Console Mode) to the more advanced, high performance access of the full Optimizer functionality through the library interface.

In the previous section we looked at the Console Optimizer interface and introduced some basic functions that all FICO Xpress Optimizer users should be familiar with. In this section we expand on the discussion and include some basic functions of the library interface.

2.1 Initialization

Before the FICO Xpress Optimization Suite can be used from any of the interfaces the Optimizer library must be initialized and the licensing status successfully verified. Details about licensing your installation can be found in FICO Xpress Installation Guide.

When the Console Optimizer is started from the command line the initialization and licensing security checks happen immediately and the results are displayed with the banner in the console window. For the library interface users, the initialization and licensing are triggered by a call to the library function `XPRSinit`, which must be made before any of the other Optimizer library routines can be successfully called. If the licensing security checks fail to *check out* a license then library users can obtain a string message explaining the issue using the function `XPRSgetlicerrmsg`.

Note that it is recommended that the users having licensing problems use the Console Optimizer as a means of checking the licensing status while resolving the issues. This is because it is the quickest and easiest way to check and display the licensing status.

Once the Optimizer functionality is no longer required the license and any system resources held by the Optimizer should be released. The Console Optimizer releases these automatically when the user exits the Console Optimizer with the `QUIT` or `STOP` command. For library users the Optimizer can be triggered to release its resources with a call to the routine `XPRSfree`, which will *free* the license checked out in the earlier call to `XPRSinit`.

```
{
  if(XPRSinit(NULL)) printf("Problem with XPRSinit\n");
  XPRSfree();
}
```

In general, library users will call `XPRSinit` once when their application starts and then call `XPRSfree` before it exits. This approach is recommended since calls to `XPRSinit` can have non-negligible (approx. 0.5 sec) overhead when using floating network licensing.

Although it is recommended that the user writes their code such that `XPRSinit` and `XPRSfree` are called only in sequence note that the routine `XPRSinit` may be called repeatedly before a call to `XPRSfree`. Each subsequent call to `XPRSinit` after the first will simply return without performing any tasks. In this case note that the routine `XPRSfree` must be called the same number of times as the calls to XPRSinit to fully release the resources held by the library. Only on the last of these calls to XPRSfree will the library be released and the license *freed*.

2.2 The Problem Pointer

The Optimizer provides a container or problem pointer to contain an optimization problem and its associated controls, attributes and any other resources the user may attach to help construct and solve the problem. Console Optimizer has one of these problem pointers that it uses to provide the user with loading and solving functionality. This problem pointer is automatically initialized when the Console Optimizer is started and release again when it is stopped.

In contrast to the Console Optimizer, library interface users can have multiple problem pointers coexisting simultaneously in a process. The library user creates and destroys a problem pointer using the routines XPRScreateprob and XPRSdestroyprob, respectively. In the C library interface, the user passes the problem pointer as the first argument in routines that are used to operate on the problem pointer's data. Note that it is recommended that the library user destroys all problem pointers before calling XPRSfree.

```
{
  XPRSprob prob;
  XPRScreateprob(&prob);
  XPRSdestroyprob(prob);
}
```

2.3 Logging

The Optimizer provides useful text logging messages for indicating progress during the optimization algorithms and for indicating the status of certain important commands such as XPRSreadprob. The messages from the optimization algorithms report information on iterations of the algorithm. The most important use of the logging, however, is to convey error messages reported by the Optimizer. Note that once a system is in production the error messages are typically the only messages of interest to the user.

Conveniently, the Console Optimizer automatically writes the logging messages for its problem pointer to the console screen. Although message management for the library users is more complicated than for Console Optimizer users, library users have more flexibility with the handling and routing of messages. The library user can route messages directly to file or they can intercept the messages via callback and marshal the message strings to appropriate destinations depending on the type of message and/or the problem pointer from which the message originates.

To write the messages sent from a problem pointer directly to file the user can call XPRSsetlogfile with specification of an output file name. To get messages sent from a problem pointer to the library user's application the user will define and then register a messaging callback function with a call to the XPRSaddcbmessage routine.

```
{
  XPRSprob prob;
  XPRScreateprob(&prob);
  XPRSsetlogfile(prob, "logfile.log");
  XPRSdestroyprob(prob);
}
```

Note that a high level messaging framework is also available — which handles messages from all problem pointers created by the Optimizer library and messages related to initialization of the library itself — by calling the XPRS_ge_addcbmsghandler function. A convenient use of this callback, particularly when developing and debugging an application, is to trap all messages to file. The following line of code shows how to use the library function XPRSlogfilehandler together with XPRS_ge_addcbmsghandler to write all library message output to the file log.txt.

```
XPRS_ge_addcbmsghandler(XPRSlogfilehandler, "log.txt", 0);
```

Details about the use of callback functions can be found in section 5.8.

2.4 Problem Loading

Once a problem pointer has been created, an optimization problem can be loaded into it. The problem can be loaded either from file or from memory via the suite of problem loading and problem manipulation routines available in the Optimizer library interface. The simplest of these approaches, and the only approach available to Console Optimizer users, is to read a matrix from an MPS or LP file using `XPRSreadprob` (`READPROB`).

```
{
  XPRSprob prob;
  XPRScreateprob(&prob);
  XPRSsetlogfile(prob, "logfile.log");
  XPRSreadprob(prob, "hpw15", "");
  XPRSdestroyprob(prob);
}
```

Library users can construct the problem in their own arrays and then load this problem specification using one of the functions `XPRSloadlp`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmip`, `XPRSloadmiqp` or `XPRSloadmiqcqp`. During the problem load routine the Optimizer will use the user's data to construct the internal problem representation in new memory that is associated with the problem pointer. Note, therefore, that the user's arrays can be freed immediately after the call. Once the problem has been loaded, any subsequent call to one of these load routines will overwrite the problem currently represented in the problem pointer.

The names of the problem loading routines indicate the type of problem that can be represented using the routine. The following table outlines the components of an optimization problem as denoted by the codes used in the function names.

Code	Problem Content
lp	Linear Program (LP) (linear constraints and linear objective)
qp	Quadratic Program (LP with quadratic objective)
mi(p)	MIP entities (optimization problem with discrete variables and/or special ordered sets)
qc	Quadratic Constraints (LP with quadratic constraints)

Many of the array arguments of the load routines can optionally take NULL pointers if the associated component of the problem is not required to be defined. Note, therefore, that the user need only use the `XPRSloadmiqcqp` routine to load any problem that can be loaded by the other routines.

Finally, note that the names of the rows and columns of the problem are not loaded together with the problem specification. These may be loaded afterwards using a call to the function `XPRSaddnames`.

2.5 Problem Solving

With a problem loaded into a problem pointer the user can run the optimization algorithms on the problem to solve it.

The two main commands to run the optimization algorithms on a problem are `XPRSmipoptimize`(`MIPOPTIMIZE`) and `XPRSlpoptimize`(`LPOPTIMIZE`) depending on whether the problem needs to be solved with or without MIP entities. The `XPRSlpoptimize` function will solve LPs, QPs and QCQPs or the initial continuous relaxation of a MIP problem, depending on the type of problem loaded in the problem pointer. The `XPRSmipoptimize` function will solve MIPs, MIQPs and MIQCQPs.

For problems with MIP entities the Optimizer can be told to stop after having solved the initial relaxation by passing the 'l' flag to the XPRSmipoptimize function. The remaining MIP search can be run by calling the XPRSmipoptimize function without the 'l' flag.

```
{
  XPRSprob prob;
  XPRScreateprob(&prob);
  XPRSsetlogfile(prob, "logfile.log");
  XPRSreadprob(prob, "hpw15", "");
  XPRSmipoptimize(prob, "");
  XPRSdestroyprob(prob);
}
```

Non-convex quadratic problems must be solved with XPRSoptimize or XPRSnlpoptimize. See the Global Solver User Guide for more information.

2.6 Interrupting the Solve

It is common that users need to interrupt iterations before a solving algorithm is complete. This is particularly common when solving MIP problems since the time to solve these to completion can be large and users are often satisfied with near-optimal solutions. The Optimizer provides for this with structured interrupt criteria using controls and with user-triggered interrupts.

As described previously in section 1.2.4, the Console Optimizer can receive a user-triggered interrupt from the keyboard Ctrl-C event. It was also described in this previous section how interrupted commands could be resumed by simply reissuing the command. Similarly, optimization runs started from the library interface and interrupted by either structured or user-triggered interrupts, will return from the call in such a state that the run may be resumed with a follow-on call.

To setup structured interrupts the user will need to set the value of certain controls. Controls are scalar values that are accessed by their name in the Console Optimizer and by their id number via the library interface using functions such as XPRSgetintcontrol and XPRSsetintcontrol. These particular library functions are used for getting and setting the values of integer controls. Similar library functions are used for accessing double precision and string type controls.

Some types of structured interrupts include limits on iterations of the solving algorithms and a limit on the overall time of the optimization run. Limits on the simplex algorithms' iterations are set using the control LPITERLIMIT. Iterations of the Newton barrier algorithm are limited using the control BARITERLIMIT. A limit on the number of nodes processed in the branch and bound search when solving MIP problems is provided with the MAXNODE control. The integer control MAXTIME is used to limit the overall run time of the optimization run.

Note that it is important to be careful when using interrupts, to ensure that the optimization run is not being unduly restricted. This is particularly important when using interrupts on MIP optimization runs. Specific controls to use as stopping criteria for the MIP search are discussed in section 4.4.7.

```
{
  XPRSprob prob;
  XPRScreateprob(&prob);
  XPRSsetlogfile(prob, "logfile.log");
  XPRSreadprob(prob, "hpw15", "");
  XPRSsetintcontrol(prob, XPRS_MAXNODE, 20000);
  XPRSmipoptimize(prob, "");
  XPRSdestroyprob(prob);
}
```

Finally note that library users can trigger an interrupt on an optimization run (in a similar way to the Ctrl-C interrupt in the Console Optimizer) using a call to the function XPRSinterrupt. It is recommended that the user call this function from a callback during the optimization run. See section

5.8 for details about using callbacks.

2.7 Results Processing

Once the optimization algorithms have completed, either a solution will be available, or else the problem will have been identified as infeasible or unbounded. In the latter case, the user might want to know what caused this particular outcome and take steps to correct it. How to identify the causes of infeasibility and unboundedness are discussed in Chapter 6. In the former case, however, the user typically wants to retrieve the solution information into the required format.

The FICO Xpress Optimizer provides a number of functions for accessing solution information. An ASCII solution file can be obtained by `XPRSwriteslxsol` (`WRITESLXSOL`). The `.slx` format is similar format to the `.mps` format for MIP models and to the `.sol` format. Files in `.slx` format can be read back into the optimizer using the `XPRSreadslxsol` function. An extended solution file with additional information per column may be obtained as an ASCII file using either of `XPRSwritesol` (`WRITESOL`) or `XPRSwriteprtsol`(`WRITEPRTSOL`).

Library interface users may additionally access the current solution information via memory using `XPRSgetsolution`. The values of the slack values, dual values and reduced costs can be obtained by calling `XPRSgetslacks`, `XPRSgetduals` and `XPRSgetredcosts`.

In addition to the arrays of solution information provided by the Optimizer, summary solution information is also available through *problem attributes*. These are named scalar values that can be accessed by their id number using the library functions `XPRSgetintattrib`, `XPRSgetdblattrib` and `XPRSgetstrattrib`. Examples of attributes include `LPOBJVAL` and `MIPOBJVAL`, which return the objective function values for the current LP, QP or QCQP solution and the last MIP solution, respectively. A full list of attributes may be found in Chapter 10.

When the optimization routine returns it is recommended that the user check the status of the run to ensure the results are interpreted correctly. For continuous optimization runs (started with `XPRSlpoptimize`) the status is available using the `LPSTATUS` integer problem attribute. For MIP optimization runs (started with `XPRSmipoptimize`) the status is available using the `MIPSTATUS` integer problem attribute. See the attribute's reference section for the definition of their values.

```
{
  XPRSprob prob;
  int nCols;
  double *x;
  XPRScreateprob(&prob);
  XPRSsetlogfile(prob, "logfile.log");
  XPRSreadprob(prob, "hpw15", "");
  XPRSgetintattrib(prob, XPRS_COLS, &nCols);
  XPRSsetintcontrol(prob, XPRS_MAXNODE, 20000);
  XPRSmipoptimize(prob, "");
  XPRSgetintattrib(prob, XPRS_MIPSTATUS, &iStatus);
  if(iStatus == XPRS_MIP_SOLUTION || iStatus == XPRS_MIP_OPTIMAL) {
    x = (double *) malloc(sizeof(double) * nCols);
    XPRSgetsolution(prob, NULL, x, 0, nCols - 1);
  }
  XPRSdestroyprob(prob);
}
```

Note that, unlike for LP, QP or QCQP solutions, dual solution information is *not* available by calling `XPRSgetduals` or `XPRSgetredcosts` and is not automatically generated with the MIP solutions. The reason for this is that MIP problems do not satisfy the theoretical conditions by which dual information is derived (i.e., Karush-Kuhn-Tucker conditions). In particular, this is because the MIP constraint functions are, in general, not continuously differentiable (indeed, the domains of integer variables are not continuous).

Despite this, some useful dual information can be generated if a MIP has continuous variables and we solve the resulting LP problem generated by fixing the non-continuous component of the problem to their solution values. Because this process can be expensive it is left to the user to perform this in a post solving phase where the user will simply call the function `XPRSfixmipentities` followed with a call to the optimization routine `XPRSlpoptimize`.

2.8 Function Quick Reference

2.8.1 Administration

`XPRSinit`	Initialize the Optimizer.
`XPRScreateprob`	Create a problem pointer.
`XPRSsetlogfile`	Direct all Optimizer output to a log file.
`XPRSaddcbmessage`	Define a message handler callback function.
`XPRSgetintcontrol`	Get the value of an integer control,
`XPRSsetintcontrol`	Set the value of an integer control.
`XPRSinterrupt`	Set the interrupt status of an optimization run.
`XPRSdestroyprob`	Destroy a problem pointer.
`XPRSfree`	Release resources used by the Optimizer.

2.8.2 Problem Loading

`XPRSreadprob`	Read an MPS or LP format file.
`XPRSloadlp`	Load an LP problem.
`XPRSloadqp`	Load a quadratic objective problem.
`XPRSloadqcqp`	Load a quadratically constrained, quadratic objective problem.
`XPRSloadmip`	Load a MIP problem.
`XPRSloadmiqp`	Load a quadratic objective MIP problem.
`XPRSloadmiqcqp`	Load a quadratically constrained, quadratic objective MIP problem.
`XPRSaddnames`	Load names for a range of rows or columns in a problem.

2.8.3 Problem Solving

`XPRSreadbasis`	Read a basis from file.
`XPRSloadbasis`	Load a basis from user arrays.
`XPRSreaddirs`	Read a directives file.
`XPRSlpoptimize`	Solve the problem without MIP entities.
`XPRSmipoptimize`	Run the problem with MIP entities.
`XPRSfixmipentities`	Fix the discrete variables in the problem to the values of the current MIP solution stored with the problem pointer.
`XPRSgetbasis`	Copy the current basis into user arrays.
`XPRSwritebasis`	Write the current basis to file.

2.8.4 Results Processing

XPRSwritesol	Write the current solution to ASCII files.
XPRSwriteprtsol	Write the current solution in printable format to file.
XPRSgetsolution	Copy the current solution values into a user array.
XPRSgetslacks	Copy the current slack values into a user array.
XPRSgetduals	Copy the current dual values into a user array.
XPRSgetredcosts	Copy the current reduced costs into a user array.
XPRSgetintattrib	Get the value of an integer problem attribute e.g., by passing the id MIPSOLS the user can get the number of MIP solutions found.
XPRSgetdblattrib	Get the value of a double problem attribute e.g., by passing the id MIPOBJVAL the user can get the objective value of the last MIP solution.
XPRSgetstrattrib	Get the value of a string problem attribute.

2.9 Summary

In the previous sections a brief introduction is provided to the most common features of the FICO Xpress Optimizer and its most general usage. The reader should now be familiar with the main routines in the Optimizer library. These routines allow a user to create problem pointers and load problems into these problem pointers. The reader should be familiar with the requirements for setting up message handling with the Optimizer library. Also the reader should know how to run the optimization algorithms on the loaded problems and be familiar with the various ways that results can be accessed.

Examples of using the Optimizer are available from a number of sources, most notably from FICO Xpress Getting Started manual. This provides a straight forward, "hands-on" approach to the FICO Xpress Optimization Suite and it is highly recommended that users read the relevant chapters before considering the reference manuals. Additionally, more advanced, examples may be downloaded from the website.

CHAPTER 3

Problem Types

The FICO Xpress Optimization Suite is a powerful optimization tool for solving Mathematical Programming problems. Users of FICO Xpress formulate real-world problems as Mathematical Programming problems by defining a set of decision variables, a set of constraints on these variables and an objective function of the variables that should be maximized or minimized. Our FICO Xpress users have applications that define and solve important Mathematical Programming problems in academia and industry, including areas such as production scheduling, transportation, supply chain management, telecommunications, finance and personnel planning.

Mathematical Programming problems are usually classified according to the types of decision variables, constraints and objective function in the problem. Perhaps the most popular application of the FICO Xpress Optimizer is for the class of Mixed Integer Programs (MIPs). In this section we will briefly introduce some important types of problem that can be solved by the FICO Xpress Optimizer.

This section does not discuss general nonlinear problems. To solve such problems see FICO Xpress Nonlinear and FICO Xpress Global. Note that solving non-convex quadratically-constrained quadratic problems is included in the Xpress Optimizer license, whereas for general nonlinear problems, you need a license for either FICO Xpress Global or FICO Xpress Nonlinear.

3.1 Linear Programs (LPs)

Linear Programming (LP) problems are a very common type of optimization problems. In this type of problem all constraints and the objective function are linear expressions of the decision variables. Each decision variable is restricted to some continuous interval (typically non-negative). Although the methods for solving these types of problems are well known (e.g., the simplex method), only a few efficient implementations of these methods (and additional specialized methods for particular classes of LPs) exists, and these are often crucial for solving the increasingly large instances of LPs arising in industry.

3.2 Mixed Integer Programs (MIPs)

Many problems can be modeled satisfactorily as Linear Programs (LPs), i.e., with variables that are only restricted to having values in continuous intervals. However, a common class of problems requires modeling using discrete variables. These problems are called Mixed Integer Programs (MIPs). MIP problems are often hard to solve and may require large amounts of computation time to obtain even satisfactory, if not optimal, results.

Perhaps the most common use of the FICO Xpress Optimization Suite is for solving MIP problems and it is designed to handle the most challenging of these problems. Besides providing solution support for MIP problems the Optimizer provides support for a variety of popular MIP modeling constructs:

Binary variables – decision variables that have value either `0` or `1`, sometimes called 0/1 variables;

Integer variables – decision variables that have integer values;

Semi-continuous variables – decision variables that either have value `0`, or a continuous value above a specified non-negative limit. Semi-continuous variables are useful for modeling cases where, for example, if a quantity is to be supplied at all then it will be supplied starting from some minimum level (e.g., a power generation unit);

Semi-continuous integer variables – decision variables that either have value `0`, or an integer value above a specified non-negative limit;

Partial integer variables – decision variables that have integer values below a specified limit and continuous values above the limit. Partial integer variables are useful for modeling cases where a supply of some quantity needs to be modeled as discrete for small values but we are indifferent whether it is discrete when the values are large (e.g., because, say, we do not need to distinguish between 10000 items and 10000.25 items);

Special ordered sets of type one (SOS1) — a set of decision variables ordered by a set of specified continuous values (or reference values) of which at most one can take a nonzero value. SOS1s are useful for modeling quantities that are taken from a specified discrete set of continuous values (e.g., choosing one of a set of transportation capacities);

Special ordered sets of type two (SOS2) – a set of variables ordered by a set of specified continuous values (or reference values) of which at most two can be nonzero, and if two are nonzero then they must be consecutive in their ordering. SOS2s are useful for modeling a piecewise linear quantity (e.g., unit cost as a function of volume supplied);

Indicator constraints – constraints each with a specified associated binary 'controlling' variable where we assume the constraint must be satisfied when the binary variable is at a specified binary value; otherwise the constraint does not need to be satisfied. Indicator constraints are useful for modeling cases where supplying some quantity implies that a fixed cost is incurred; otherwise if no quantity is supplied then there is no fixed cost (e.g., starting up a production facility to supply various types of goods and the total volume of goods supplied is bounded above).

Piecewise linear constraints – constraints that define a piecewise linear relationship between two variables. These are defined via a set of breakpoints with linearly interpolated values between and beyond them (with the slope before the first and after the last point continuing the slope between the first/last two points). The piecewise linear functions are allowed to be discontinuous by defining multiple points with the same value of the input variable `x`, in which case the output variable `y` is allowed to take any value between the corresponding `y`-values of these breakpoints, while the first of them will define the slope before and the last will define the slope after this `x`-value. Piecewise linear constraints are useful to model e.g. transportation costs that are constant/linear in specific intervals but may jump between the different brackets.

General constraints – specific type of MIP constraints to model `min`, `max`, `and`, `or`, and `absolute value` relationships between two or more variables.

All of the above MIP variable types are collectively referred to as **MIP entities**.

3.3 Quadratic Programs (QPs)

Quadratic Programming (QP) problems are an extension of Linear Programming (LP) problems where the objective function may include a second order polynomial. An example of this is where the user wants to minimize the statistical variance (a quadratic function) of the solution values.

The FICO Xpress Optimizer can be used directly for solving QP problems with support for quadratic objectives in the MPS and LP file formats and library routines for loading QPs and manipulating quadratic objective functions.

3.4 Quadratically Constrained Quadratic Programs (QC-QPs)

Quadratically Constrained Quadratic Programs (QCQPs) are an extension of the Quadratic Programming (QP) problem where the constraints may also include second order polynomials.

A QCQP problem may be written as:

$$
\begin{array}{lccccccccc}
\text{minimize} & c_1x_1 & + & \dots & + & c_nx_n & + & x^TQ_0x & & \\
\text{subject to} & a_{11}x_1 & + & \dots & + & a_{1n}x_n & + & x^TQ_1x & \leqslant & b_1 \\
 & & & \vdots & & & & & & \\
 & a_{m1}x_1 & + & \dots & + & a_{mn}x_n & + & x^TQ_mx & \leqslant & b_m \\
 & l_1 \leqslant x_1 \leqslant u_1, & & \dots & & l_n \leqslant x_n \leqslant u_n & & & &
\end{array}
$$

where any of the lower or upper bounds l_i or u_i may be infinite.

The FICO Xpress Optimizer can be used directly for solving QCQP problems with support for quadratic constraints and quadratic objectives in the MPS and LP file formats and library routines for loading QCQPs and manipulating quadratic objective functions and the quadratic component of constraints.

Properties of QCQP problems are discussed in the following few sections.

3.4.1 Algebraic and matrix form

Each second order polynomial can be expressed as x^TQx where Q is an appropriate symmetric matrix: the quadratic expressions are generally either given in the algebraic form

$$a_{11}x_1^2 + 2a_{12}x_1x_2 + 2a_{13}x_1x_3 + \dots + a_{22}x_2^2 + 2a_{23}x_2x_3 + \dots$$

like in LP files, or in the matrix form x^TQx where

$$Q = \begin{pmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & & \\ \vdots & & \ddots & \\ a_{n1} & & & a_{nn} \end{pmatrix}$$

like in MPS files. As symmetry is always assumed, $a_{ij} = a_{ji}$ for all index pairs (i, j).

3.4.2 Convexity

A fundamental property for nonlinear optimization problems, thus in QCQP as well, is convexity. A region is called *convex* if for any two points from the region the connecting line segment is also part of the region.

The lack of convexity may give rise to several unfavorable model properties. Lack of convexity in the objective may introduce the phenomenon of locally optimal solutions that are not global ones (a local optimal solution is one for which a neighborhood in the feasible region exists in which that solution is the best). While the lack of convexity in constraints can also give rise to local optima, they may even introduce non-connected feasible regions as shown in Figure 3.1.

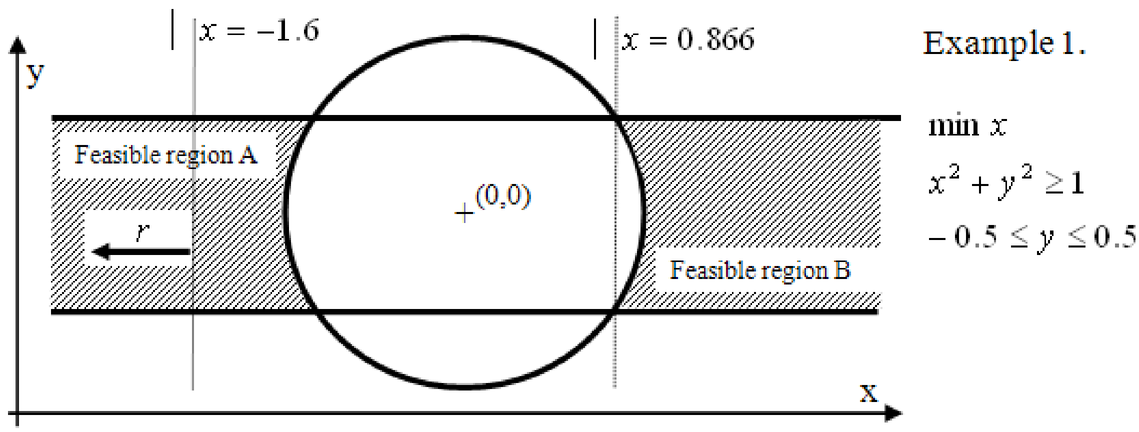

Figure 3.1: Non-connected feasible regions

In this example, the feasible region is divided into two parts. In region B, the objective function has two alternative locally optimal solutions, while in region A the objective function is not even bounded.

For convex problems, each locally optimal solution is a global one, making the characterization of the optimal solution efficient.

3.4.3 Characterizing Convexity in Quadratic Constraints

A quadratic constraint of the form

$$a_1x_1 + ... + a_nx_n + x^TQx \leqslant b$$

defines a convex region if and only if Q is a so-called *positive semi-definite* (PSD) matrix.

A square matrix Q is PSD by definition if for any vector (not restricted to the feasible set of a problem) x it holds that $x^TQx \geqslant 0$.

It follows that for greater-than or equal constraints

$$a_1x_1 + ... + a_nx_n + x^TQx \geqslant b$$

the negative of Q must be PSD.

A nontrivial quadratic equality constraint (one for which not every coefficient is zero) always defines a non-convex region (or in other words, if both Q and its negative is PSD, then Q must be equal to the 0 matrix). Therefore, quadratic equality constraints require the functionality of FICO Xpress Global, which can solve non-convex optimization problems.

Determining whether a matrix is PSD is not always obvious nor trivial. There are certain constructs, however, that can easily be recognized as being non convex:

1. the product of two variables, say xy, without having both x^2 and y^2 present;
2. having $-x^2$ in any quadratic expression in a less than or equal constraint, or having x^2 in any greater than or equal constraint.

3.5 Second Order Cone problems (SOCPs)

Second order cone problems (SOCP) are a special class of quadratically constrained problems, where the quadratic matrix Q is not required to be semi-definite.

The FICO Xpress Optimizer supports (mixed integer) second order cone problems that satisfy the following requirements.

Each quadratic constraint satisfies one of the following two forms:

1. **Second order (or Lorentz) cone**: $x_1^2 + x_2^2 + \cdots + x_k^2 - t^2 \leqslant 0$ where $t \geqslant 0$
2. **Rotated second order (or Lorentz) cone**: $x_1^2 + x_2^2 + \cdots + x_k^2 - 2t_1t_2 \leqslant 0$ where $t_1, t_2 \geqslant 0$

All of the cone coefficients must be exactly one, except for the coefficient of 2 for the t_1t_2-product. Constants or linear terms are not allowed.

For barrier solves, the cones should not overlap, i.e., any variable x may only appear in a single cone. Otherwise a transformation will be be applied that allows the problem to be solved with barrier, but it might not be possible to obtain the optimal dual values in this case.

Second order cone problems are loaded using the same API functions as for quadratic constraints, and the conic constraints are auto-detected by the optimizer at run time.

CHAPTER 4

Solution Methods

The FICO Xpress Optimization Suite provides four fundamental optimization algorithms for LP or QP problems: the *primal simplex*, the *dual simplex*, the *hybrid gradient* and the *Newton barrier* algorithm (QCQP and SOCP problems are always solved with the *Newton barrier* algorithm). Using these algorithms the Optimizer implements solving functionality for the various types of continuous problems the user may want to solve.

Typically the user will allow the Optimizer to choose what combination of methods to use for solving their problem. For example, by default, the FICO Xpress Optimizer uses the dual simplex method for solving LP problems and the barrier method for solving QP problems. For the initial continuous relaxation of a MIP, the defaults will be different and depends both on the number of solver threads used, the type of the problem and the MIP technique selected.

For most users the default behavior of the Optimizer will provide satisfactory solution performance and they need not consider any customization. However, if a problem seems to be taking an unusually long time to solve or if the solving performance is critical for the application, the user may consider, as a first attempt, to force the Optimizer to use an algorithm other than the default.

The main points where the user has a choice of what algorithm to use are (i) when the user calls the optimization routine `XPRSlpoptimize` (`LPOPTIMIZE`) and (ii) when the Optimizer solves the node relaxation problems during the branch and bound search. The user may force the use of a particular algorithm by specifying flags to the optimization routine `XPRSlpoptimize` (`LPOPTIMIZE`). If the user specifies flags to `XPRSmipoptimize` (`MIPOPTIMIZE`) to select a particular algorithm then this algorithm will be used for the initial relaxation only. To specify what algorithm to use when solving the node relaxation problems during branch and bound use the special control parameter, `DEFAULTALG`.

As a guide for choosing optimization algorithms other than the default consider the following. As a general rule, the dual simplex is usually much faster than the primal simplex if the problem is neither infeasible nor near-infeasible. If the problem is likely to be infeasible or if the user wishes to get diagnostic information about an infeasible problem then the primal simplex is the best choice. This is because the primal simplex algorithm finds a basic solution that minimizes the sum of infeasibilities and these solutions are typically helpful in identifying causes of infeasibility. The Newton barrier algorithm can perform much better than the simplex algorithms on certain classes of problems. The barrier algorithm will, however, likely be slower than the simplex algorithms if, for a problem coefficient matrix A, $A^{\mathsf{T}}A$ is large and dense.

The hybrid gradient algorithm can be used only for LP problems. It is especially useful on very degenerate problems, and, thanks to its low memory usage, on extremely large problems.

In the following few sections, performance issues relating to these methods will be discussed in more detail. Performance issues relating to the search for MIP solutions will also be discussed.

4.1 Simplex Method

The simplex method was the first efficient method devised for solving Linear Programs (LPs). This method is still commonly used today and there are efficient implementations of the primal and dual simplex methods available in the Optimizer. We briefly outline some basic simplex theory to give the user a general idea of the simplex algorithm's behavior and to define some terminology that is used in the reference sections.

A region defined by a set of constraints is known in Mathematical Programming as a *feasible region*. When these constraints are linear the feasible region defines the solution space of a Linear Programming (LP) problem. Each value of the objective function of an LP defines a hyperplane or a *level set*. A fundamental result of simplex algorithm theory is that an optimal value of the LP objective function will occur when the level set grazes the boundary of the feasible region. The optimal level set either intersects a single point (or *vertex*) of the feasible region (if such a point exists), in which case the solution is unique, or it intersects a boundary set of the feasible region in which case there is an infinite set of solutions.

In general, vertices occur at points where as many constraints and variable bounds as there are variables in the problem intersect. Simplex methods usually only consider solutions at vertices, or *bases* (known as *basic solutions*) and proceed or iterate from one vertex to another until an optimal solution has been found, or the problem proves to be infeasible or unbounded. The number of iterations required increases with model size, and typically goes up slightly faster than the increase in the number of constraints.

The primal and dual simplex methods differ in which vertices they consider and how they iterate. The dual is the default for LP problems, but may be explicitly invoked using the `d` flag with `XPRSlpoptimize` (`LPOPTIMIZE`).

4.1.1 Output

While the simplex methods iterate, the Optimizer produces iteration logs. The Console Optimizer writes these logging messages to the screen. Library users can setup logging management using the various relevant functions in the Optimizer library e.g., `XPRSsetlogfile`, `XPRSaddcbmessage` or `XPRSaddcblplog`. The simplex iteration log is produced at regular time intervals, determined by an internal deterministic. When `LPLOG` is set to `0`, a log is displayed only when the optimization run terminates. If it is set to a positive value, a summary style log is output; otherwise, a detailed log is output.

4.2 Newton Barrier Method

In contrast to the simplex methods that iterate through boundary points (vertices) of the feasible region, the Newton barrier method iterates through solutions in the interior of the feasible region and will typically find a close approximation of an optimal solution. Consequently, the number of barrier iterations required to complete the method on a problem is determined more so by the required proximity to the optimal solution than the number of decision variables in the problem. Unlike the simplex method, therefore, the barrier often completes in a similar number of iterations regardless of the problem size.

The barrier solver can be invoked on a problem by using the '`b`' flag with `XPRSlpoptimize` (`LPOPTIMIZE`). This is used by default for QP problems, whose quadratic objective functions in general result in optimal solutions that lie on a face of the feasible region, rather than at a vertex.

4.2.1 Crossover

Typically the barrier algorithm terminates when it is within a given tolerance of an optimal solution. Since this solution will not lie exactly on the boundary of the feasible region, the Optimizer can be optionally made to perform a, so-called, 'crossover' phase to obtain an optimal solution on the boundary. The nature of the 'crossover' phase results in a basic optimal solution, which is at a vertex of the feasible region. In the crossover phase the simplex method is used to continue the optimization from the solution found by the barrier algorithm. The `CROSSOVER` control determines whether the Optimizer performs crossover. When set to `1` (the default for LP problems), crossover is performed. If `CROSSOVER` is set to `0`, no crossover will be attempted and the solution provided will be that determined purely by the barrier method. Note that if a basic optimal solution is required, then the `CROSSOVER` option must be activated before optimization starts.

4.2.2 Output

While the barrier method iterates, the Optimizer produces iteration log messages. The Console Optimizer writes these log messages to the screen. Library users can setup logging management using the various relevant functions in the Optimizer library, e.g. `XPRSsetlogfile`, `XPRSaddcbmessage` or `XPRSaddcbbarlog`. Note that the amount of barrier iteration logging is dependent on the value of the `BAROUTPUT` control.

4.3 Hybrid Gradient Method

Similar to the Newton barrier algorithm, the hybrid gradient method iterates through solutions in the interior of the feasible region and will typically find a close approximation of an optimal solution. The number of hybrid gradient iterations required to complete the method is typically large, but each iteration is very fast.

The hybrid gradient solver can be invoked on a problem by using the 'b' flag with `XPRSlpoptimize` (`LPOPTIMIZE`) and setting `BARALG` to 4.

The hybrid gradient method is the implementation of a first-order primal-dual alternating directions method. As such, its rate of convergence is only linear, thus it is not designed to achieve the accuracy of some of the other algorithms, most notably that of the Newton barrier methods. For this reason it is advised to relax the feasibility (BARPRIMALSTOP) and optimality tolerances (BARDUALSTOP). By default they are 100 times the defaults we would have in the Newton barrier algorithm.

4.3.1 Crossover

The hybrid gradient algorithm terminates when the solution it has found is within a given tolerance of an optimal solution. This solution is typically not on the boundary of the feasible set. A crossover phase, similar to the one applied after the Newton barrier method, can be used to arrive at a basic solution.

4.3.2 Output

The format and content of the hybrid gradient method log is identical to the Newton barrier log with a few notable differences: since no Cholesky decomposition is performed the table about matrix density is not printed. Also, upper bound infeasibility is always 0.

4.4 Branch and Bound

The FICO Xpress Optimizer uses the approach of LP based branch and bound with cutting planes for solving Mixed Integer Programming (MIP) problems. That is, the Optimizer solves the optimization problem (typically an LP problem) resulting from relaxing the discreteness constraints on the variables and then uses branch and bound to search the relaxation space for MIP solutions. It combines this with heuristic methods to quickly find good solutions, and cutting planes to strengthen the LP relaxations.

The Optimizer's MIP solving methods are coordinated internally by sophisticated algorithms so the Optimizer will work well on a wide range of MIP problems with a wide range of solution performance requirements without any user intervention in the solving process. Despite this the user should note that the formulation of a MIP problem is typically not unique and the solving performance can be highly dependent on the formulation of the problem. It is recommended, therefore, that the user undertake careful experimentation with the problem formulation using realistic examples before committing the formulation for use on large production problems. It is also recommended that users have small scale examples available to use during development.

Because of the inherent difficulty in solving MIP problems and the variety of requirements users have on the solution performance on these problems it is not uncommon that users would like to improve over the default performance of the Optimizer. In the following sections we discuss aspects of the branch and bound method for which the user may want to investigate when customizing the Optimizer's MIP search.

4.4.1 Theory

In this section we present a brief overview of branch and bound theory as a guide for the user on where to look to begin customizing the Optimizer's MIP search and also to define the terminology used when describing branch and bound methods.

To simplify the text in the following, we limit the discussion to MIP problems with linear constraints and a linear objective function. Note that it is not difficult to generalize the discussion to problems with quadratic constraints and a quadratic objective function.

The branch and bound method has three main concepts: relaxation, branching and fathoming.

The relaxation concept relates to the way discreteness or integrality constraints are dropped or 'relaxed' in the problem. The initial relaxation problem is a Linear Programming (LP) problem which we solve resulting in one of the following cases:

(a) The LP is infeasible so the MIP problem must also be infeasible;

(b) The LP has a feasible solution, but some of the integrality constraints are not satisfied – the MIP has not yet been solved;

(c) The LP has a feasible solution and all the integrality constraints are satisfied so the MIP has also been solved;

(d) The LP is unbounded.

Case (d) is a special case. It can only occur when solving the initial relaxation problem and in this situation the MIP problem itself is not well posed (see Chapter 6 for details about what to do in this case). For the remaining discussion we assume that the LP is not unbounded.

Outcomes (a) and (c) are said to 'fathom' the particular MIP, since no further work on it is necessary. For case (b) more work is required, since one of the unsatisfied integrality constraints must be selected and the concept of separation applied.

To illustrate the branching concept suppose, for example, that the optimal LP value of an integer variable x is 1.34, a value which violates the integrality constraint. It follows that in any solution to the original problem either $x \leqslant 1.0$ or $x \geqslant 2.0$. If the two resulting MIP problems are solved (with the integrality constraints), all integer values of x are considered in the combined solution spaces of the two MIP problems and no solution to one of the MIP problems is a solution to the other. In this way we have branched the problem into two disjoint *sub-problems*.

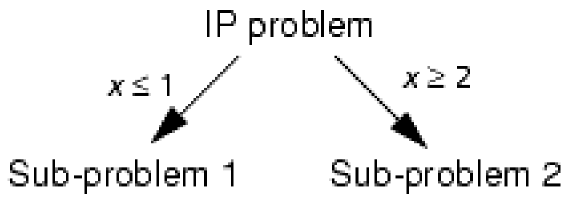

If both of these sub-problems can be solved and the better of the two is chosen, then the MIP is solved. By recursively applying this same strategy to solve each of the sub-problems and given that in the limiting case the integer variables will have their domains divided into fixed integer values then we can guarantee that we solve the MIP problem.

Branch and bound can be viewed as a *tree search* algorithm. Each *node* of the tree is a MIP problem. A MIP node is relaxed and the LP relaxation is solved. If the LP relaxation is not fathomed, then the node MIP problem is partitioned into two more sub-problems, or *child* nodes. Each child MIP will have the same constraints as the *parent* node MIP, plus one additional inequality constraint. Each node is therefore either fathomed or has two children or *descendants*.

We now introduce the concept of a *cutoff*, which is an extension of the fathoming concept. To understand the cutoff concept we first make two observations about the behavior of the node MIP problems. Firstly, the optimal MIP objective of a node problem can be no better than the optimal objective of the LP relaxation. Secondly, the optimal objective of a child LP relaxation can be no better than the optimal objective of its parent LP relaxation. Now assume that we are exploring the tree and we are keeping the value of the best MIP objective found so far. Assume also that we keep a 'cutoff value' equal to the best MIP objective found so far. To use the cutoff value we reason that if the optimal LP relaxation objective is no better than the cutoff then any MIP solution of a descendant can be no better than the cutoff and the node can be fathomed (or cutoff) and need not be considered further in the search.

The concept of a cutoff can be extended to apply even when no integer solution has been found in situations where it is known, or may be assumed, from the outset that the optimal solution must be better than some value. If the relaxation is worse than this cutoff, then the node may be fathomed. In this way the user can reduce the number of nodes processed and improve the solution performance. Note that there is the possibility, however, that all MIP solutions, including the optimal one, may be missed if an overly optimistic cutoff value is chosen.

The cutoff concept may also be extended in a different way if the user intends only to find a solution within a certain tolerance of the overall optimal MIP solution. Assume that we have found a MIP solution to our problem and assume that the cutoff is maintained at a value 100 objective units better than the current best MIP solution. Proceeding in this way we are guaranteed to find a MIP solution within 100 units of the overall MIP optimal since we only cutoff nodes with LP relaxation solutions worse than 100 units better than the best MIP solution that we find.

If the MIP problem contains SOS entities then the nodes of the branch and bound tree are determined by branching on the sets. Note that each member of the set has a double precision reference row entry and the sets are ordered by these reference row entries. Branching on the sets is done by choosing a position in the ordering of the set variables and setting all members of the set to 0 either above or below the chosen point. The optimizer used the reference row entries to decide on the branching position and so it is important to choose the reference row entries which reflect the cost of setting the set member to 0. In some cases it maybe better to model the problem with binary variables instead of special ordered sets. This is especially the case if the sets are small.

4.4.2 Variable Selection and Cutting

The branch and bound technique leaves many choices open to the user. In practice, the success of the technique is highly dependent on several key choices.

(a) Deciding which variable to branch on is known as the *variable selection problem* and is often the most critical choice.

(b) *Cutting planes* are used to strengthen the LP relaxation of a subproblem, and can often bring a significant reduction in the number of sub-problems that must be solved

The Optimizer incorporates a default strategy for both choices which has been found to work adequately on most problems. Several controls are provided to tailor the search strategy to a particular problem.

4.4.3 Variable Selection for Branching

Each MIP entity has a priority for branching, or one set by the user in the directives file. A *low* priority value means that the variable is *more* likely to be selected for branching. The Optimizer uses a priority range of 400–500 by default. To guarantee that a particular MIP entity is always branched first, the user should assign a priority value less than 400. Likewise, to guarantee that a MIP entity is only branched on when it is the only candidate left, a priority value above 500 should be used.

The Optimizer uses a wide variety of information to select among those entities that remain unsatisifed and which belong to the lowest valued priority class. A *pseudo cost* is calculated for each candidate entity, which is typically an estimate of how much the LP relaxation objective value will change (degradationas a result of branching on this particular candidate. Estimates are calculated separately for the up and down branches and combined according to the strategy selected by the `VARSELECTION` control.

The default strategy is based on calculating pseudo costs using the method of *strong branching*. With strong brancing, the LP relaxations of the two potential subproblems that would result from branching on a candidate MIP entity, are solved partially. Dual simplex is applied for a limited number of iterations and the change in objective value is recorded as a pseudo cost. This can be very expensive to apply to every candidate for every node of the branch and bound search, which is why the Optimizer by default will reuse pseudo costs collected from one node, on subsequent nodes of the search.

Selecting a MIP entity for branching is a multi-stage process, which combines estimates that are cheap to compute, with the more expensive strong branching based pseudo costs. The basic selection process is given by the following outline, together with the controls that affect each step:

1. Pre-filter the set of candidate entities using very cheap estimates.
   `SBSELECT`: determine the filter size.
2. Calculate simple estimates based on local node information and rank the selected candidates.
   `SBESTIMATE`: local ranking function.
3. Calculate strong branching pseudo costs for candidates lacking such information.
   `SBBEST`: number of variables to strong branch on.
   `SBITERLIMIT`: LP iteration limit for strong branching.
4. Select the best candidate using a combination of pseudo costs and the local ranking functions.

The overall amount of effort put into this process can be adjusted using the `SBEFFORT` control.

4.4.4 Cutting Planes

Cutting planes are valid constraints used for tightening the LP relaxation of a MIP problem, without affecting the MIP solution space. They can be very effective at reducing the amount of subproblems that the branch and bound search has to solve. The Optimizer will automatically create many different well-known classes of cutting planes, such as *mixed integer Gomory cuts*, *lift-and-project cuts*, *mixed integer rounding* (MIR) *cuts*, *clique cuts*, *implied bound cuts*, *flow-path cuts*, *zero-half cuts*, etc. These classes of cuts are grouped together into two groups that can be controlled separately. The following table lists the main controls and the related cut classes that are affected by those control:

COVERCUTS	Mixed integer rounding cuts
TREECOVERCUTS	Lifted cover cuts
	Clique cuts
	Implied bound cuts
	Flow-path cuts
	Zero-half cuts
GOMCUTS	Mixed integer Gomory cuts
TREEGOMCUTS	Lift-and-project cuts

The controls COVERCUTS and GOMCUTS sets an upper limit on the number of rounds of cuts to create for the root problem, for their respective groups. Correspondingly, TREECOVERCUTS and TREEGOMCUTS sets an upper limit on the number of rounds of cuts for any subproblem in the tree.

An important aspect of cutting is the choice of how many cuts to add to a subproblem. The more cuts are added, the harder it becomes to solve the LP relaxation of the node problem. The tradeoff is therefore between the additional effort in solving the LP relaxation versus the strengthening of the subproblem. The CUTSTRATEGY control sets the general level of how many cuts to add, expressed as a value from 0 (no cutting at all) to 3 (high level of cuts).

Another important aspect of cutting is how often cuts should be created and added to a subproblem. The Optimizer will automatically decide on a frequency that attempts to balance the effort of creating cuts versus the benefits they provide. It is possible to override this and set a fixed strategy using the CUTFREQ control. When set to a value k, cutting will be applied to every k'th level of the branch and bound tree. Note that setting CUTFREQ = 0 will disable cutting on subproblems completely, leaving only cutting on the root problem. This might be advantageous for some problems and the Optimizer uses an ML-based strategy to detect such cases automatically. This feature can be (de)activated by the AUTOCUTTING control.

4.4.5 Node Selection

The Optimizer applies a search scheme involving best-bound first search combined with dives. Subproblems that have not been fathomed or which have not been branched further into new subproblems are referred to as *active nodes* of the branch and bound search tree. Such activate nodes are maintained by the Optimizer in a pool.

The search process involves selecting a subproblem (or node) from this active nodes pool and commencing a dive. When the Optimizer branches on a MIP entity and creates the two subproblems, it has a choice of which of the two subproblems to work on next. This choice is determined by the BRANCHCHOICE control. The dive is a recursive search, where it selects a child problem, branches on it to create two new child problems, and repeats with one of the new child problems, until it ends with a subproblem that should not be branched further. At this point it will go back to the active nodes pool and pick a new subproblem to perform a dive on. This is called a *backtrack* and the choice of node is determined by the BACKTRACK control. The default backtrack strategy will select the active node with the best bound.

4.4.6 Adjusting the Cutoff Value

The parameter MIPADDCUTOFF determines the cutoff value set by the Optimizer when it has identified a new MIP solution. The new cutoff value is set as the objective function value of the MIP solution plus the value of MIPADDCUTOFF. If MIPADDCUTOFF has not been set by the user, the value used by the Optimizer will be calculated after the initial LP optimization step as:

$$\max(\texttt{MIPADDCUTOFF}, 0.01 \cdot \texttt{MIPRELCUTOFF} \cdot \mathit{LP\_value})$$

using the initial values for MIPADDCUTOFF and MIPRELCUTOFF, and where *LP_value* is the optimal objective value of the initial LP relaxation.

4.4.7 Stopping Criteria

Often when solving a MIP problem it is sufficient to stop with a good solution instead of waiting for a potentially long solve process to find an optimal solution. The Optimizer provides several stopping criteria related to the solutions found, through the MIPRELSTOP and MIPABSSTOP parameters. If MIPABSSTOP is set for a minimization problem, the Optimizer will stop when it finds a MIP solution with an objective value equal to or less than MIPABSSTOP. The MIPRELSTOP parameter can be used to stop the solve process when the found solution is sufficiently close to optimality, as measure relative to the best available bound. The optimizer will stop due to MIPRELSTOP when the following is satisfied:

$$|\texttt{MIPOBJVAL} - \texttt{BESTBOUND}| \leqslant \texttt{MIPRELSTOP} \cdot \max(|\texttt{BESTBOUND}|, |\texttt{MIPOBJVAL}|)$$

It is also possible to set limits on the solve process, such as number of nodes (MAXNODE), time limit (MAXTIME) or on the number of solutions found (MAXMIPSOL). If the solve process is interrupted due to any of these limits, the problem will be left in its unfinished state. It is possible to resume the solve from an unfinished state by calling XPRSmipoptimize (MIPOPTIMIZE) again.

To return an unfinished problem to its starting state, where it can be modified again, the user should use the function XPRSpostsolve (POSTSOLVE). This function can be used to restore a problem from an interrupted tree search even if the problem is not in a presolved state.

4.4.8 Integer Preprocessing

If MIPPRESOLVE has been set to a nonzero value before solving a MIP problem, integer preprocessing will be performed at each node of the branch and bound tree search (including the root node). This incorporates reduced cost tightening of bounds and tightening of implied variable bounds after branching. If a variable is fixed at a node, it remains fixed at all its child nodes, but it is not deleted from the matrix (unlike the variables fixed by presolve).

MIPPRESOLVE is a bitmap whose values are acted on as follows:

Bit	Value	Action
0	1	Reduced cost fixing;
1	2	Integer implication tightening.
2	4	*Unused*
3	8	Tightening of implied continuous variables.
4	16	Fixing of variables based on dual (i.e. optimality) implications.

So a value of 1+2=3 for MIPPRESOLVE causes reduced cost fixing and tightening of implied bounds on integer variables.

4.5 Convex QCQP and SOCP Methods

Convex continuous QCQP and SOCP problems are solved by the Xpress Newton-barrier solver. For

QCQP, SOCP and QP problems, there is no solution purification method applied after the barrier (such as the crossover for linear problems). This means that solutions tend to contain more active variables than basic solutions, and fewer variables will be at or close to one of their bounds.

When solving a linearly constrained quadratic program (QP) from scratch, the Newton barrier method is usually the algorithm of choice. In general, the quadratic simplex methods are better if a solution with a low number of active variables is required, or when a good starting basis is available (e.g., when reoptimizing).

Non-convex QCQP problems are solved by a branch-and-bound approach, see the FICO Xpress Global userguide for more details.

4.5.1 Convexity Checking

Problems with quadratic constraints or quadratic objective are automatically checked for convexity to determine whether the global solver is needed. The Newton-barrier and quadratic simplex solvers require that the quadratic coefficient matrix in each constraint or in the objective function is either positive semi-definite or negative semi-definite, depending on the sense of for constraints or the direction of optimization for the objective. The only exception is when a quadratic constraint describes a second order cone. Note that the convexity checker will reject any problem where this requirement is violated by more than a small tolerance.

Each constraint is checked individually for convexity. In certain cases it is possible that the problem itself is convex, but the representation of it is not. A simple example would be

minimize:	x		
subject to:	$x^2 - y^2 + 2xy$	$\leqslant$	1
	$y = 0$		

The global solver will be selected to solve this problem if the automatic convexity check is on, although the problem is clearly convex. The reason is that convexity of QCQPs is checked before any presolve takes place. To understand why, consider the following example:

minimize:	y		
subject to:	$y - x^2$	$\leqslant$	1
	$y = 2$		

This problem is clearly feasible, and an optimal solution is $(x, y) = (1, 2)$. However, when presolving the problem, it will be found infeasible, since assuming that the quadratic part of the first constraint is convex the constraint cannot be satisfied (remember that if a constraint is convex, then removing the quadratic part is always a relaxation). Thus since presolve makes use of the assumption that the problem is convex, convexity must be checked before presolve.

Note that for quadratic programming (QP) and mixed integer quadratic programs (MIQP) where the quadratic expressions appear only in the objective, the convexity check takes place after presolve, making it possible to accept matrices that are not PSD, but define a convex function over the feasible region (note that this is only a chance).

4.5.2 Quadratically Constrained and Second Order Cone Problems

Quadratically constrained and second order cone problems are solved by the barrier algorithm.

Mixed integer quadratically constrained (MIQCQP) and mixed integer second order problems (MISOCP) are solved using traditional branch and bound using the barrier to solve the node problems, or by means of outer approximation, as defined by control `MIQCPALG`.

It is sometimes beneficial to solve the root node of an MIQCQP or MISOCP by the barrier, even if outer approximation is used later; controlled by the `QCROOTALG` control. The number of cut rounds on the root for outer approximation is defined by `QCCUTS`.

CHAPTER 5

Advanced Usage

5.1 Problem Names

Problems loaded in the Optimizer have a name. The name is either taken from the file name if the problem is read into the optimizer or it is specified as a string in a function call when a problem is loaded into the Optimizer using the library interface. Once loaded the name of the problem can be queried and modified. For example, the library provides the function `XPRSsetprobname` for changing the name of a problem.

When reading a problem from a matrix file the user can optionally specify a file extension. The search order used for matrix files in the case where the file extension is not specified is described in the reference for the function `XPRSreadprob`. In this case, the problem name becomes the file name, including the full path, but without the file extension.

Note that matrix files can be read directly from a gzip compressed file. Recognized names of matrix files stored with gzip compression have an extension that is one of the usual matrix file format extensions followed by the `.gz` extension. For example, `hpw15.mps.gz`.

The problem name is used as a default base name for the various file system interactions that the Optimizer may make when handling a problem. For example, when commanded to read a basis file for a problem and the basis file name is not supplied with the read basis command the Optimizer will try to open a file with the problem name appended with the `.bss` extension.

It is useful to note that the problem name can include file system path information. For example, `c:/matrices/hpw15`. Note the use of forward slashes in the Windows path string. It is recommended that Windows users use forward slashes as path delimiters in all file name specifications for the Optimizer since (i) this will work in all situations and (ii) it avoids any problems with the back slash being interpreted as the escape character.

5.2 Manipulating the Matrix

In general, the basic usage of the FICO Xpress Optimizer described in the previous chapters will be sufficient for most users' requirements. Using the Optimizer in this way simply means load the problem, solve the problem, get the results and finish.

In some cases, however, it is required that the problem is first solved, then modified, and solved again. We may want to do this, for example, if a problem was found to be infeasible. In this case, to find a feasible subset of constraints we iteratively remove some constraints and re-solve the problem. Another example is when a user wants to 'generate' columns using the optimal duals of a 'restricted' LP problem. In this case we will first need to load a problem and then we will need to add columns to this problem after it has been solved.

For library users, FICO Xpress provides a suite of functions providing read and modify access to the matrix.

5.2.1 Reading the Matrix

The Optimizer provides a suite of routines for read access to the optimization problem including access to the objective coefficients, constraint right hand sides, decision variable bounds and the matrix coefficients.

It is important to note that the information returned by these functions will depend on whether or not the problem has been run through an optimization algorithm or if the problem is currently being solved using an optimization algorithm, in which case the user will be calling the access routines from a callback (see section 5.8 for details about callbacks). Note that the dependency on when the access routine is called is mainly due to the way "presolve" methods are applied to modify the problem. How the presolve methods affect what the user accesses through the read routines is discussed in section 5.3.

The user can access the names of the problem's constraints, or 'rows', as well as the decision variables, or 'columns', using the `XPRSgetnamelist` routine.

The linear coefficients of the problem constraints can be read using `XPRSgetrows`. Note that for the cases where the user requires access to the linear matrix coefficients in the column-wise sense the Optimizer includes the `XPRSgetcols` function. The type of the constraint, the right hand side and the right hand side range are accessed using the functions `XPRSgetrowtype`, `XPRSgetrhs` and `XPRSgetrhsrange`, respectively.

The coefficients of the objective function can be accessed using the `XPRSgetobj` routine, for the linear coefficients, and `XPRSgetqobj` for the quadratic objective function coefficients. The type of a column (or decision variable) and its upper and lower bounds can be accessed using the routines `XPRSgetcoltype`, `XPRSgetub` and `XPRSgetlb`, respectively.

The quadratic coefficients in constraints can be accessed either in matrix form, using the `XPRSgetqrowqmatrix` routine, or as a list of quadratic coefficients with the `XPRSgetqrowqmatrixtriplets`.

Note that the reference section in Chapter 8 of this manual provides details on the usage of these functions.

5.2.2 Modifying the Matrix

The Optimizer provides a set of routines for manipulating the problem data. These include a set of routines for adding and deleting problem constraints ('rows') and decision variables ('columns'). A set of routines is also provided for changing individual coefficients of the problem and for changing the types of decision variables in the problem.

Rows and columns can be added to a problem together with their linear coefficients using `XPRSaddrows` and `XPRSaddcols`, respectively. Rows and columns can be deleted using `XPRSdelrows` and `XPRSdelcols`, respectively.

The Optimizer provides a suite of routines for modifying the data for existing rows and columns. The linear matrix coefficients can be modified using `XPRSchgcoef` (or use `XPRSchgmcoef` if a batch of coefficients are to be changed). Row and column types can be changed using the routines `XPRSchgrowtype` and `XPRSchgcoltype`, respectively. Right hand sides and their ranges may be changed with `XPRSchgrhs` and `XPRSchgrhsrange`. The linear objective function coefficients may be changed with `XPRSchgobj` while the quadratic objective function coefficients are changed using `XPRSchgqobj` (or use `XPRSchgmqobj` if a batch of coefficients are to be changed). Likewise, quadratic coefficients in constraints are changed with `XPRSchgqrowcoeff`.

Examples of the usage of all the above functions and their syntax may be found in the reference section of this manual in Chapter 8.

Finally, it is important to note that it is not possible to modify a matrix when it has been 'presolved' and has not been subsequently 'postsolved'. The following section 5.3 discusses some important points concerning reading and modifying a problem that is ”presolved”.

5.3 Working with Presolve

The Optimizer provides a number of algorithms for simplifying a problem prior to the optimization process. This elaborate collection of procedures, known as *presolve*, can often greatly improve the Optimizer’s performance by modifying the problem matrix, making it easier to solve. The presolve algorithms identify and remove redundant rows and columns, reducing the size of the matrix, for which reason most users will find it a helpful tool in reducing solution times. However, presolve is included as an option and can be disabled if not required by setting the `PRESOLVE` control to `0`. Usually this is set to `1` and presolve is called by default.

For some users the presolve routines can result in confusion since a problem viewed in its presolved form will look very different to the original model. Under standard use of the Optimizer this may cause no difficulty. On a few occasions, however, if errors occur or if a user tries to access additional properties of the matrix for certain types of problem, the presolved values may be returned instead. In this section we provide a few notes on how such confusion may be best avoided. If you are unsure if the matrix is in a presolved state or not, check the `PRESOLVESTATE` attribute

It is important to note that when solving a problem with presolve on, the Optimizer will take a copy of the matrix and modify the copy. The original matrix is therefore preserved, but will be inaccessible to the user while the presolved problem exists. Following optimization, the whole matrix is automatically *postsolved* to recover a solution to the original problem and restoring the original matrix. Consequently, either before or after, but not during, a completed optimization run, the full matrix may be viewed and altered as described above, being in its original form.

A problem might be left in a presolved state if the solve was interrupted, for example due to the Ctrl-C key combination, or if a time limit (set by `MAXTIME`) was reached. In such a case, the matrix can always be returned to its original state by calling `XPRSpostsolve` (`POSTSOLVE`). If the matrix is already in the original state then `XPRSpostsolve` (`POSTSOLVE`) will return without doing anything.

While a problem is in a presolved state it is not possible to make any modifications to it, such as adding rows or columns. The problem must first be returned to its original state by calling `XPRSpostsolve` before it can be changed.

5.3.1 (Mixed) Integer Programming Problems

If a model contains MIP entities, integer presolve methods such as bound tightening and coefficient tightening are applied to tighten the LP relaxation. As a simple example of this might be if the matrix has a binary variable x and one of the constraints of the matrix is $x \leqslant 0.2$. It follows that x can be fixed at zero since it can never take the value 1. If presolve uses the MIP entities to alter the matrix in this way, then the LP relaxation is said to have been *tightened*. For Console users, notice of this is sent to the screen; for library users it may be sent to a callback function, or printed to the log file if one has been set up. In such circumstances, the optimal objective function value of the LP relaxation for a presolved matrix may be different from that for the unpresolved matrix.

The strict LP solution to a model with MIP entities can be obtained by calling the `XPRSlpoptimize` (`LPOPTIMIZE`) command. This removes the discrete restrictions from the variables, preventing the LP relaxation from being tightened and solves the resulting matrix. In the example above, x would not be fixed at 0, but allowed to range between 0 and 0.2.

When `XPRSmipoptimize` (`MIPOPTIMIZE`) finds an integer solution, it is postsolved and saved in memory. The solution can be read with the `XPRSgetsolution` function. A permanent copy can be saved to a solution file by calling `XPRSwritebinsol` (`WRITEBINSOL`), or `XPRSwriteslxsol` (`WRITESLXSOL`) for a simpler text file. This can be retrieved later by calling `XPRSreadbinsol` (`READBINSOL`) or `XPRSreadslxsol` (`READSLXSOL`), respectively.

After calling `XPRSmipoptimize` (MIPOPTIMIZE), the matrix will be postsolved whenever the MIP search has completed. If the MIP search hasn't completed the matrix can be postsolved by calling the `XPRSpostsolve` (`POSTSOLVE`) function.

5.4 Working with LP Folding

In addition to presolve procedures, the Optimizer provides an algorithm called LP folding that can further simplify LP problems. The LP folding is applicable to LP problems that can be partitioned into *equitable partitions*, and it works by aggregating matrix columns of equitable partitions and then reducing the problem size.

Solutions for the folded problem are also valid for the original problem. While it is straightforward to transfer a solution from the folded problem to the original problem, it is non-trivial to do so for the basis. When an LP problem is solved to optimality and a basis is needed, the LP unfolding will use the crossover algorithm to provide one. When the folded LP problem is unbounded or infeasible, or when the solving process is stopped due to time or iteration limit, the basis will not be available. Please note that LP folding tends to provide solutions with a larger support (number of variables that are not at any of their bounds).

LP folding is applied automatically when appropriate. It can be enabled or disabled by setting the `LPFOLDING` control.

5.5 Working with Heuristics

The Optimizer contains several primal heuristics that help to find feasible solutions during a tree search. These heuristics fall broadly into one of three classes:

1. Simple rounding heuristics
   These take the continuous relaxation solution to a node and, through simple roundings of the solution values for MIP entities, try to construct a feasible MIP solution. These are typically run on every node.

2. Diving heuristics
   These start from the continuous relaxation solution to a node and combines rounding and fixing of MIP entities with occasional reoptimization of the continuous relaxation to construct a better quality MIP solution. They are run frequently on both the root node and during the branch and bound tree search.

3. Local search heuristics
   The local search heuristics are generally the most expensive heuristics and involve solving one or more smaller MIPs whose feasible regions describe a neighborhood around a candidate MIP solution. These heuristics are run at the end of the root solve and typically on every 500–1000 nodes during the tree search.

Some simple heuristics and a few fast diving heuristics, which do not require a starting solution, will be tried before the initial continuous relaxation of a MIP is solved. On very simple problems, it is possible that an optimal MIP solution will be found at this point, which can lead to the initial relaxation being cut off. These heuristics can be enabled or disabled using the `HEURBEFORELP` control.

In contrast to the simple heuristics mentioned above, the Pre-root parallel heuristic phase is a separate, optional phase which can be enabled to execute more heuristics before the initial continuous relaxation. This phase is optional and not enabled by default. If enabled, it orchestrates different heuristics in parallel for a user-defined amount of algorithmic work. The phase is inspired by Alternating Criteria Search and can find both first feasible as well as improving solutions. To this end, it combines constructive heuristics such as Shift-and-propagate or Feasibility Jump from the literature with restricted sub-MIP searches, all of which are run in parallel. The work spent during this phase is either set through the PREROOTWORKLIMIT control, or can be indirectly dialed up or down using the PREROOTEFFORT control, which decides the work limit based on problem characteristics. Setting any of those two controls to 0 disables this phase. Pre-root parallel heuristics are mainly aimed at optimization problems with very time-consuming continuous relaxations. On such problems, the pre-root parallel phase can provide high-quality solutions much earlier than most of the relaxation-based heuristic approaches. The pre-root parallel heuristics employ a diversification scheme that scales with the number of available CPU threads. By default, all available threads are used. For very large problems, the required memory may exceed the system memory. In this case, the control PREROOTTHREADS can be used to reduce the number of parallel heuristic tasks. The Pre-root parallel phase heuristics currently only support MIP optimization problems without indicator constraints.

The most important control for steering the overall heuristic behavior of the Optimizer is HEUREMPHASIS Setting this control to 1 will trigger many additional heuristic calls. This setting typically leads to a quicker reduction of the optimality gap at the beginning of the search. However, it often comes at the expense of the running time to proven optimality increasing. Consequently, we recommend this setting for use cases that aim to find a reasonably good solution quickly but for which proving optimality is out of scope.

Setting HEUREMPHASIS to 2 will enable a very aggressive heuristic behavior. While this can be beneficial for some problems, it is on average inferior to the more balanced setting of 1. Setting HEUREMPHASIS to 0 will disable all heuristics. This replaces the deprecated HEURSTRATEGY=0 setting.

In addition to running the solver with a heuristic emphasis from the beginning, a secondary application of HEUREMPHASIS is solution polishing. Therefore, the user would alter the heuristic emphasis only when the solve comes close to the intended time limit or the progress in the optimality gap stalls. This can, e.g., happen from a callback. Note that this is most promising when the solver has already made some significant progress in tree search. Especially for problems that spend a majority of the targeted solution time in the root node, we recommend setting HEUREMPHASIS=1 from the beginning.

There are a few other controls that affect heuristics. The diving heuristics have the following controls:

HEURFREQ	The frequency at which to run a diving heuristic during the branch and bound tree search. If HEURFREQ=k, a diving heuristic will be applied when at least k nodes of the tree search have been solved since the last run. Set this control to zero to disable diving heuristics during the tree search. With a default setting of -1, the Optimizer will automatically select a frequency that depends on how expensive it is to run and how many integer variables need to be rounded. Typically, this results in a diving heuristic being run for every 10–50 nodes.
HEURDIVESTRATEGY	Can be used to select one specific out of 10 predefined diving strategies, otherwise the Optimizer will automatically select which appears to work best. Set this control to zero to disable the diving heuristic.
HEURDIVERANDOMIZE	How much randomization to introduce into the diving heuristics.
HEURDIVESPEEDUP	The amount of effort to put into the individual dives. This essentially determines how often the continuous relaxation is reoptimized during a dive.

The local search heuristics have the following controls:

HEURSEARCHFREQ	The frequency at which to run the local search heuristics during the branch and bound tree search. If HEURSEARCHFREQ=k, the local search heuristics will be run when at least k nodes of the tree search have been solved since the last run.

`HEURSEARCHEFFORT`	Determines the complexity of the local search MIP problems solved and, if `HEURSEARCHFREQ=-1`, also how often they are applied.
`HEURSEARCHROOTSELECT`	Selects which local search heuristics are allowed to be run on the root node. Each bit of this integer control represents an individual heuristic.
`HEURSEARCHTREESELECT`	Selects which local search heuristics are allowed to be run during the branch and bound tree search.

The simple rounding heuristics do not have any further controls associated with them.

5.6 Analyzing and Handling Numerical Issues

There are many optimization applications which give rise to numerically challenging models. You might notice that the Optimizer takes unexpectedly long for simplex reoptimization, that minimal changes in the models lead to an unexpectedly large change in the optimal solution value or that the optimal solution shows a certain amount of violation in the postsolved state. The Optimizer provides various tools to analyze whether a model is numerically challenging and to handle numerical issues when they occur.

5.6.1 Analyzing Models for Numerical Issues

There are two main reasons which can make models numerically challenging: Firstly, using coefficients that span many orders of magnitude, e.g., using numbers as large as 100 million mixed with numbers as small as 1 over 100 million. Those span 16 orders of magnitude. A double-precision floating point number, however, can only represent 15 precise digits. Thus, round-off errors are inevitable. Secondly, if a model contains structures that amplify the effect of numeric error propagation, e.g., when the result of subtracting two almost identical values is scaled up and then used for further computations.

To both ends, Xpress provides features to analyze models for numerical stability. Addressing the first issue, Xpress provides the user with information on the coefficient ranges in both the original problem and the problem that is solved after presolving and scaling has been applied. In the log, the minimum and maximum absolute values of the matrix coefficients, the right-hand side/bounds and the objective are printed. The relevant part for the numerical behavior of the solution process are the coefficient ranges in the solved model. The difference between the exponents of the min and max values tells you how many orders of magnitude are covered. As a rule of thumb, those should not be more than nine (and not more than six in an individual row or column of the original matrix). For MIP solves, Xpress will notify the user after the solution of the root LP when the coefficient ranges and other stability measures indicate that the solve might become numerically cumbersome. In such a case, it will print a warning "High attention level predicted from matrix features" to the log.

The second issue, error propagation, is a bit trickier to trace. The most important source to consider for this is the multiplication of a vector with the constraint matrix, which gets stored in a factorized fashion. Hence, it makes sense to consider the condition number of the basis inverse matrix. Computing this can be expensive and is hence not done by default. You can activate it by setting the `MIPKAPPAFREQ` control to one. When setting this control, you will get a final statistic report that summarizes the condition numbers collected during search. Besides the percentage of stable, suspicious, unstable, and ill-posed basis inverse matrices, the Optimizer will report a quantity called the attention level after the solve. The attention level takes values between zero and one. It is equal to zero if all basis inverse matrices are stable, and one if all basis inverse matrices are ill-posed. The higher the attention level, the more likely are numerical errors. As a rule of thumb, matrices with an attention level larger than 0.1 should be investigated further. The attention level is available as an attribute: `ATTENTIONLEVEL`.

After having solved the root LP relaxation of a MIP solve, the Optimizer applies a Machine Learning model to predict the attention level of the current MIP solve. If the prediction is larger than 0.1, it will print a message to the log: "High attention level predicted from matrix features". The predicted attention level is available as an attribute: PREDICTEDATTLEVEL. Finally, if the Optimizer undergoes numerical failures during the optimization process, it will report these at the end of the solve. If you see dual, primal or barrier failures, or single inverts being reported, it might be worthwhile to try some of the methods described in the following sections.

5.6.2 Scaling

Scaling is a widely used preconditioning technique that aims at reducing the condition number of the constraint matrix, at reducing error propagation, and at reducing the number of LP iterations required to solve the problem. In Xpress, both columns and rows are scaled, however, only by powers of 2 to avoid round-off errors. By default, Xpress applies a machine learning algorithm to choose a scaling variant that is predicted to give the most stable performance. Although this prediction is correct in most of the cases, one can try the opposite setting, i.e., setting SCALING to 163 when autoscaling selected Curtis-Reid scaling and setting scaling to 16 when autoscaling selected standard scaling. Furthermore, disabling special handling of big-M rows and conduction scaling before presolving, represented by bits 6 and 9 of the SCALING control, is useful for some problems.

5.6.3 Solution Refinement

The Optimizer offers two methods of refining solutions, both are independent and complement each other. The first is called LP Refinement and aims at providing LP solutions of a higher precision, i.e., with more significant bits. It consists of two parts. Standard LP Refinement iteratively attempts to increase the accuracy of the solution until either both FEASTOLTARGET and OPTIMALITYTOLTARGET are satisfied, or accuracy cannot further be increased, or some effort limit is exhausted. It is applied by default both to LP solutions and to MIP solutions. Iterative refinement has the same goal, but uses more expensive, but also more promising measures of doing so, e.g., quad precision computing. If the postsolved LP solution is slightly infeasible, setting bits 5 and 6 of the REFINEOPS control aims at reducing those infeasibilities.

The second refinement scheme is called MIP Refinement and aims at providing MIP solutions which are truly integral and will not lead to infeasibilities when fixing integer variables in the original space. Note that both Iterative Refinement and MIP Refinement can lead to a slowdown of the solution process which is more considerable the more numerically challenging the matrix is.

5.6.4 Other Ways to Handle Numerical Issues

In addition to the methods named above, the Optimizer gives the user the possibility to change the numerical tolerances, such as FEASTOL and MATRIXTOL, but caution is advised here. Finally, if the numerical issues mainly come from the behavior of the simplex algorithm, setting DUALSTRATEGY to values 7 or 32 might help, or even using only barrier for solving LPs during a MIP solve, achieved by changing DEFAULTALG to 4.

In any case, it is best practice to reconsider the model. If you have very small and/or very large values in there — are those really necessary? Or could they be adapted to some significantly more stable value while still representing the same logic? Can you determine places where large values might cancel each other out and the residual is used for further computations? Have you tried using indicator instead of big-M formulations?

5.7 Common Causes of Confusion

It should be noted that most of the library routines described above and in chapter 8, which modify the matrix will not work on a presolved matrix. The only exception is inside a callback for a MIP solve, where cuts may be added or variable bounds tightened (using XPRSchgbounds). Any of these functions expect references to the presolved problem. If one tries to retrieve rows, columns, bounds or the number of these, such information will come from the presolved matrix and not the original. A few functions exist which are specifically designed to work with presolved and scaled matrices, although care should be exercised in using them. Examples of these include the commands `XPRSgetpresolvesol`, `XPRSgetpresolvebasis`, `XPRSgetscaledinfeas`, `XPRSloadpresolvebasis` and `XPRSloadpresolvedirs`.

5.8 Using the Callbacks

Console users are constantly provided with information on the standard output device by the Optimizer as it searches for a solution to the current problem. The same output is also available to library users if a log file has been set up using `XPRSsetlogfile`. However, whilst Console users can respond to this information as it is produced and allow it to influence their session, the same is not immediately true for library users, since their program must be written and compiled before the session is initiated. For such users, a more interactive alternative to the above forms of output is provided by the use of *callback functions*.

The library *callbacks* are a collection of functions which allow user-defined routines to be specified to the Optimizer. In this way, users may define their own routines which should be called at various stages during the optimization process, prompting the Optimizer to return to the user's program before continuing with the solution algorithm. Perhaps the three most general of the callback functions are those associated with the search for an LP solution. However, the vast majority of situations in which such routines might be called are associated with the tree search, and will be addressed below.

5.8.1 Output Callbacks

Instead of catching the standard output from the Optimizer and saving it to a log file, the callback `XPRSaddcbmessage` allows the user to define a routine which should be called every time a text line is output by the Optimizer. Since this returns the status of each message output, the user's routine could test for error or warning messages and take appropriate action accordingly.

5.8.2 LP Callbacks

The functions `XPRSaddcblplog` and `XPRSaddcbbarlog` allow the user to respond after each iteration of either the simplex or barrier algorithms, respectively. The controls `LPLOG` and `BAROUTPUT` may additionally be set to reduce the frequency at which these routines should be called.

5.8.3 MIP Tree Search Callbacks

When a problem with MIP entities is to be optimized, a large number of subproblems, called *nodes*, must typically be solved as part of the branch and bound tree search. At various points in this process user-defined routines can be called, depending on the callback that is used to specify the routine to the Optimizer.

In a branch and bound tree search, the Optimizer starts by selecting an active node amongst all candidates (known as a *full backtrack*) and then proceed with solving it, which can lead to new descendent nodes being created. If there is a descendent node, the optimizer will by default select one of these next to solve and repeat this iterative descend while new descendent nodes are being created. This *dive* stops when it reaches a node that is found to be infeasible or cutoff, at which point the Optimizer will perform a *full backtrack* again and repeat the process with a new active node.

A routine may be called whenever a node is selected by the optimizer during a *full backtrack*, using `XPRSaddcbchgnode`. This will also allow a user to directly select the active node for the optimizer. Whenever a new node is created, a routine set by `XPRSaddcbnewnode` will be called, which can be used to record the identifier of the new node, e.g. for use with XPRSaddcbchgnode.

When the Optimizer solves a new node, it will first call any routine set by XPRSaddcbprenode, which can be used to e.g. tighten bounds on columns (with XPRSchgbounds) as part of a user node presolve. Afterwards, the LP relaxation of the node problem is solved to obtain a feasible solution and a best bound for the node. This might be followed by one or more rounds of cuts. Before any potential round of cuts, it will first call any routine set by XPRSaddcbcutround. This function allows a user to provide their own cutting planes with XPRSaddmanagedcuts, where the Optimizer will manage the automatic loading of violated cuts and the removal of inactive cuts. The callback can also be used for overriding the Optimizer's decision on when to apply a round of cutting planes. If the node problem is found to be infeasible or cutoff during this process, a routine set by XPRSaddcbinfnode will be called. Otherwise, a routine set by XPRSaddcboptnode will be called to let the user know that the optimizer now has an optimal solution to the LP relaxation of the node problem. In this routine, the user is allowed to add cuts (see section 5.9) and tighten bounds to tighten the node problem, or apply branching objects (see XPRS_bo_create) to separate on the current node problem. If the user modifies the problem inside this *optnode* callback routine, the optimizer will automatically resolve the node LP and, if the LP is still feasible, call the *optnode* routine again.

If the LP relaxation solution to the node problem also satisfies all MIP entities and the user has not added any branching objects, i.e., if it is a MIP solution, the Optimizer will call a routine set by XPRSaddcbpreintsol *before* saving the new solution, and call a routine set by XPRSaddcbintsol *after* saving the solution. These two routines will also be called whenever a new MIP solution is found using one of the Optimizer heuristics.

Otherwise, if the node LP solution does not satisfy the MIP entities (or any user branching objects), the Optimizer will proceed with branching. After the optimizer has selected the candidate entity for branching, a routine set by XPRSaddcbchgbranch will be called, which also allows a user to change the selected candidate. If, during the candidate evaluation the optimizer discovers that e.g. bounds can be tightened, it will tighten the node problem and go back to resolving the node LP, followed by the callback routines explained above.

When the Optimizer finds a better MIP solution, it is possible that some of the nodes in the active nodes pool are cut off due to having an LP solution bound that is worse than the new cutoff value. For such nodes, a routine set by XPRSaddcbnodecutoff will be called and the node will be dropped from the active nodes pool.

The final MIP-related callback, XPRSaddcbmiplog, is more similar to the LP search callbacks, allowing a user's routine to be called whenever a line of the MIP log is printed. The frequency with which this occurs is set by the control MIPLOG.

5.9 Working with the Cut Manager

5.9.1 Cuts and the Cut Pool

Solving the LP relaxations during a branch and bound tree search is often made more efficient by supplying additional rows (constraints) to the matrix which reduce the size of the feasible region, whilst ensuring that it still contains an optimal integer solution. Such additional rows are called *cutting planes*, or *cuts*.

By default, cuts are automatically added to the matrix by the Optimizer during a tree search to speed up the solution process. However, for advanced users, the Optimizer library provides greater freedom, allowing the possibility of choosing which cuts are to be added at particular nodes, or removing cuts entirely. The cutting planes themselves are held in a *cut pool*, which may be manipulated using library functions.

Cuts may be handled directly by the user (see section 5.9.2) or can be added to the Optimizer to be managed automatically (see section 5.9.3). In the first case, cuts can be added directly to the matrix at a particular node or stored in the cut pool first before subsequently being loaded into the matrix. It often makes little difference which of these two approaches is adopted, although as a general rule if cuts are cheap to generate, it may be preferable to add the cuts directly to the matrix and delete any redundant cuts after each sub-problem (node) has been optimized. Any cuts added to the matrix at a node and not deleted at that node will automatically be added to the cut pool. If you wish to save all the cuts that are generated, it is better to add the cuts to the cut pool first. Cuts can then be loaded into the matrix from the cut pool. This approach has the advantage that the cut pool routines can be used to identify duplicate cuts and save only the stronger cuts.

To help track the cuts that have been added to the matrix at different nodes, the cuts can be classified according to a user-defined *cut type*. The cut type can either be a number such as the node number or it can be a bit-vector (see section 9.2). In the latter case each bit of the cut type may be used to indicate a property of the cut. For example, cuts could be classified as local cuts applicable at the current node and its descendants, or as global cuts applicable at all nodes. If the first bit of the cut type is set this could indicate a local cut and if the second bit is set this could indicate a global cut. Other bits of the cut type could then be used to signify other properties of the cuts. The advantage of using bit maps is that all cuts with a particular property can easily be selected, for example all local cuts.

5.9.2 Cut Management Routines

Cuts may be added directly into the matrix at the current node using `XPRSaddcuts`. Any cuts added to the matrix at a node will be automatically added to the cut pool and hence restored at descendant nodes unless specifically deleted at that node, using `XPRSdelcuts`. Cuts may be deleted from a parent node which have been automatically restored, as well as those added to the current node using `XPRSaddcuts`, or loaded from the cut pool using `XPRSloadcuts`.

It is recommended to delete only those cuts with basic slacks. Otherwise, the basis will no longer be valid and it may take many iterations to recover an optimal basis. If the second argument to `XPRSdelcuts` is set to `1`, this will ensure that cuts with non-basic slacks will not be deleted, even if the other controls specify that they should be. It is highly recommended that this is always set to `1`.

Cuts may be saved directly to the cut pool using the function `XPRSstorecuts`. Since cuts added to the cut pool are *not* automatically added to the matrix at the current node, any such cut must be explicitly loaded into the matrix using `XPRSloadcuts` before it can become active. If the third argument of `XPRSstorecuts` is set to `1`, the cut pool will be checked for duplicate cuts with a cut type identical to the cuts being added. If a duplicate cut is found, the new cut will only be added if its right hand side value makes the cut stronger. If the cut in the cut pool is weaker than the added cut, it will be removed unless it has already been applied to active nodes of the tree. If, instead, this argument is set to `2`, the same test is carried out on all cuts, ignoring the cut type. The routine `XPRSdelcpcuts` allows the user to remove cuts from the cut pool, unless they have already been applied to active nodes in the branch and bound tree.

A list of cuts in the cut pool may be obtained using the command `XPRSgetcpcuts`, whilst `XPRSgetcpcutlist` returns a list of their indices. A list of those cuts which are active at the current node is returned using `XPRSgetcutlist`.

5.9.3 Optimizer Managed Cuts

An alternative to managing cuts directly is to let the Optimizer manage the provided cuts, by loading them using the function `XPRSaddmanagedcuts` within a cutround callback set by `XPRSaddcbcutround`. The Optimizer will automatically load any such cuts that are violated into a node and remove cuts that are no longer violated, as it deems necessary. These cuts are not stored in the user cut pool and cannot be accessed using the functions `XPRSgetcutlist` and `XPRSgetcpcutlist`.

5.9.4 User Cut Manager Routines

Users may also write their own cut manager routines to be called during the branch and bound search. Cuts can be added or removed on any node of the branch and bound search using a callback function set by the routine `XPRSaddcboptnode` (see section 5.8.3).

Further details of these functions may be found in chapter 8 within the functional reference which follows.

5.10 Solving Problems Using Multiple Threads

It is possible to use multiple processors when solving any type of problem with the Optimizer. On the more common processor types, such as those from Intel or AMD, the Optimizer will detect how many logical processors are available in the system and attempt to solve the problem in parallel using as many threads as possible. The number detected can be read through the `CORESDETECTED` integer attribute. It is also possible to adjust the number of threads to use by setting the integer parameter `THREADS`.

By default a problem will be solved deterministically, in the sense that the same solution path will be followed each time the problem is solved when given the same number of threads. For an LP this means that the number of iterations and the optimal, feasible solution returned will always be the same.

When solving a MIP deterministically, each node of the branch-and-bound tree will always be solved the same. Each node of the branch-and-bound tree can be identified by a unique number, available through the attribute `CURRENTNODE`. The tree will always have the same parent/child relationship in terms of these identifiers. A deterministic MIP solve will always find integer solutions on the same nodes and the attributes and solutions on a node will always be returned the same from one run to another. Since nodes will be solved in parallel the order in which nodes are solved can vary. There is an overhead in synchronizing the threads to make the parallel runs deterministic and it can be faster to run in non-deterministic mode. This can be done by setting the `DETERMINISTIC` control to 0.

For an LP problem (or the initial continuous relaxation of a MIP), there are several choices of parallelism. The Newton barrier, the hybrid gradient and the dual simplex algorithms support multiple threads. The number of threads to use can be set with `BARTHREADS` or `DUALTHREADS`, respectively. It is also possible to run some or all of primal simplex, dual simplex and the Newton barrier or hybrid gradient algorithm side-by-side in separate threads, known as a *concurrent* LP solve. (The Newton barrier and the hybrid gradient algorithm cannot be run concurrently.) This can be useful when none of the methods is the obvious choice. In this mode, the Optimizer will stop with the first algorithm to solve the problem. The number of threads for the concurrent LP solve can be set using `CONCURRENTTHREADS`. The algorithms to use for the concurrent solve can be specified by concatenating the required "d", "p", "n" and "b" flags when calling `XPRSlpoptimize` (`LPOPTIMIZE`) or `XPRSmipoptimize` (`MIPOPTIMIZE`); please refer to section 5.10.1 for more details.

When solving a MIP problem, the Optimizer will try to run the branch and bound tree search in parallel. Use the `MIPTHREADS` control to set the number of threads specifically for the tree search.

The operation of the optimizer for MIPs is fairly similar in serial and parallel mode. The MIP callbacks can still be used in parallel and callbacks are called when each MIP worker problem is created and destroyed. The mipthread callback (declared with `XPRSaddcbmipthread`) is called whenever a MIP worker problem is created and the callback declared with `XPRSaddcbdestroymt` is called whenever the worker problem is destroyed. Each worker problem has a unique ID which can be obtained from the `MIPTHREADID` integer attribute. When an executing thread solves a branch-and-bound node, it will also do so on a worker problem assigned to it. Note that a given worker problem can be assigned to different threads during its lifetime and the threads might differ from one run to another.

When the MIP callbacks are called they are MUTEX protected to allow non threadsafe user callbacks. If a significant amount of time is spent in the callbacks then it is worth turning off the automatic MUTEX protection by setting the `MUTEXCALLBACKS` control to 0. It this is done then the user must ensure that their callbacks are threadsafe.

On some problems it is also possible to obtain a speedup by using multiple threads for the MIP solve process between the initial LP relaxation solve and the branch and bound search. The default behavior here is for the Optimizer to use a single thread to create its rounds of cuts and to run its heuristic methods to obtain MIP solutions. Extra threads can be started, dedicated to running the heuristics only, by setting the `HEURTHREADS` control. By setting `HEURTHREADS` to a non-zero value, the heuristics will be run in separate threads, in parallel with cutting.

When a MIP solve is terminated early, due to e.g. a time or node limit, it is possible to select between two different termination behaviors. This has implications for the determinism of callbacks called near termination and how quickly the Optimizer stops. In the default behavior, when termination is detected, all work is immediately stopped and any partial node solves are discarded. It is therefore possible that some callbacks will have been called for nodes that are discarded at termination. Note that this termination method does not affect the final state the problem is left in after termination and that any integer solution for which the *preintsol* and *intsol* callbacks are called will never be dropped. By setting the control `MIPTERMINATIONMETHOD` to 1, the termination behavior will be changed such that partial work is never discarded. Instead, all worker threads will be allowed to complete their current work before the solve stops. This termination behavior might cause a longer delay between termination is detected and the Optimizer stops, but it will ensure that work is never dropped for any callbacks that have already been called.

5.10.1 The concurrent solver

The concurrent solve is activated either by passing multiple algorithm flags to `XPRSlpoptimize` (e.g. ”pb” for running primal and the barrier) or by setting `CONCURRENTTHREADS` to a positive number. The order in which threads are allocated to the algorithms is not affected by the order of the flags provided.

Instead of passing flags, you can alternatively set the `LPFLAGS` control.

To activate the hybrid gradient method you can use the ”b” flag and set `BARALG` to 4. This will start the hybrid gradient method instead of the Newton barrier algorithm.

If algorithm flags are specified or the `LPFLAGS` control is set, then concurrent will run the specified algorithms, provided that the setting of `CONCURRENTTHREADS` allows for a sufficient number of threads. When no flags are specified, the automatic order of selecting algorithms starts with dual, followed by barrier. Primal simplex will be added as a third solver only in a minority of cases, mainly when the problem has much more columns than rows. The network solver is only used if specified by flags.

`CONCURRENTTHREADS` represents the total target number of threads that can be used by concurrent. The optimizer will then first start dual then barrier (if `CONCURRENTTHREADS` >1), in some cases, see above, followed by primal (if `CONCURRENTTHREADS` >2). Any remaining threads will be allocated to parallel barrier.

If manual algorithm selection has been made using algorithm flags, then `CONCURRENTTHREADS` will limit the number of algorithms started (if smaller than the number of algorithms provided), in which case the number of algorithms started will be the first `CONCURRENTTHREADS` in the dual → barrier or hybrid gradient → primal → network order.

Once an algorithm is started, the direct thread controls `BARTHREADS` and `DUALTHREADS` are respected. Note that due to the latter controls the total number of theads may exceed `CONCURRENTTHREADS`.

In case a single algorithm is started and relevant controls are on automatic, the value of the `THREADS` control is used.

If multiple algorithms have been started and `CONCURRENTTHREADS` is on automatic, then `THREADS` will be used as the overall number of threads used in concurrent (unless overwritten by the relevant algorithm specific control on a per-algorithm basis).

5.11 Solving Large Models (the 64-bit Functions)

The size of the models that can be loaded into the optimizer using the standard optimizer functions is limited by the largest number that can be held in a 32-bit integer. This means that it is not possible to load any problem with more than 2147483648 matrix elements with the standard optimizer functions. On 64-bit machines, it is possible to use the optimizer 64-bit functions to load problems with a larger number of elements (these functions have 64 appended to the standard optimizer function names). For example, it is possible to load a problem with a large number of elements with the `XPRSloadlp64` function. The only difference between `XPRSloadlp64` and `XPRSloadlp` is that the `mstart` array containing the starting points of the elements in each column that is passed to `XPRSloadlp64` is a pointer to an array of 64-bit integers. Typically, the declaration and allocation of space for the 64-bit `mstart` array would be as follows:

```
XPRSint64 *mstart = malloc( ncol * sizeof( *mstart) );
```

The starting points of the elements in `mstart` can then exceed the largest 32-bit integer.

Wherever there is a need to pass a large array to an optimizer subroutine there is a corresponding 64-bit function. Once a large model has been loaded into the optimizer then all the standard optimizer functions (such as `XPRSlpoptimize`) can be used.

Note that although the 64-bit functions allow very large models to be loaded, there is no guarantee that such large problems can be solved in a reasonable time. Also if the machine doesn't have sufficient memory, the matrix will be constantly swapped in and out of memory which can slow the solution process considerably.

5.12 Multi-objective Optimization

Xpress supports problems with more than one objective function. Each objective function has a priority and a weight. Xpress will solve the problem once for each distinct objective priority that is defined, optimizing in each iteration a linear combination of the objective functions with the same priority. In the second and subsequent iterations, all objectives from previous iterations are fixed to their optimal values within some tolerance. Xpress supports a blended approach, where all objectives have equal priority but different weights; a lexicographic approach, where each objective has a different priority (and unit weight); and a hybrid approach, where several iterations of the problem are solved, each iteration consisting of solving a linear combination of several objectives.

A problem always has at least one objective, accessed with `XPRSgetobj` and `XPRSchgobj`. Additional objectives are added using `XPRSaddobj`, updated using `XPRSchgobjn`, queried using `XPRSgetobjn` and deleted using `XPRSdelobj`. Adding or deleting an objective updates the `OBJECTIVES` attribute. Objectives can be assigned names using `XPRSaddnames`. Calling `XPRSgetobjn` / `XPRSchgobjn` with an objective index of zero is equivalent to calling `XPRSgetobj` / `XPRSchgobj`.

Each objective has several controls which modify its behaviour. These controls can be modified using `XPRSsetobjintcontrol` and `XPRSsetobjdblcontrol`, and queried using `XPRSgetobjintcontrol` and `XPRSgetobjdblcontrol`.

Name	Type	Default value	Meaning
XPRS_OBJECTIVE_PRIORITY	integer	0	The priority of the objective. Objectives with higher prio
XPRS_OBJECTIVE_WEIGHT	double	1	The weight of the objective when blending it with other
XPRS_OBJECTIVE_ABSTOL	double	0.001	The absolute tolerance used when fixing the objective to
XPRS_OBJECTIVE_RELTOL	double	0.001	The relative tolerance used when fixing the objective to
XPRS_OBJECTIVE_RHS	double	0	The constant part of the objective.

The sense of the problem as defined by `OBJSENSE` holds for all objectives, but the sense of each objective can be reversed by assigning a negative weight. The absolute and relative tolerances are added together when computing the value to which a solved objective will be fixed in the next iteration. For minimization problems:

```
objective <= optimal_value * (1 + reltol) + abstol
```

For maximization problems:

```
objective >= optimal_value * (1 - reltol) - abstol
```

Note that when reduced cost fixing is used (the default when solving LPs), the tolerances have a different meaning, as discussed below.

A multi-objective problem may only contain linear objective functions. The problem itself may be any kind supported by Xpress, but any nonlinear objective terms must be modeled using a transfer variable. The problem is solved by calling the appropriate function, depending on the problem type: `XPRSlpoptimize`, `XPRSmipoptimize`, `XPRSnlpoptimize` or `XPRSoptimize`.

After solving a multi-objective problem, `SOLVEDOBJS` is set to the number of solves that were performed. If all objectives are solved successfully, this will be equal to the number of distinct priorities assigned to objectives, excluding objectives with weight equal to zero (which are considered to be disabled). If problems were encountered, such as infeasibility due to numerical issues, then `SOLVEDOBJS` is set to the number of solves that were attempted, including the failed solve. The order in which the objectives are solved can be calculated as follows:

- Remove all objectives with zero weight;
- sort all remaining objectives by priority in decreasing order;
- for each distinct priority level, combine all objectives with that priority using their weights.

The attributes associated with each solve can be queried using `XPRSgetobjintattrib` and `XPRSgetobjdblattrib`. For example, `XPRSgetobjintattrib` can be used with `MIPSOLS` to query the number of MIP solutions that were found when optimizing each objective. The final value of each objective function can be queried using `XPRScalcobjn`.

Two additional callbacks are available during a multi-objective solve: the `beforeobjective` callback is fired before each optimization, and the `afterobjective` is fired after each optimization. See `XPRSaddcbbeforeobjective` and `XPRSaddcbafterobjective` for more information.

For all problem types except LPs, each objective is fixed to its optimal value in subsequent iterations by adding an additional row during presolve. For LPs, the default behaviour is to instead fix all variables whose reduced costs exceed the absolute tolerance associated with the objective. In this case, the relative tolerance is ignored. Reduced cost fixing can be disabled using `MULTIOBJOPS`.

Multi-objective problems can be loaded from MPS and LP files. See A.2.3 and A.3.8 for more information.

5.13 Using the Tuner

For a given optimization problem, setting suitable control parameters frequently results in improved performance such as solution time reduction. The Xpress Optimizer built-in tuner can help a user to identify such set of control settings that allows the Xpress Optimizer, Xpress SLP, or Xpress Global to solve problems faster than by using defaults.

For instances that should be tuned for heuristic performance, we recommend trying `HEUREMPHASIS`=1 first. This setting addresses instances that will likely not solve to proven optimality within a given time limit. It aims at reducing the primal-dual gap in an early stage of the phase, primarily by running more aggressive heuristics.

5.13.1 Basic Usage

With a loaded problem, the tuner can be started by simply calling `TUNE` from the console, or `XPRStune` from a user application. The tuner will then search for better control settings from a list of controls (called the tuner method). To achieve this, the tuner will solve the problem with its default baseline control settings and then solve the problem multiple times with each individual control and certain combinations of these controls.

As the tuner works by solving a problem mutiple times, it is important and recommended to set time limits. Setting `MAXTIME` will limit the effort spent on each individual solve and setting `TUNERMAXTIME` will limit the overall effort of the tuner.

The tuner works with LP and MIP problems. It automatically determines the problem type by examining the characteristics of the current problem. It is possible to tune a MIP problem as an LP or vice versa by passing the flag `l` or `g` to `XPRStune` or `TUNE`.

The tuner can also work with nonlinear problems when Xpress Nonlinear or Xpress Global is available. For tuning a problem with Xpress Nonlinear, you need to either set the `NLPSOLVER` control to `XSLP_NLPSOLVER_LOCAL` or pass the flag `s` to `XPRStune` or `TUNE`.

5.13.2 The Tuner Method

A tuner method consists of a list of controls for the tuner to try with. It is possible to run the tuner with different pre-defined lists of controls, so-called factory tuner methods, or with a user-defined list of controls. When using the tuner, it will automatically choose a default factory tuner method according to the problem type. A non-default factory tuner method can be selected by setting the `TUNERMETHOD` control. There are several choices available for factory tuner methods, among them:

- A simple MIP method, which only features a few controls and can be used in situations where tuning with the default method would take too long, e.g., because the instance to be tuned takes a long time for each individual solve
- A comprehensive MIP method, which features a larger list of controls (and control settings) and can be used when individual instance solves are relatively fast or the default method could not reveal a satisfying improvement
- A root-focus method, which only considers controls that affect the root node processing of the MIP solve. It can either be used when root and tree behavior should be tuned independently, in a two stage process, or when it is evident that improvements have to come from root node processing. When tuning with a root-focus, it might make sense to choose minimizing the primal dual integral as a tuner target.
- A tree-focus method, which only considers controls that affect the tree search behavior of the MIP solve. It can either be used when root and tree behavior should be tuned independently, in a two stage process, or when it is evident that improvements have to come from the tree search, e.g., because the dual bound needs better branching.
- A method for tuning primal heuristics, which should be used when finding a better MIP solutions is the only focus and improving the best bound can be neglected. In this case, it might make sense to choose improvement of the primal bound also as a tuner target.
- Methods which specifically focus on SLP, MISLP, or global (MI)NLP solves. Those will be selected automatically for nonlinear problems when Xpress Nonlinear or Xpress Global are available.

Please also refer to the documentation of the TUNERMETHOD control.

A tuner method can be written out using XPRStunerwritemethod. This function will create a file in XTM format, that is effectively a list of Xpress Optimizer controls, each with a set of possible settings to try in tuning. When writing out one of the factory methods, it is recommended to first select the tuner method by setting TUNERMETHOD, or to load a targeting problem, so that the tuner can write out suitable tuner methods for the respective problem types.

Users can provide their own method to the tuner by setting up an XTM file (or editing one that has been written out). This can be read into the tuner with XPRStunerreadmethod.

An alternate way to load a user-defined tuner method is to set the TUNERMETHODFILE control to the file name. This will only work when no tuner method has been loaded by explicitly calling XPRStunerreadmethod. If a user-defined method is successfully loaded, the tuner will use it and not load any factory tuner method.

Please refer to Appendix A.8 for the format of tuner method file.

5.13.3 The Tuner Output

While the tuner examines various control settings, it prints a progress report to the console. At the same time, it writes out the result and individual logs to the file system.

On the console, the tuner will print a one-line summary for each finished run. When a new better control setting is identified, it will be highlighted with an asterisk (*) at the begining of its log line, and followed by details of the control setting and its log file name. The console progress logging can be switched off by disabling the OUTPUTLOG control. Please refer to Appendix A.14 for a more detailed description of the tuner logging.

In the background, the tuner will output the result and individual logs to the file system. By default, all the output files will be stored in the directory tuneroutput/probname/. The root folder path can be changed by setting the TUNEROUTPUTPATH control. This is the central folder in which all subfolders for the results and logs of different problems will be stored. The subfolders themselves are automatically named using the current problem name. They can be manually given a different name by setting the TUNERSESSIONNAME control. The subfolder contains one result file in XML format, and many log files, one for each evaluated control setting. The XML result file consists of the control settings, solution results and pointers to the log files of all finished tuner runs.

The file output can be turned off completely by disabling the TUNEROUTPUT control.

5.13.4 The Tuner Target

A tuner target defines how to compare two finished runs with different control settings.

A common usage of the tuner is to pursue a solution time reduction, where two runs will be compared by their solution time, the faster one is considered the better. However, when both of the runs time out, it will be more meaningful to compare other attributes of the two runs, for example the final gap or the best integer solution for MIP problems.

The tuner will choose a default tuner target according to problem types. For instance, comparing the time firstly and then the gap is the default tuner target for MIP problems. A user can select a different target by setting the TUNERTARGET control. Please refer to the documentation of TUNERTARGET for a list of supported tuner targets.

5.13.5 Restarting the Tuner

When tuning the same problem again, the tuner will attempt to pick up results from previous tuner runs so that it can avoid testing with the same control settings again. For this, it checks whether an XML result file is available in the directory `tuneroutput/probname/`, see Section 5.13.3. Reusing of the history results even works when a user changes the baseline settings or uses a different tuner method. In this case, the tuner will only pick up history results which match the new control combinations. By default, when a new control setting is evaluated, the result will be appended to the existing result file from the previous tuner session.

This feature of reusing and appending to previous results can be switched off by setting the `TUNERHISTORY` control. This control has the default value `2`, which allows both, reusing and appending. Setting it to `1` will switch off reusing of the results, while still allowing to append new result to the XML result file. Setting it to `0` will switch off appending as well; consequently, the old result file will be overwritten. Note that all log files from previous tuner session will always be kept even if they run with identical settings. This is realized by having a time stamp and a unique number in the file name. Log files can only be removed manually.

5.13.6 Tuner with Multiple Threads

The tuner can work in parallel, i.e., it can run several evaluations of different control settings simultaneously. When setting `TUNERTHREADS` larger than 1, the tuner will start in parallel mode with the given number of threads. Setting the tuner threads won't affect the number of threads used by each individual run. However, it is natural that, when solving different control settings in parallel, each of the runs may slow down.

When using the parallel tuner, it is worth considering to set the `THREADS` control as well; ideally such that the product of `THREADS` and `TUNERTHREADS` is at most the number of system threads.

5.13.7 Tuner with Problem Permutations

For a certain problem, there may exist several "lucky" controls, that show a better performance by coincidence and not due to structural reasons. Such lucky controls will typically not work with other problems of the same type, or when a user modifies the problem slightly or updates Xpress. They can be thought of as false positives of tuning.

To address this issue, the tuner can exploit a phenomenon known as performance variability and solve the problem with multiple random permutations. When setting `TUNERPERMUTE` to a positive number, for each control setting, the tuner will solve the original problem and the corresponding number of permuted problems and finally aggregate their results as one. Generally, tuner results with permutations are expected to be more stable.

5.13.8 Tuning a Set of Problems

The tuner can tune a set of problems to search for an overall best control setting for all the problems in the set. Tuning a problem set can be started from the optimizer console with the command

```
tune probset problem.set,
```

where the problem.set is a plain text file which contains a list of problem files in MPS or LP format.

The tuner starts by checking all the problems defined in the problem set file. It will read in each problem to find out its type (one of LP, MIP, SLP, MISLP or Global) and optimization direction. When there are mixed problem types, the tuner will quit with a warning message. The tuner can work with mixed optimization directions and it will treat the whole problem set as a minimization problem. For a given problem set, it is possible to force the tuner to tune the problem set as LP or MIP problems with the command

`tune lpset problem.set` or `tune mipset problem.set` respectively.

For a problem set, the tuner works by solving each individual problem in the set for each specific combination of control settings separately. When all the problems in the set are solved for a specific control setting, the tuner combines the individual problem results into a consolidated one and reports it on the console. During the solve, for each problem in the set, the tuner will output its result and log files to a path defined by `TUNEROUTPUTPATH/PROBLEMNAME`. For the main problem set, the tuner will write the consolidated results to the main output path, together with a concatenated copy of all the individual problem logs.

When tuning a problem set again, the tuner can pick up the result of existing runs for the main problem set and for each separate problem in the set as well. If the full problem set can be recovered from the existing tuning records, the tuner will omit solving them as usual. Otherwise, the tuner will go through all the problems in the set. For each problem in the set, the tuner will also check whether it is possible to pick up an existing result with the specific control setting and omit solving for existing ones when possible.

5.13.9 Advanced Topics

Besides explicitly calling `TUNE` or `XPRStune`, the tuner can also be started by enabling the `TUNERMODE` control. When enabling this control (setting to `1`), all the optimization such as `XPRSmipoptimize` or `XPRSlpoptimize` will be carried out as a tuned optimization. The Optimizer will first use the tuner to find the best setting and then apply the best setting to solve the problem. On the other hand, a user can disable this control (setting to `0`) to always disable the tuner, such that a call to `XPRStune` will have no effect. This `TUNERMODE` has a default value of `-1`, which won't affect the behaviour of any of the above mentioned functions.

When using the tuner from a user application with callbacks, the callbacks will also be passed on to each individual runs. A user needs to keep in mind that these callbacks may be called mutiple times from the tuner, as the tuner will solve the problem mutiples times. Moreover, when using the parallel tuner, it is the user's responsibility to ensure that callbacks are thread-safe.

Though the tuner is built-in with the Xpress Optimizer, it can tune nonlinear problems when Xpress Nonlinear or Xpress Global is available. Currently, parallel tuning and permutations will be disabled in this case.

5.14 Remote Solving with Xpress Insight Compute Interface

The Xpress Optimizer libraries can be configured to outsource optimization computation to a remote Insight server that supports the Compute Interface. Software applications which depend on the Optimizer libraries for optimization computation therefore inherit the ability to transparently send jobs to Insight. This includes the Xpress applications (Optimizer Console, Mosel, Workbench).

When a solve is started, the Optimizer library directs any operations that can be solved remotely to the remote server. Some features such as callbacks, multi-start, and the solution enumerator have restrictions applied which are documented here.

To integrate Xpress with an Insight Compute Server you must provide some configuration. Please see Chapter 3 of the Xpress Insight Compute Interface guide here:
(https://www.fico.com/fico-xpress-optimization/docs/latest/insight5/compute/)

The single solve operations XPRSlpoptimize, XPRSmipoptimize, XPRSnlpoptimize, XPRSiis are supported. Calls to XPRSrepairinfeas and XPRStune which generate multiple problem solves are also supported and each solve will be outsourced to the remote Insight server. The number of parallel solves in the tuner is driven by the `TUNERTHREADS` control.

The Xpress Insight execution service and the local client application must be using the same major version of Xpress. Remote solves by Insight are supported by Xpress v8.10 and higher. Note: If you get a solve path difference, update the version of Xpress to match the version on the server with that on the client machine and check hardware controls, in particular threading controls. Solves will use the default execution service unless you specify one using the `COMPUTEEXECSERVICE` control or the configuration file as described in Appendix A.15.3.

Compute solves do not support the continuation of solves once they are interrupted, nor the multistart nonlinear algorithm.

A remote solve can be terminated by calling XPRSinterrupt. When called from the supported callbacks - with the exception of the message callback - this will stop the optimizer the same way as for a local solve. Otherwise, calling XPRSinterrupt outside of the supported callbacks will terminate the solve at the earliest opportunity, and no results will be generated.

The remote solve is resilient to a temporary loss of connection between client and server. Xpress will try to reconnect for a period of time and a message will appear in the run log if this is successful, or the solve will terminate with an error if it is not. If the connection between client and server is established when the connection between server and the executing worker is lost then the solve will be restarted to maintain determinism, and a computerestart callback will be fired to notify the calling application. Any work done by the disconnected remote worker, including any integer solution callbacks already fired, will be repeated.

Support for the following features are disabled when solving remotely and calling the related API methods will cause a runtime error:

- multiple solution pools,
- solution enumeration,
- callbacks not listed as supported above.

5.14.1 Authentication

Please refer to the Insight Compute Interface guide Chapter 2 and 3 for details of connecting Xpress to a remote Insight server. (https://www.fico.com/fico-xpress-optimization/docs/latest/insight5/compute/)

5.14.2 Callbacks

Callbacks are supported. When submitting a job to a remote machine, these callbacks are restricted to the `message`, `barlog`, `miplog`, `lplog`, `cutlog`, `gapnotify`, and `intsol` callbacks. Attempting to set any other callbacks will cause a runtime error. Controls can be changed in the usual way in all the supported callbacks with the exception of the message callback. Note: when solving remotely, the value of control `SLPAUTOSAVE` is always ignored.

Within the supported callbacks, calls can be made to functions that retrieve attributes and setting control values. The incumbent solution can be obtained by calling `XPRSgetsolution`. Calling any other API function will cause a runtime error, including any nonlinear and BCL API calls.

Any job that features callbacks which return hardware related attributes will use values from the remote server. For example, XPRS_CORESDETECTED will reflect the hardware on which the problem is being solved, not the hardware of the local client.

5.14.3 Licensing

When an Xpress application or an application embedding the Optimizer library is started with remote solving configured, the local license check is omitted and no local license is required to execute the application.

When a solve is started, the Optimizer instance will direct any operations that can be solved remotely to the remote server. This will also be the case if additional Optimizer instances are initiated as separate threads of the same process.

5.14.4 Advanced Configuration

There are some advanced settings that can be set using the Remote Solving Configuration file; this is described in section A.15.

CHAPTER 6

Infeasibility, Unboundedness and Instability

All users will, generally, encounter occasions in which an instance of the model they are developing is solved and found to be *infeasible* or *unbounded*. An infeasible problem is a problem that has no solution while an unbounded problem is one where the constraints do not restrict the objective function and the objective goes to infinity. Both situations often arise due to errors or shortcomings in the formulation or in the data defining the problem. When such a result is found it is typically not clear what it is about the formulation or the data that has caused the problem.

Problem instability arises when the coefficient values of the problem are such that the optimization algorithms find it difficult to converge to a solution. This is typically because of large ratios between the largest and smallest coefficients in the rows or columns and the handling of the range of numerical values in the algorithm is causing floating point accuracy issues. Problem instability generally manifests in either long run times or spurious infeasibilities.

It is often difficult to deal with these issues since it is often difficult to diagnose the cause of the problems. In this chapter we discuss the various approaches and tools provided by the Optimizer for handling these issues.

6.1 Infeasibility

A problem is said to be *infeasible* if no solution exists which satisfies all the constraints. The FICO Xpress Optimizer provides functionality for diagnosing the cause of infeasibility in the user's problem.

Before we discuss the infeasibility diagnostics of the Optimizer we will define some types of infeasibility in terms of the type of problem it relates to and how the infeasibility is detected by the Optimizer.

We will consider two basic types of infeasibility. The first we will call continuous infeasibility and the second discrete or integer infeasibility. Continuous infeasibility is where a non-MIP problem is infeasible. In this case the feasible region defined by the intersecting constraints is empty. Discrete or integer infeasibility is where a MIP problem has a feasible relaxation (a relaxation of a MIP is the problem we get when we drop the discreteness requirement on the variables) but the feasible region of the relaxation contains no solution that satisfies the discreteness requirement.

Either type of infeasibility may be detected at the presolve phase of an optimization run. Presolve is the analysis and processing of the problem before the problem is run through the optimization algorithm. If continuous infeasibility is not detected in presolve then the optimization algorithm will detect the infeasibility. If integer infeasibility is not detected in presolve, a branch and bound search will be necessary to detect the infeasibility. These scenarios are discussed in the following sections.

6.1.1 Diagnosis in Presolve

The presolve processing, if activated (see section 5.3), provides a variety of checks for infeasibility. When presolve detects infeasibility, it is possible to "trace" back the implications that determined an inconsistency and identify a particular cause. This diagnosis is carried out whenever the control parameter `TRACE` is set to `1` before the optimization routine `XPRSlpoptimize` (`LPOPTIMIZE`) is called. In such a situation, the cause of the infeasibility is then reported as part of the output from the optimization routine.

6.1.2 Diagnosis using Primal Simplex

The trace presolve functionality is typically useful when the infeasibility is simple, such that the sequence of bound implications that explains the infeasibility is short. If, however, this sequence is long or there are a number of sequences on different sets of variables, it might be useful to try forcing presolve to continue processing and then solve the problem using the primal simplex to get the, so called, 'phase 1' solution. To force presolve to continue even when an infeasibility is discovered the user can set the control `PRESOLVE` to `-1`. The 'phase 1' solution is useful because the sum of infeasibilities is minimized in the solution and the resulting set of violated constraints and violated variable bounds provides a clear picture of what aspect of the model is causing the infeasibility.

6.1.3 Irreducible Infeasible Sets

A general technique to analyze infeasibility is to find a small subset of the matrix that is itself infeasible. The Optimizer does this by finding *irreducible infeasible sets* (IISs). An IIS is a minimal set of constraints and variable bounds which is infeasible, but becomes feasible if any constraint or bound in it is removed.

A model may have several infeasibilities. Repairing a single IIS may not make the model feasible, for which reason the Optimizer can attempt to find an IIS for each of the infeasibilities in a model. The IISs found by the optimizer are independent in the sense that each constraint and variable bound may only be present in at most one IIS. In some problems there are overlapping IISs. The number of all IISs present in a problem may be exponential, and no attempt is made to enumerate all. If the infeasibility can be represented by several different IISs the Optimizer will attempt to find the IIS with the smallest number of constraints in order to make the infeasibility easier to diagnose (the Optimizer tries to minimize the number of constraints involved, even if it means that the IIS will contain more bounds).

Using the library functions IISs can be generated iteratively using the `XPRSiisfirst` and `XPRSiisnext` functions. All (a maximal set of independent) IISs can also be obtained with the `XPRSiisall` function. Note that if the problem is modified during the iterative search for IISs, the process has to be started from scratch. After a set of IISs is identified, the information contained by any one of the IISs (size, constraint and bound lists, duals, etc.) may be retrieved with the function `XPRSgetiisdata`. A summary on the generated IISs is provided by function `XPRSiisstatus`, while it is possible to save the IIS data or the IIS subproblem directly into a file in MPS or LP format using `XPRSiiswrite`. The information about the IISs is available while the problem remains unchanged. The information about an IIS may be obtained at any time after it has been generated. Function `XPRSiisclear` clears the information already stored about IISs.

On the console, all the IIS functions are available by passing different flags to the `IIS` console command. A single IIS may be found with the command `IIS`. If further IISs are required (e.g., if trying to find the smallest one) the `IIS -n` command may be used to generate subsequent IISs, or the `IIS -a` to generate all independent IISs, until no further independent IIS exists. These functions display the constraints and bounds that are identified to be in an IIS as they are found. If further information is required, the `IIS -p` *num* command may be used to retrieve all the data for a given IIS, or the `IISw` and `IISe` functions to create an LP/MPS or CSV containing the IIS subproblem or the additional information about the IIS in a file.

Once an IIS has been found it is useful to know if dropping a single constraint or bound in the IIS will completely remove the infeasibility represented by the IIS, thus an attempt is made to identify a subset of the IIS called a `sub-IIS isolation`. A `sub-IIS isolation` is a special constraint or bound in an IIS. Removing an IIS isolation constraint or bound will remove all infeasibilities in the IIS without increasing the infeasibilities outside the IIS, that is, in any other independent IISs.

The IIS isolations thus indicate the likely cause of each independent infeasibility and give an indication of which constraint or bound to drop or modify. This procedure is computationally expensive, and is carried out separately by function XPRSiisisolations (IIS-i) for an already identified IIS. It is not always possible to find IIS isolations.

After an optimal but infeasible first phase primal simplex, it is possible to identify a subproblem containing all the infeasibilities (corresponding to the given basis) to reduce the IIS work-problem dramatically. Rows with zero duals (thus with slack variables having zero reduced cost) and columns that have zero reduced costs may be excluded from the search for IISs. Moreover, for rows and columns with nonzero costs, the sign of the cost is used to relax equality rows either to less then or greater than equal rows, and to drop either possible upper or lower bounds on variables. This process is referred to as sensitivity filter for IISs.

The IIS algorithm works for all problem types, including nonconvex quadratic and fully nonlinear problems, general constraints (MIN, MAX, ABS, etc...) and discrete entities. Using the IIS algorithm on a nonlinear or nonconvex quadratic problem requires a FICO Xpress Global license.

The identification of an IIS, especially if the isolations search is also performed, may take a very long time. For this reason, using the sensitivity filter for IISs, it is possible to find only an approximation of the IISs, which typically contains all the IISs (and may contain several rows and bounds that are not part of any IIS). This approximation is a sub-problem identified at the beginning of the search for IISs, and is referred to as the initial infeasible sub-problem. Its size is typically crucial to the running time of the IIS procedure. This sub-problem is accessible by setting the input parameters of XPRSiisfirst or by calling (IIS -f) on the console. Note that the IIS approximation and the IISs generated so far are always available.

The XPRSgetiisdata function also returns dual multipliers. These multipliers are associated with Farkas' lemma of linear optimization. Farkas' lemma in its simplest form states that if $Ax = b, x \geqslant 0$ has no solution, then there exists a y for which $y^T A \geqslant 0$ and $y^T b < 0$. In other words, if the constraints and bounds are contradictory, then an inequality of form $d^T x < 0$ may be derived, where d is a constant vector of nonnegative values. The vector y, i.e., the multipliers with which the constraints and bounds have to be combined to get the contradiction is called dual multipliers. For each IIS identified, these multipliers are also provided. For an IIS all the dual multipliers should be nonzero.

The IISOPS control can be used to limit which restrictions (variable bounds, constraints, discrete entities) can be dropped during the IIS procedure. This is useful if some restrictions cannot be relaxed: either because of some physical limitation or because they are known to be correct. The deletion filter will always keep these restrictions in the problem and will try to remove only the remaining ones.

The IISLOG control can be used to decide how much information should be printed during the IIS procedure. Please refer to Appendix A.13 for a more detailed description of the IIS logging format.

6.1.4 The Infeasibility Repair Utility

In some cases, identifying the cause of infeasibility, even if the search is based on IISs may prove very demanding and time consuming. In such cases, a solution that violates the constraints and bounds minimally can greatly assist modeling. This functionality is provided by the XPRSrepairweightedinfeas function.

Based on preferences provided by the user, the Optimizer relaxes the constraints and bounds in the problem by introducing penalized deviation variables associated with selected rows and columns. Then a weighted sum of these variables (sometimes referred to as infeasibility breakers) is minimized, resulting in a solution that violates the constraints and bounds minimally regarding the provided preferences. The preference associated with a constraint or bound reflects the modeler's will to relax the corresponding right-hand side or bound. The higher the preference, the more willing the modeler is to relax (the penalty value associated is the reciprocal of the preference). A zero preference reflects that the constraint or bound cannot be relaxed. It is the responsibility of the modeler to provide preferences that yield a feasible relaxed problem. Note, that if all preferences are nonzero, the relaxed problem is always feasible (with the exception of problems containing binary or semi-continuous variables, since because of their special associated modeling properties, such variables are not relaxed).

Note, that this utility does not repair the infeasibility of the original model, but based on the preferences provided by the user, it introduces extra freedom into it to make it feasible, and minimizes the utilization of the added freedom.

The magnitude of the preferences does not affect the quality of the resulting solution, and only the ratios of the individual preferences determine the resulting solution. If a single penalty value is used for each constraint and bound group (less than and greater than or equal constraints, as well as lower and upper bounds are treated separately) the `XPRSrepairinfeas` (`REPAIRINFEAS`) function may be used, which provides a simplified interface to `XPRSrepairweightedinfeas`.

Using the new variables introduced, it is possible to warm start the primal simplex algorithm with a basic solution. However, based on the value of the control `KEEPBASIS`, the function may modify the actual basis to produce a warm start basis for the solution process. An infeasible, but first phase optimal primal solution typically speeds up the solution of the relaxed problem.

Once the optimal solution to the relaxed problem is identified (and is automatically projected back to the original problem space), it may be used by the modeler to modify the problem in order to become feasible. However, it may be of interest to know what the optimal objective value will be if the original problem is relaxed according to the solution found be the infeasibility repair function.

In order to provide such information, the infeasibility repair tool may carry out a second phase, in which the weighted violation of the constraints and bounds are restricted to be no greater than the optimum of the first phase in the infeasibility repair function, and the original objective function is minimized or maximized.

It is possible to slightly relax the restriction on the weighted violation of the constraints and bounds in the second phase by setting the value of the parameter `delta` in `XPRSrepairweightedinfeas`, or using the `-delta` option with the Console Optimizer command `REPAIRINFEAS`. If the minimal weighted violation in the first phase is `p`, a nonzero `delta` would relax the restriction on the weighted violations to be less or equal than `(1+delta)p`. While such a relaxation allows considering the effect of the original objective function in more detail, on some problems the trade-off between increasing `delta` to improve the objective can be very large, and the modeler is advised to carefully analyze the effect of the extra violations of the constraints and bounds to the underlying model.

Note, that it is possible that an infeasible problem becomes unbounded in the second phase of the infeasibility repair function. In such cases, the cause of the problem being unbounded is likely to be independent from the cause of its infeasibility.

When not all constraints and bounds are relaxed it is possible for the relaxed problem to remain infeasible. In such cases it is possible to run the IIS tool on the relaxed problem, which can be used to identify why it is still infeasible.

It is also possible to limit the amount of relaxation allowed on a per constraint side or bound by using `XPRSrepairweightedinfeasbounds`.

It can sometimes be desired to achieve an even distribution of relaxation values. This can be achieved by using quadratic penalties on the added relaxation variables, and is indicated to the optimizer by specifying a negative preference value for the constraint or bound on which a quadratic penalty should be added.

6.1.5 Integer Infeasibility

In rare cases a MIP problem can be found to be infeasible although its LP relaxation was found to be feasible. In such circumstances the feasible region for the LP relaxation, while nontrivial, contains no solutions which satisfy the various integrality constraints. These are perhaps the worst kind of infeasibilities as it can be hard to determine the cause. In such cases it is recommended that the user try to introduce some flexibility into the problem by adding slack variables to all of the constraints each with some moderate penalty cost. With the solution to this problem the user should be able to identify, from the non-zero slack variables, where the problem is being overly restricted and with this decide how to modify the formulation and/or the data to avoid the problem.

6.2 Unboundedness

A problem is said to be *unbounded* if the objective function may be improved indefinitely without violating the constraints and bounds. This can happen if a problem is being solved with the wrong optimization sense, e.g., a maximization problem is being minimized. However, when a problem is unbounded and the problem is being solved with the correct optimization sense then this indicates a problem in the formulation of the model or the data. Typically, the problem is caused by missing constraints or the wrong signs on the coefficients. Note that unboundedness is often diagnosed by presolve.

6.3 Instability

6.3.1 Scaling

When developing a model and the definition of its input data users often produce problems that contain constraints and/or columns with large ratios in the absolute values of the largest and smallest coefficients. For example:

maximize:	$10^6x + 7y$	$=$	z
subject to:	$10^6x + 0.1y$	$\leqslant$	100
	$10^7x + 8y$	$\leqslant$	500
	$10^{12}x + 10^6y$	$\leqslant$	$50*10^6$

Here the objective coefficients, constraint coefficients, and right-hand side values range between 0.1 and 10^{12}. We say that the model is *badly scaled*.

During the optimization process, the Optimizer must perform many calculations involving subtraction and division of quantities derived from the constraints and the objective function. When these calculations are carried out with values differing greatly in magnitude, the finite precision of computer arithmetic and the fixed tolerances employed by FICO Xpress result in a build up of rounding errors to a point where the Optimizer can no longer reliably find the optimal solution.

To minimize undesirable effects, when formulating your problem try to choose units (or equivalently scale your problem) so that objective coefficients and matrix elements do not range by more than 10^6, and the right-hand side and non-infinite bound values do not exceed 10^6. One common problem is the use of large finite bound values to represent infinite bounds (i.e., no bounds) — if you have to enter explicit infinite bounds, make sure you use values greater than 10^{20} which will be interpreted as infinity by the Optimizer. Avoid having large objective values that have a small relative difference — this makes it hard for the dual simplex algorithm to solve the problem. Similarly, avoid having large right-hand side or bound values that are close together, but not identical.

In the above example, both the coefficient for *x* and the last constraint might be better scaled. Issues arising from the first may be overcome by *column scaling*, effectively a change of coordinates, with the replacement of 10^6x by some new variable. Those from the second may be overcome by *row scaling*. If we set $x = 10^6x'$ and scale the last row by 10^{-6}, the example becomes the much better scaled problem:

maximize:	$x' + 7y$	$=$	z
subject to:	$x' + 0.1y$	$\leqslant$	100
	$10x' + 8y$	$\leqslant$	500
	$x' + y$	$\leqslant$	50

FICO Xpress also incorporates a number of automatic scaling options to improve the scaling of the matrix. However, the general techniques described below cannot replace attention to the choice of units specific to your problem. The best option is to scale your problem following the advice above, and use the automatic scaling provided by the Optimizer.

The form of scaling provided by the Optimizer depends on the setting of the bits of the control parameter `SCALING`. To get a particular form of scaling, set `SCALING` to the sum of the values corresponding to the scaling required. For instance, to get row scaling, column scaling and then row scaling again, set `SCALING` to 1+2+4=7. The scaling processing is applied after presolve and before the optimization algorithm. The most important of the defined bits are given in the following table. For a full list, refer to `SCALING` in Chapter 9

Bit	Value	Type of Scaling
0	1	Row scaling.
1	2	Column scaling.
2	4	Row scaling again.
3	8	Maximin.
4	16	Curtis-Reid.
5	32	`0`– scale by geometric mean; `1`– scale by maximum element (not applicable if maximin or Curtis-Reid is specified).
7	128	Objective function scaling.
8	256	Exclude the quadratic part of constraint when calculating scaling factors.
9	512	Scale the problem before presolve is applied.

If scaling is not required, `SCALING` should be set to `0`.

If the user wants to get quick results when attempting to solve a badly scaled problem it may be useful to try running customized scaling on a problem before calling the optimization algorithm. To run the scaling process on a problem the user can call the routine `XPRSscale(SCALE)`. The `SCALING` control determines how the scaling will be applied.

If the user is applying customized scaling to their problem and they are subsequently modifying the problem, it is important to note that the addition of new elements in the matrix can cause the problem to become badly scaled again. This can be avoided by reapplying their scaling strategy after completing their modifications to the matrix.

Finally, note that the scaling operations are determined by the matrix elements only. The objective coefficients, right hand side values and bound values do not influence the scaling. Only continuous variables (i.e., their bounds and coefficients) and constraints (i.e., their right-hand sides and coefficients) are scaled. Discrete entities such as integer variables are not scaled so the user should choose carefully the scaling of these variables.

6.3.2 Accuracy

The accuracy of the computed variable values and objective function value is affected in general by the various tolerances used in the Optimizer. Of particular relevance to MIP problems are the accuracy and cut off controls. The `MIPRELCUTOFF` control has a non-zero default value, which will prevent solutions very close but better than a known solution being found. This control can of course be set to zero if required.

When the LP solver stops at an optimal solution, the scaled constraints will be violated by no more than `FEASTOL` and the variables will be within `FEASTOL` of their scaled bounds. However once the constraints and variables have been unscaled the constraint and variable bound violation can increase to more than `FEASTOL`. If this happens then it indicates the problem is badly scaled. Reducing `FEASTOL` can help however this can cause the LP solve to be unstable and reduce solution performance.

However, for all problems it is probably ambitious to expect a level of accuracy in the objective of more than 1 in 1,000,000. Bear in mind that the default feasibility and optimality tolerances are 10^{-6}. It is often not practially possible to compute the solution values and reduced costs from a basis, to an accuracy better than 10^{-8} anyway, particularly for large models. It depends on the condition number of the basis matrix and the size of the right-hand side and cost coefficients. Under reasonable assumptions, an upper bound for the computed variable value accuracy is $4 x K x \|RHS\| / 10^{16}$, where $\|RHS\|$ denotes the L-infinity norm of the right-hand side and *K* is the basis condition number. The basis condition number can be found using the `XPRSbasiscondition` (`BASISCONDITION`) function.

You should also bear in mind that the matrix is scaled, which would normally have the effect of increasing the apparent feasibility tolerance.

CHAPTER 7

Reference Documentation by Topic/Functionality

This section lists all functions, controls, and attributes of the Optimizer by topics. The topics comprise problem creation, modification, and the solution process itself as well as querying the solution status and values to quickly get started with the Optimizer. Other intermediate and advanced topics include, but are not limited to Infeasibility analysis, solver customization through callback functions, or solving optimization problems with multiple objectives.

Every function, control or attribute in this section is displayed together with its main and perhaps a secondary topic label.

7.1 Barrier

Reference section for functions, controls, and attributes related to the Barrier Algorithm, which the Optimizer uses to solve linear programming problems (LPs) as standalone optimization problems or as relaxations during the Branch and Bound Search.

7.1.1 Barrier library functions

`XPRSaddcbbariteration`	Declares a barrier iteration callback function, called after each iteration during the interior point algorithm, with the ability to access the current barrier solution/slack/duals or reduced cost values, and to ask barrier to stop. [Barrier, Callback]	p. 183
`XPRSaddcbbarlog`	Declares a barrier log callback function, called at each iteration during the interior point algorithm. [Barrier, Callback]	p. 185
`XPRScrossoverlpsol`	Provides a basic optimal solution for a given solution of an LP problem. [Barrier, LP]	p. 267
`XPRSgetlastbarsol`	Used to obtain the last barrier solution values following optimization that used the barrier solver. [Barrier, Solution]	p. 329
`XPRSremovecbbariteration`	Removes a barrier iteration callback function previously added by `XPRSaddcbbariteration`. [Barrier, Callback]	p. 451

`XPRSremovecbbarlog` — Removes a Newton barrier log callback function previously added by `XPRSaddcbbarlog`.
[Barrier, Callback] p. 454

7.1.2 Barrier controls

`ALGAFTERCROSSOVER` — The algorithm to be used for the final clean up step after the crossover.
[Barrier, LP] p. 527

`BARALG` — This control determines which barrier algorithm is used to solve the problem.
[Barrier, LP, Primal Dual Hybrid Gradient] p. 531

`BARCORES` — If set to a positive integer it determines the number of physical CPU cores assumed to be present in the system by the barrier and hybrid gradient algorithms.
[Barrier, Parallel, Primal Dual Hybrid Gradient] p. 532

`BARCRASH` — Newton barrier and hybrid gradient: This determines the type of crash used for the crossover.
[Barrier, Primal Dual Hybrid Gradient] p. 532

`BARDUALSTOP` — Newton barrier and hybrid gradient: This is a convergence parameter, representing the tolerance for dual infeasibilities.
[Barrier, Primal Dual Hybrid Gradient, Tolerances] p. 532

`BARFAILITERLIMIT` — Newton barrier: The maximum number of consecutive iterations that fail to improve the solution in the barrier algorithm.
[Barrier, Limits] p. 533

`BARFREESCALE` — Defines how the barrier algorithm scales free variables.
[Barrier] p. 533

`BARGAPSTOP` — Newton barrier and hybrid gradient: This is a convergence parameter, representing the tolerance for the relative duality gap.
[Barrier, Primal Dual Hybrid Gradient] p. 533

`BARGAPTARGET` — Newton barrier: The target tolerance for the relative duality gap.
[Barrier] p. 534

`BARINDEFLIMIT` — Newton Barrier.
[Barrier, Limits] p. 535

`BARITERATIVE` — The maximum number of barrier iterations in which an iterative solver is used instead of the Cholesky decomposition.
[Barrier] p. 535

`BARITERLIMIT` — Newton barrier: The maximum number of iterations.
[Barrier, Limits] p. 536

`BARKERNEL` — Newton barrier: Defines how centrality is weighted in the barrier algorithm.
[Barrier] p. 536

`BARLARGEBOUND` — Threshold for the barrier to handle large bounds.
[Barrier] p. 537

BAROBJPERTURB	Defines how the barrier perturbs the objective. [Barrier]	p. 537
BAROBJSCALE	Defines how the barrier scales the objective. [Barrier, Numerics]	p. 537
BARORDER	Newton barrier: This controls the Cholesky factorization in the Newton-Barrier. [Barrier]	p. 538
BARORDERTHREADS	If set to a positive integer it determines the number of concurrent threads for the sparse matrix ordering algorithm in the Newton-barrier method. [Barrier, Parallel]	p. 538
BAROUTPUT	Newton barrier and hybrid gradient: This specifies the level of solution output provided. [Barrier, Logging, Primal Dual Hybrid Gradient]	p. 538
BARPERTURB	Newton barrier: In numerically challenging cases it is often advantageous to apply perturbations on the KKT system to improve its numerical properties. [Barrier, Numerics]	p. 539
BARPRESOLVEOPS	Newton barrier: This bit-vector (see Section 9.2) controls the Newton-Barrier specific presolve operations. [Barrier, Bit-vector, Presolve]	p. 539
BARPRIMALSTOP	Newton barrier and hybrid gradient: This is a convergence parameter, indicating the tolerance for primal infeasibilities. [Barrier, Primal Dual Hybrid Gradient, Tolerances]	p. 539
BARREFITER	Newton barrier: After terminating the barrier algorithm, further refinement steps can be performed. [Barrier]	p. 540
BARREGULARIZE	This bit-vector control (see Section 9.2) determines how the barrier algorithm applies regularization on the KKT system. [Barrier, Bit-vector]	p. 540
BARRHSSCALE	Defines how the barrier scales the right hand side. [Barrier, Numerics]	p. 541
BARSOLUTION	This determines whether the barrier has to decide which is the best solution found or return the solution computed by the last barrier iteration. [Barrier]	p. 541
BARSTART	Controls the computation of the starting point and warm-starting for the Newton barrier and the hybrid gradient algorithms. [Barrier, Primal Dual Hybrid Gradient]	p. 541
BARSTARTWEIGHT	Newton barrier: This sets a weight for the warm-start point when warm-start is set for the barrier algorithm. [Barrier]	p. 542
BARSTEPSTOP	Newton barrier: A convergence parameter, representing the minimal step size. [Barrier]	p. 542

BARTHREADS	If set to a positive integer it determines the number of threads implemented to run the Newton-barrier and hybrid gradient algorithms. [Barrier, Parallel, Primal Dual Hybrid Gradient]	p. 543
CHOLESKYALG	Newton barrier: type of Cholesky factorization used; bit-vector-control (see Section 9.2). [Barrier, Bit-vector]	p. 548
CHOLESKYTOL	Newton barrier: The tolerance for pivot elements in the Cholesky decomposition of the normal equations coefficient matrix, computed at each iteration of the barrier algorithm. [Barrier]	p. 549
CROSSOVER	Newton barrier and hybrid gradient: This control determines whether the barrier method will cross over to the simplex method when at optimal solution has been found, to provide an end basis (see XPRSgetbasis, XPRSwritebasis) and advanced sensitivity analysis information (see XPRSobjsa, XPRSrhssa, XPRSbndsa). [Barrier, Primal Dual Hybrid Gradient]	p. 554
CROSSOVERACCURACYTOL	Newton barrier: This control determines how crossover adjusts the default relative pivot tolerance. [Barrier, Primal Dual Hybrid Gradient]	p. 555
CROSSOVERITERLIMIT	Newton barrier and hybrid gradient: The maximum number of iterations that will be performed in the crossover procedure before the optimization process terminates. [Barrier, Primal Dual Hybrid Gradient]	p. 555
CROSSOVEROPS	Newton barrier and hybrid gradient: a bit-vector (see Section 9.2) for adjusting the behavior of the crossover procedure. [Barrier, Bit-vector, Primal Dual Hybrid Gradient]	p. 555
CROSSOVERTHREADS	Determines the maximum number of threads that parallel crossover is allowed to use. [Barrier, Parallel, Primal Dual Hybrid Gradient]	p. 556
DENSECOLLIMIT	Newton barrier: Columns with more than DENSECOLLIMIT elements are considered to be dense. [Barrier, LP]	p. 559
PREANALYTICCENTER	Determines if analytic centers should be computed and used for variable fixing and the generation of alternative reduced costs (-1: Auto 0: Off, 1: Fixing, 2: Redcost, 3: Both) [Barrier, Presolve]	p. 616

7.1.3 Barrier attributes

BARAASIZE	Number of nonzeros in AA^{T}. [Barrier]	p. 657
BARCGAP	Convergence criterion for the Newton barrier algorithm. [Barrier]	p. 658
BARCONDA	Absolute condition measure calculated in the last iteration of the barrier algorithm. [Barrier]	p. 658

BARCONDD	Condition measure calculated in the last iteration of the barrier algorithm. [Barrier]	p. 658
BARCROSSOVER	Indicates whether or not the basis crossover phase has been entered. [Barrier]	p. 658
BARDENSECOL	Number of dense columns found in the matrix. [Barrier]	p. 659
BARDUALINF	Sum of the dual infeasibilities for the Newton barrier algorithm. [Barrier]	p. 659
BARDUALOBJ	Dual objective value calculated by the Newton barrier algorithm. [Barrier]	p. 659
BARITER	Number of Newton barrier iterations. [Barrier]	p. 659
BARLSIZE	Number of nonzeros in L resulting from the Cholesky factorization. [Barrier]	p. 659
BARPRIMALINF	Sum of the primal infeasibilities for the Newton barrier algorithm. [Barrier]	p. 660
BARPRIMALOBJ	Primal objective value calculated by the Newton barrier algorithm. [Barrier]	p. 660
BARSING	Number of linearly dependent binding constraints at the optimal barrier solution. [Barrier]	p. 660
BARSINGR	Regularized number of linearly dependent binding constraints at the optimal barrier solution. [Barrier]	p. 660
CROSSOVERITER	Number of simplex iterations performed in crossover. [Barrier, LP]	p. 665

7.2 Bit-vector

Reference section for all bit-vector controls of the Optimizer. See Section 9.2 for more details on bit-vector controls.

7.2.1 Bit-vector controls

BACKGROUNDSELECT	Bit-vector control (see Section 9.2) to select which tasks to run in background jobs (for example in parallel to the root cut loop). [Bit-vector, Heuristics, Root]	p. 531
BARHGOPS	Bit-vector control (see Section 9.2) options for the hybrid gradient algorithm. [Bit-vector, Primal Dual Hybrid Gradient]	p. 535

BARPRESOLVEOPS	Newton barrier: This bit-vector (see Section 9.2) controls the Newton-Barrier specific presolve operations. [Barrier, Bit-vector, Presolve]	p. 539
BARREGULARIZE	This bit-vector control (see Section 9.2) determines how the barrier algorithm applies regularization on the KKT system. [Barrier, Bit-vector]	p. 540
CHOLESKYALG	Newton barrier: type of Cholesky factorization used; bit-vector-control (see Section 9.2). [Barrier, Bit-vector]	p. 548
CLAMPING	This bit-vector control (see Section 9.2) allows for the adjustment of returned solution values such that they are always within bounds. [Bit-vector, Numerics, Solution]	p. 549
CRASH	Simplex: This determines the type of crash used when the algorithm begins. [Bit-vector, LP, Simplex]	p. 554
CROSSOVEROPS	Newton barrier and hybrid gradient: a bit-vector (see Section 9.2) for adjusting the behavior of the crossover procedure. [Barrier, Bit-vector, Primal Dual Hybrid Gradient]	p. 555
CUTSELECT	A bit-vector (see Section 9.2) providing detailed control of the cuts created for the root node of a MIP solve. [Bit-vector, Cuts, Root]	p. 557
DUALIZEOPS	Bit-vector control (see Section 9.2) for adjusting the behavior when a problem is dualized. [Bit-vector, LP]	p. 561
DUALSTRATEGY	This bit-vector control (see Section 9.2) specifies the dual simplex strategy. [Bit-vector, LP, Simplex]	p. 561
HEURSEARCHBACKGROUNDSELECT	Bit-vector control (see Section 9.2) to select which large neighborhood searches to run in the background (for example in parallel to the root cut loop). [Bit-vector, Heuristics, Root]	p. 577
HEURSEARCHROOTSELECT	A bit-vector control (see Section 9.2) for selecting which local search heuristics to apply on the root node of a MIP solve. [Bit-vector, Heuristics, Root]	p. 578
HEURSEARCHTREESELECT	A bit-vector control (see Section 9.2) for selecting which local search heuristics to apply during the tree search of a MIP solve. [Bit-vector, Branch and Bound Search, Heuristics]	p. 579
IISOPS	Selects which part of the restrictions (bounds, constraints, entities) should always be kept in an IIS. [Bit-vector, Infeasibility]	p. 582
LPFLAGS	A bit-vector control (see Section 9.2) which defines the algorithm for solving an LP problem or the initial LP relaxation of a MIP problem. [Bit-vector, LP, Root]	p. 587
MIPPRESOLVE	Branch and Bound: Type of integer processing to be performed. [Bit-vector, Branch and Bound Search, Presolve]	p. 603

7.2.2 Bit-vector attributes

7.3 Branch and Bound Search

Reference section for functions, controls, and attributes related to the Branch and Bound Search.

7.3.1 Branch and Bound Search library functions

7.3.2 Branch and Bound Search controls

`BACKTRACKTIE`	Branch and Bound: Specifies how to break ties when selecting the next node to work on when a full backtrack is performed. [Branch and Bound Search] p. 530
`BREADTHFIRST`	The number of nodes to include in the best-first search before switching to the local first search (`NODESELECTION` = `4`). [Branch and Bound Search] p. 545
`CONFLICTCUTS`	Branch and Bound: Specifies how cautious or aggressive the optimizer should be when searching for and applying conflict cuts. [Branch and Bound Search] p. 551
`CUTDEPTH`	Branch and Bound: Sets the maximum depth in the tree search at which cuts will be generated. [Branch and Bound Search, Cuts] p. 556
`CUTFACTOR`	Limit on the number of cuts and cut coefficients the optimizer is allowed to add to the matrix during tree search. [Branch and Bound Search, Cuts] p. 557
`CUTFREQ`	Branch and Bound: This specifies the frequency at which cuts are generated in the tree search. [Branch and Bound Search, Cuts] p. 557
`CUTSTRATEGY`	Branch and Bound: This specifies the cut strategy. [Branch and Bound Search, Cuts] p. 558
`DEFAULTALG`	This selects the algorithm that will be used to solve LPs, standalone or during MIP optimization. [Branch and Bound Search, LP] p. 558
`HEURDEPTH`	Branch and Bound: Sets the maximum depth in the tree search at which heuristics will be used to find MIP solutions. [Branch and Bound Search, Heuristics] p. 572
`HEURFREQ`	Branch and Bound: This specifies the frequency at which heuristics are used in the tree search. [Branch and Bound Search, Heuristics] p. 575
`HEURSEARCHFREQ`	Branch and Bound: This specifies how often the local search heuristic should be run in the tree. [Branch and Bound Search, Heuristics] p. 577
`HEURSEARCHTREESELECT`	A bit-vector control (see Section 9.2) for selecting which local search heuristics to apply during the tree search of a MIP solve. [Bit-vector, Branch and Bound Search, Heuristics] p. 579
`LOCALCHOICE`	Controls when to perform a local backtrack between the two child nodes during a dive in the branch and bound tree. [Branch and Bound Search] p. 586
`MAXLOCALBACKTRACK`	Branch-and-Bound: How far back up the current dive path the optimizer is allowed to look for a local backtrack candidate node. [Branch and Bound Search] p. 592
`MAXNODE`	Branch and Bound: The maximum number of nodes that will be explored. [Branch and Bound Search, Limits] p. 595

`NODESELECTION`	Branch and Bound: This determines which nodes will be considered for solution once the current node has been solved. [Branch and Bound Search]	p. 612
`REFACTOR`	Indicates whether the optimization should restart using the current representation of the factorization in memory. [Branch and Bound Search, Simplex]	p. 634
`TREECOMPRESSION`	When writing nodes to the gloal file, the optimizer can try to use data-compression techniques to reduce the size of the tree file on disk. [Branch and Bound Search, File IO, Memory]	p. 645
`TREECOVERCUTS`	Branch and Bound: The number of rounds of lifted cover inequalities generated at nodes other than the root node in the tree. [Branch and Bound Search, Cuts]	p. 645
`TREECUTSELECT`	A bit-vector (see Section 9.2) providing detailed control of the cuts created during the tree search of a MIP solve. [Bit-vector, Branch and Bound Search, Cuts]	p. 646
`TREEDIAGNOSTICS`	A bit-vector (see Section 9.2) providing control over how various tree-management-related messages get printed in the tree log file during the branch-and-bound search. [Bit-vector, Branch and Bound Search, Logging]	p. 646
`TREEFILELOGINTERVAL`	This control sets the interval between progress messages output while writing tree data to the tree file, in seconds. [Branch and Bound Search, Logging]	p. 569
`TREEGOMCUTS`	Branch and Bound: The number of rounds of Gomory cuts generated at nodes other than the first node in the tree. [Branch and Bound Search, Cuts]	p. 647
`TREEMEMORYSAVINGTARGET`	When the memory used by the branch-and-bound search tree exceeds the limit specified by the `TREEMEMORYLIMIT` control, the optimizer will try to save memory by writing lower-rated sections of the tree to the tree file. [Branch and Bound Search, Memory]	p. 648
`TREEQCCUTS`	Branch and Bound: Limit on the number of rounds of outer approximation cuts generated for nodes other than the root node, when solving a mixed integer quadratic constrained or mixed integer second order conic problem with outer approximation. [Branch and Bound Search, Quadratic]	p. 648

7.3.3 Branch and Bound Search attributes

`ACTIVENODES`	Number of outstanding nodes. [Branch and Bound Search]	p. 656
`ATTENTIONLEVEL`	A measure between 0 and 1 for how numerically unstable the problem is. [Branch and Bound Search, Numerics]	p. 657
`BESTBOUND`	Value of the best bound determined so far by the MIP search. [Branch and Bound Search]	p. 661

7.4 Branching Object

Reference section for functions, controls, and attributes related to branching objects. Branching objects can be used to implement custom user branching strategies.

7.4.1 Branching Object library functions

7.5 Branching

Reference section for functions, controls, and attributes related to Branching. All of them affect how problems are subdivided to resolve infeasibilities during the Branch and Bound search.

7.5.1 Branching library functions

7.5.2 Branching controls

BRANCHCHOICE	Once a MIP entity has been selected for branching, this control determines which of the branches is solved first. [Branching] p. 544
BRANCHDISJ	Branch and Bound: Determines whether the optimizer should attempt to branch on general split disjunctions during the branch and bound search. [Branching] p. 544
BRANCHSTRUCTURAL	Branch and Bound: Determines whether the optimizer should search for special structure in the problem to branch on during the branch and bound search. [Branching] p. 545
GLOBALSPATIALBRANCHCUTTINGEFFORT	Limits the effort that is spent on creating cuts during spatial branching. [Branching, Global] p. 570
GLOBALSPATIALBRANCHIFPREFERORIG	Whether spatial branchings on original variables should be preferred over branching on auxiliary variables that were introduced by the reformulation of the global solver. [Branching, Global] p. 571
HISTORYCOSTS	Branch and Bound: How to update the pseudo cost for a MIP entity when a strong branch or a regular branch is applied. [Branching] p. 581
PSEUDOCOST	Branch and Bound: The default pseudo cost used in estimation of the degradation associated with an unexplored node in the tree search. [Branching] p. 631
SBBEST	Number of infeasible MIP entities to initialize pseudo costs for on each node. [Branching] p. 638
SBEFFORT	Adjusts the overall amount of effort when using strong branching to select an infeasible MIP entity to branch on. [Branching] p. 638
SBESTIMATE	Branch and Bound: How to calculate pseudo costs from the local node when selecting an infeasible MIP entity to branch on. [Branching] p. 639
SBITERLIMIT	Number of dual iterations to perform the strong branching for each entity. [Branching, Limits] p. 639
SBSELECT	The size of the candidate list of MIP entities for strong branching. [Branching] p. 639
VARSELECTION	Branch and Bound: This determines the formula used to calculate the estimate of each integer variable, and thus which integer variable is selected to be branched on at a given node. [Branching] p. 654

7.5.3 Branching attributes

`BRANCHVALUE`	The value of the branching variable at a node of the Branch and Bound tree. [Branching] p. 661
`BRANCHVAR`	The branching variable at a node of the Branch and Bound tree. [Branching] p. 661

7.6 Callback

Reference section for functions, controls, and attributes related to the use of callback functions within the Optimizer. Callbacks enable the user to interact with the Optimizer at all stages of the solution process. For example, use callbacks to interrupt the search when a special condition is satisfied, to query or reject certain solutions while the solution process is still running, or to make custom problem modifications.

7.6.1 Callback library functions

`XPRS_ge_addcbmsghandler`	Declares an output callback function in the global environment, called every time a line of message text is output by any data in the library. [Callback, Global Environment] p. 167
`XPRS_ge_removecbmsghandler`	Removes a message callback function previously added by `XPRS_ge_addcbmsghandler`. [Callback, Global Environment] p. 170
`XPRSaddcbafterobjective`	Declares a callback which will be called after each objective in a multi-objective problem is solved. [Callback, Multiobjective] p. 186
`XPRSaddcbbariteration`	Declares a barrier iteration callback function, called after each iteration during the interior point algorithm, with the ability to access the current barrier solution/slack/duals or reduced cost values, and to ask barrier to stop. [Barrier, Callback] p. 183
`XPRSaddcbbarlog`	Declares a barrier log callback function, called at each iteration during the interior point algorithm. [Barrier, Callback] p. 185
`XPRSaddcbbeforeobjective`	Declares a callback which will be called before each objective in a multi-objective problem is solved. [Callback, Multiobjective] p. 187
`XPRSaddcbchecktime`	Declares a callback function which is called every time the Optimizer checks if the time limit has been reached. [Callback, Limits] p. 190
`XPRSaddcbchgbranchobject`	Declares a callback function that will be called after the selection of a MIP entity to branch on. [Branching, Callback] p. 192
`XPRSaddcbcomputerestart`	Declares a callback to be called when a solve executed in compute mode needs to be restarted. [Callback, Compute Interface] p. 188

`XPRSremovecbmessage`	Removes a message callback function previously added by `XPRSaddcbmessage`. [Callback, Logging]	p. 472
`XPRSremovecbmiplog`	Removes a MIP log callback function previously added by `XPRSaddcbmiplog`. [Callback, Logging]	p. 468
`XPRSremovecbmipthread`	Removes a callback function previously added by `XPRSaddcbmipthread`. [Callback, Parallel]	p. 473
`XPRSremovecbnewnode`	Removes a new-node callback function previously added by `XPRSaddcbnewnode`. [Branch and Bound Search, Callback]	p. 474
`XPRSremovecbnodecutoff`	Removes a node-cutoff callback function previously added by `XPRSaddcbnodecutoff`. [Branch and Bound Search, Callback]	p. 475
`XPRSremovecbnodelpsolved`	Removes a node callback function previously added by `XPRSaddcbnodelpsolved`. [Branch and Bound Search, Callback]	p. 457
`XPRSremovecboptnode`	Removes a node-optimal callback function previously added by `XPRSaddcboptnode`. [Branch and Bound Search, Callback]	p. 476
`XPRSremovecbpreintsol`	Removes a pre-integer solution callback function previously added by `XPRSaddcbpreintsol`. [Callback, Solution]	p. 477
`XPRSremovecbprenode`	Removes a preprocess node callback function previously added by `XPRSaddcbprenode`. [Branch and Bound Search, Callback]	p. 478
`XPRSremovecbpresolve`	Removes a presolve callback function previously added by `XPRSaddcbpresolve`. [Callback, Presolve]	p. 453
`XPRSremovecbusersolnotify`	Removes a user solution notification callback previously added by `XPRSaddcbusersolnotify`. [Callback, Solution]	p. 479

7.6.2 Callback controls

`CALLBACKCHECKTIMEDELAY`	Minimum delay in milliseconds between two consecutive executions of the CHECKTIME callback in the same solution process [Callback]	p. 546
`CALLBACKCHECKTIMEWORKDELAY`	Minimum delay in work units between two consecutive executions of the CHECKTIME callback in the same solution process [Callback, Determinism]	p. 547
`CALLBACKFROMMASTERTHREAD`	Branch and Bound: specifies whether the MIP callbacks should only be called on the master thread. [Callback]	p. 546

`MIPABSGAPNOTIFY`
Branch and bound: if the `gapnotify` callback has been set using `XPRSaddcbgapnotify`, then this callback will be triggered during the tree search when the absolute gap reaches or passes the value you set of the `MIPRELGAPNOTIFY` control.
[Callback] p. 597

`MIPABSGAPNOTIFYBOUND`
Branch and bound: if the `gapnotify` callback has been set using `XPRSaddcbgapnotify`, then this callback will be triggered during the tree search when the best bound reaches or passes the value you set of the `MIPRELGAPNOTIFYBOUND` control.
[Callback] p. 598

`MIPABSGAPNOTIFYOBJ`
Branch and bound: if the `gapnotify` callback has been set using `XPRSaddcbgapnotify`, then this callback will be triggered during the tree search when the best solution value reaches or passes the value you set of the `MIPRELGAPNOTIFYOBJ` control.
[Callback] p. 598

`MIPRELGAPNOTIFY`
Branch and bound: if the `gapnotify` callback has been set using `XPRSaddcbgapnotify`, then this callback will be triggered during the branch and bound tree search when the relative gap reaches or passes the value you set of the `MIPRELGAPNOTIFY` control.
[Callback] p. 605

`MUTEXCALLBACKS`
Branch and Bound: This determines whether the callback routines are mutexed from within the optimizer.
[Branch and Bound Search, Callback] p. 611

`SERIALIZEPREINTSOL`
Setting `SERIALIZEPREINTSOL` to 1 will ensure that the `preintsol` callback is always fired in a deterministic order during a parallel MIP solve.
[Callback, Determinism] p. 641

7.6.3 Callback attributes

`CALLBACKCOUNT_CUTMGR`
This attribute counts the number of times the cut manager callback set by `XPRSaddcbcutmgr` has been called for the current node, including the current callback call.
[Callback, Cuts] p. 661

`CALLBACKCOUNT_OPTNODE`
This attribute counts the number of times the optimal node callback set by `XPRSaddcboptnode` has been called for the current node, including the current callback call.
[Branch and Bound Search, Callback] p. 662

7.7 Compute Interface

Reference section for functions, controls, and attributes related to the Compute Interface.

7.7.1 Compute Interface library functions

`XPRS_ge_getcomputeallowed`
Query whether the current application is allowed to use the Insight Compute interface.
[Compute Interface, Global Environment] p. 168

`XPRS_ge_setcomputeallowed` Set whether the current application is allowed to use the Insight Compute interface.
[Compute Interface, Global Environment] p. 172

`XPRSaddcbcomputerestart` Declares a callback to be called when a solve executed in compute mode needs to be restarted.
[Callback, Compute Interface] p. 188

`XPRSremovecbcomputerestart` Removes a computerestart callback function previously added by `XPRSaddcbcomputerestart`.
[Callback, Compute Interface] p. 452

7.7.2 Compute Interface controls

`COMPUTE` Controls whether the next solve is performed directly or on an Insight Compute Interface.
[Compute Interface] p. 549

`COMPUTEEXECSERVICE` Selects the Insight execution service that will be used for solving remote optimizations.
[Compute Interface] p. 550

`COMPUTEJOBPRIORITY` Selects the priority that will be used for remote optimization jobs.
[Compute Interface] p. 550

`COMPUTELOG` Controls how the run log is fetched when a solve is performed on an Insight Compute Interface.
[Compute Interface] p. 551

7.7.3 Compute Interface attributes

`COMPUTEEXECUTIONS` The number of solves executed on a compute server.
[Compute Interface] p. 663

7.8 Controls and Attributes

Reference section for functions related to setting and querying Controls and Attributes of the Optimizer. The Optimizer provides various problem and solution statistics in the form of user attributes. User controls govern the execution of the solution algorithms within the Optimizer.

7.8.1 Controls and Attributes library functions

`XPRScopycontrols` Copies controls defined for one problem to another.
[Controls and Attributes] p. 264

`XPRSdumpcontrols` Displays the list of controls and their current value for those controls that have been set to a non default value.
[Controls and Attributes] p. 279

`XPRSgetattribinfo` Accesses the id number and the type information of an attribute given its name.
[Controls and Attributes] p. 288

`XPRSgetcontrolinfo` Accesses the id number and the type information of a control given its name.
[Controls and Attributes] p. 304

7.8.2 Controls and Attributes controls

7.9 Cuts

Reference section for functions, controls, and attributes related to cutting plane separation. Separation denotes the process of deriving new valid inequalities to strengthen the linear (LP) relaxation during the Branch and Bound Search.

7.9.1 Cuts library functions

XPRSloadcuts	Loads cuts from the cut pool into the matrix. [Cuts, Data Input]	p. 393
XPRSloaddelayedrows	Specifies that a set of rows in the matrix will be treated as delayed rows during a tree search. [Cuts, Problem Input]	p. 394
XPRSloadmodelcuts	Specifies that a set of rows in the matrix will be treated as model cuts. [Cuts, Data Input]	p. 405
XPRSremovecbcutround	Removes a cut round callback function previously added by XPRSaddcbcutround. [Callback, Cuts]	p. 464
XPRSstorecuts, XPRSstorecuts64	Stores cuts into the cut pool, but does not apply them to the current node. [Cuts]	p. 507

7.9.2 Cuts controls

AUTOCUTTING	Should the Optimizer automatically decide whether to generate cutting planes at local nodes in the tree or not? If the CUTFREQ control is set, no automatic selection will be made and local cutting will be enabled. [Cuts]	p. 528
COVERCUTS	Branch and Bound: The number of rounds of lifted cover inequalities at the root node. [Cuts, Root]	p. 552
CUTDEPTH	Branch and Bound: Sets the maximum depth in the tree search at which cuts will be generated. [Branch and Bound Search, Cuts]	p. 556
CUTFACTOR	Limit on the number of cuts and cut coefficients the optimizer is allowed to add to the matrix during tree search. [Branch and Bound Search, Cuts]	p. 557
CUTFREQ	Branch and Bound: This specifies the frequency at which cuts are generated in the tree search. [Branch and Bound Search, Cuts]	p. 557
CUTSELECT	A bit-vector (see Section 9.2) providing detailed control of the cuts created for the root node of a MIP solve. [Bit-vector, Cuts, Root]	p. 557
CUTSTRATEGY	Branch and Bound: This specifies the cut strategy. [Branch and Bound Search, Cuts]	p. 558
GLOBALNLPCUTS	Limit on the number of rounds of outer approximation and convexification cuts generated for the root node, when solving an (MI)NLP to global optimality. [Cuts, Global]	p. 570
GLOBALNUMINITNLPCUTS	Specifies the maximum number of tangent cuts when setting up the initial relaxation during a global solve. [Cuts, Global]	p. 570

GLOBALTREENLPCUTS	Limit on the number of rounds of outer approximation and convexification cuts generated for each node in the tree, when solving an (MI)NLP to global optimality. [Cuts, Global]	p. 571
GOMCUTS	Branch and Bound: The number of rounds of Gomory or lift-and-project cuts at the root node. [Cuts, Root]	p. 572
LNPBEST	Number of infeasible MIP entities to create lift-and-project cuts for during each round of Gomory cuts at the root node (see GOMCUTS). [Cuts, Root]	p. 586
LNPITERLIMIT	Number of iterations to perform in improving each lift-and-project cut. [Cuts, Limits]	p. 586
MAXCUTTIME	The maximum amount of time allowed for generation of cutting planes and reoptimization. [Cuts, Limits]	p. 591
MCFCUTSTRATEGY	Level of Multi-Commodity Flow (MCF) cutting planes separation: This specifies how much aggresively MCF cuts should be separated. [Cuts]	p. 596
QCCUTS	Branch and Bound: Limit on the number of rounds of outer approximation cuts generated for the root node, when solving a mixed integer quadratic constrained or mixed integer second order conic problem with outer approximation. [Cuts, Quadratic]	p. 632
RLTCUTS	Determines whether RLT cuts should be separated in the Xpress Global Solver. [Cuts, Quadratic]	p. 637
TREECOVERCUTS	Branch and Bound: The number of rounds of lifted cover inequalities generated at nodes other than the root node in the tree. [Branch and Bound Search, Cuts]	p. 645
TREECUTSELECT	A bit-vector (see Section 9.2) providing detailed control of the cuts created during the tree search of a MIP solve. [Bit-vector, Branch and Bound Search, Cuts]	p. 646
TREEGOMCUTS	Branch and Bound: The number of rounds of Gomory cuts generated at nodes other than the first node in the tree. [Branch and Bound Search, Cuts]	p. 647

7.9.3 Cuts attributes

CALLBACKCOUNT_CUTMGR	This attribute counts the number of times the cut manager callback set by XPRSaddcbcutmgr has been called for the current node, including the current callback call. [Callback, Cuts]	p. 661

CUTROUNDS	Number of rounds of cuts applied to the current node of a branch-and-bound search. [Cuts]	p. 666
CUTS	Number of cuts being added to the matrix. [Cuts]	p. 667

7.10 Data Input

Reference section for functions, controls, and attributes related to the input of auxiliary data into the Optimizer. While not strictly necessary, auxiliary data can be used to customize the solution process.

7.10.1 Data Input library functions

XPRSloadbasis	Loads a basis from the user's areas. [Data Input, LP]	p. 391
XPRSloadbranchdirs	Loads directives into the current problem to specify which MIP entities the Optimizer should continue to branch on when a node solution is integer feasible. [Branching, Data Input]	p. 392
XPRSloadcuts	Loads cuts from the cut pool into the matrix. [Cuts, Data Input]	p. 393
XPRSloaddirs	Loads directives into the matrix. [Data Input]	p. 395
XPRSloadlpsol	Loads an LP solution for the problem into the Optimizer. [Data Input, LP, Solution]	p. 403
XPRSloadmipsol	Loads a starting MIP solution for the problem into the Optimizer. [Data Input, MIP Entities, Solution]	p. 404
XPRSloadmodelcuts	Specifies that a set of rows in the matrix will be treated as model cuts. [Cuts, Data Input]	p. 405
XPRSloadpresolvebasis	Loads a presolved basis from the user's areas. [Data Input, LP]	p. 406
XPRSloadpresolvedirs	Loads directives into the presolved matrix. [Branching, Data Input]	p. 407
XPRSloadsecurevecs	Allows the user to mark rows and columns in order to prevent the presolve removing these rows and columns from the matrix. [Data Input, Presolve]	p. 423

7.11 Determinism

Reference section for functions, controls, and attributes affecting determinism. In its default settings, the Optimizer behaves deterministically - every invocation of the solution routines for the same input data results not only in the same final solution, but also in a reproducible solution path. This is true even in parallel and independent of the system specs (CPU architecture, number of threads, Operating System) where the process is launched. The controls in this group can be used to drop this reproducibility and should hence be used with utmost care.

7.11.1 Determinism controls

7.11.2 Determinism attributes

7.12 File IO

Reference section for functions, controls, and attributes related to File IO. Use these functions to write and read data from disk into the Optimizer.

7.12.1 File IO library functions

7.12.2 File IO controls

CHECKINPUTDATA	Check input arrays for bad data. [File IO, Problem Creation]	p. 547
ESCAPENAMES	If characters illegal to an mps or lp file should be escaped to guarantee readability, and whether escaped characters should be transformed back when reading such a file. [File IO]	p. 563
FORCEOUTPUT	Certain names in the problem object may be incompatible with different file formats (such as names containing spaces for LP files). [File IO]	p. 567
INPUTTOL	The tolerance on input values elements. [File IO, Problem Creation, Tolerances]	p. 583
IOTIMEOUT	The maximum number of seconds to wait for an I/O operation before it is cancelled. [File IO, Limits]	p. 584
KEEPNROWS	How nonbinding rows should be handled by the MPS reader. [File IO, Problem Creation]	p. 585
MAXTREEFILESIZE	The maximum size, in megabytes, to which the tree file may grow, or 0 for no limit. [File IO, Limits]	p. 591
MPS18COMPATIBLE	Provides compatibility of MPS file output for older MPS readers. [File IO]	p. 608
MPSBOUNDNAME	When reading an MPS file, this control determines which entries from the BOUNDS section will be read. [File IO]	p. 609
MPSECHO	Determines whether comments in MPS matrix files are to be printed out during matrix input. [File IO]	p. 609
MPSFORMAT	Specifies the format of MPS files. [File IO]	p. 609
MPSOBJNAME	When reading an MPS file, this control determines which neutral row will be read as the objective function. [File IO]	p. 610
MPSRANGENAME	When reading an MPS file, this control determines which entries from the RANGES section will be read. [File IO]	p. 610
MPSRHSNAME	When reading an MPS file, this control determines which entries from the RHS section will be read. [File IO]	p. 610
OUTPUTMASK	Mask to restrict the row and column names written to file. [File IO]	p. 615
OUTPUTTOL	Zero tolerance on print values. [File IO, Tolerances]	p. 615

`TREECOMPRESSION`	When writing nodes to the gloal file, the optimizer can try to use data-compression techniques to reduce the size of the tree file on disk. [Branch and Bound Search, File IO, Memory]	p. 645
`TUNERMETHODFILE`	Tuner: Defines a file from which the tuner can read user-defined tuner method. [File IO, Tuner]	p. 650
`TUNEROUTPUT`	Tuner: Whether to output tuner results and logs to the file system. [File IO, Tuner]	p. 651
`TUNEROUTPUTPATH`	Tuner: Defines a root path to which the tuner writes the result file and logs. [File IO, Tuner]	p. 651

7.12.3 File IO attributes

`RANGENAME`	Active range name. [File IO]	p. 690
`RHSNAME`	Active right hand side name. [File IO]	p. 690

7.13 Global Environment

Reference section for functions, controls, and attributes related to the Global Environment, which acts mainly as a license guard around all Optimizer functionality.

7.13.1 Global Environment library functions

`XPRS_ge_addcbmsghandler`	Declares an output callback function in the global environment, called every time a line of message text is output by any data in the library. [Callback, Global Environment]	p. 167
`XPRS_ge_getcomputeallowed`	Query whether the current application is allowed to use the Insight Compute interface. [Compute Interface, Global Environment]	p. 168
`XPRS_ge_getlasterror`	Returns the last error encountered during a call to the Xpress global environment. [Global Environment]	p. 169
`XPRS_ge_removecbmsghandler`	Removes a message callback function previously added by `XPRS_ge_addcbmsghandler`. [Callback, Global Environment]	p. 170
`XPRS_ge_setarchconsistency`	Sets whether to force the same execution path on various CPU architecture extensions, in particular (pre-)AVX and AVX2. [Global Environment]	p. 171
`XPRS_ge_setcomputeallowed`	Set whether the current application is allowed to use the Insight Compute interface. [Compute Interface, Global Environment]	p. 172

7.14 Global

Reference section for functions, controls, and attributes related to the Global MINLP solver.

7.14.1 Global controls

7.14.2 Global attributes

7.15 Heuristics

Reference section for functions, controls, and attributes related to Primal Heuristics, which are auxiliary search algorithms for quickly finding improving solutions during the Branch and Bound Search.

7.15.1 Heuristics controls

`BACKGROUNDSELECT`	Bit-vector control (see Section 9.2) to select which tasks to run in background jobs (for example in parallel to the root cut loop). [Bit-vector, Heuristics, Root]	p. 531
`FEASIBILITYJUMP`	MIP: Decides if the Feasibility Jump heuristic should be run. [Heuristics]	p. 565
`FEASIBILITYPUMP`	Branch and Bound: Decides if the Feasibility Pump heuristic should be run at the root node. [Heuristics]	p. 566
`GLOBALLSHEURSTRATEGY`	When integer-feasible (for MINLP, any solution for NLP) but nonlinear-infeasible solutions are encountered within a global solve, the integer variables can be fixed and a local solver (as defined by the `LOCALSOLVER` control) can be called on the remaining continuous problem. [Global, Heuristics]	p. 569
`HEURBEFORELP`	Branch and Bound: Determines whether primal heuristics should be run before the initial LP relaxation has been solved. [Heuristics]	p. 572
`HEURDEPTH`	Branch and Bound: Sets the maximum depth in the tree search at which heuristics will be used to find MIP solutions. [Branch and Bound Search, Heuristics]	p. 572
`HEURDIVEITERLIMIT`	Branch and Bound: Simplex iteration limit for reoptimizing during the diving heuristic. [Heuristics, Limits]	p. 573
`HEURDIVERANDOMIZE`	The level of randomization to apply in the diving heuristic. [Heuristics]	p. 573
`HEURDIVESOFTROUNDING`	Branch and Bound: Enables a more cautious strategy for the diving heuristic, where it tries to push binaries and integer variables to their bounds using the objective, instead of directly fixing them. [Heuristics]	p. 573
`HEURDIVESPEEDUP`	Branch and Bound: Changes the emphasis of the diving heuristic from solution quality to diving speed. [Heuristics]	p. 574
`HEURDIVESTRATEGY`	Branch and Bound: Chooses the strategy for the diving heuristic. [Heuristics]	p. 574
`HEUREMPHASIS`	Branch and Bound: This control specifies an emphasis for the search w.r.t. primal heuristics and other procedures that affect the speed of convergence of the primal-dual gap. For problems where the goal is to achieve a small gap but not neccessarily solving them to optimality, it is recommended to set	

	HEUREMPHASIS to 1. This setting triggers many additional heuristic calls, aiming for reducing the gap at the beginning of the search, typically at the expense of an increased time for proving optimality. [Heuristics]	p. 574
HEURFORCESPECIALOBJ	Branch and Bound: This specifies whether local search heuristics without objective or with an auxiliary objective should always be used, despite the automatic selection of the Optimiezr. [Heuristics]	p. 575
HEURFREQ	Branch and Bound: This specifies the frequency at which heuristics are used in the tree search. [Branch and Bound Search, Heuristics]	p. 575
HEURMAXSOL	Branch and Bound: This specifies the maximum number of heuristic solutions that will be found in the tree search. [Heuristics]	p. 576
HEURNODES	Branch and Bound: This specifies the maximum number of nodes at which heuristics are used in the tree search. [Heuristics]	p. 576
HEURSEARCHBACKGROUNDSELECT	Bit-vector control (see Section 9.2) to select which large neighborhood searches to run in the background (for example in parallel to the root cut loop). [Bit-vector, Heuristics, Root]	p. 577
HEURSEARCHCOPYCONTROLS	Select how user-set controls should affect local search heuristics. [Controls and Attributes, Heuristics]	p. 578
HEURSEARCHEFFORT	Adjusts the overall level of the local search heuristics. [Heuristics]	p. 576
HEURSEARCHFREQ	Branch and Bound: This specifies how often the local search heuristic should be run in the tree. [Branch and Bound Search, Heuristics]	p. 577
HEURSEARCHROOTCUTFREQ	How frequently to run the local search heuristic during root cutting. [Heuristics, Root]	p. 577
HEURSEARCHROOTSELECT	A bit-vector control (see Section 9.2) for selecting which local search heuristics to apply on the root node of a MIP solve. [Bit-vector, Heuristics, Root]	p. 578
HEURSEARCHTREESELECT	A bit-vector control (see Section 9.2) for selecting which local search heuristics to apply during the tree search of a MIP solve. [Bit-vector, Branch and Bound Search, Heuristics]	p. 579
HEURSHIFTPROP	Determines whether the Shift-and-propagate primal heuristic should be executed. [Heuristics]	p. 580
HEURTHREADS	Branch and Bound: The number of threads to dedicate to running heuristics during the root solve. [Heuristics, Parallel]	p. 580
PREROOTEFFORT	Dial for the work spent during the Pre-root parallel heuristic phase. [Heuristics, Limits, Parallel]	p. 625

`PREROOTTHREADS`	Specifies an explicit number of threads that should be used for the Pre-root parallel heuristic phase. [Heuristics, Memory, Parallel]	p. 626
`PREROOTWORKLIMIT`	Set an explicit work limit in work units for the Pre-root parallel heuristic phase. [Heuristics, Limits, Parallel]	p. 626
`USERSOLHEURISTIC`	Determines how much effort to put into running a local search heuristic to find a feasible integer solution from a partial or infeasible user solution. [Heuristics]	p. 653

7.16 Infeasibility

Reference section for functions, controls, and attributes related to Infeasibility handling. Whenever a problem is infeasible, IIS routines can help to isolate information about the infeasibility. Further API functions help in repairing an infeasibility.

7.16.1 Infeasibility library functions

`XPRSgetiisdata`	Returns information for an Irreducible Infeasible Set: size, variables and constraints (row and column vectors), and conflicting sides of the variables. [Infeasibility]	p. 321
`XPRSiisall`	Performs an automated search for independent Irreducible Infeasible Sets (IIS) in an infeasible problem. [Infeasibility]	p. 381
`XPRSiisclear`	Resets the search for Irreducible Infeasible Sets (IIS). [Infeasibility]	p. 382
`XPRSiisfirst`	Initiates a search for an Irreducible Infeasible Set (IIS) in an infeasible problem. [Infeasibility]	p. 383
`XPRSiisisolations`	Performs the isolation identification procedure for an Irreducible Infeasible Set (IIS). [Infeasibility]	p. 384
`XPRSiisnext`	Continues the search for further Irreducible Infeasible Sets (IIS), or calls `XPRSiisfirst` (`IIS`) if no IIS has been identified yet. [Infeasibility]	p. 385
`XPRSiisprint`	Prints a given Irreducible Infeasible Set (IIS) in the log. [Infeasibility, Logging]	p. 386
`XPRSiisstatus`	Returns statistics on the Irreducible Infeasible Sets (IIS) found so far by `XPRSiisfirst` (`IIS`), `XPRSiisnext` (`IIS -n`) or `XPRSiisall` (`IIS -a`). [Infeasibility]	p. 387
`XPRSiiswrite`	Writes an LP/MPS/CSV file containing a given Irreducible Infeasible Set (IIS). [File IO, Infeasibility]	p. 388

7.16.2 Infeasibility console functions

7.16.3 Infeasibility controls

7.16.4 Infeasibility attributes

7.17 Licensing

Reference section for functions, controls, and attributes related to Licensing.

7.17.1 Licensing library functions

7.18 Limits

Reference section for controls and attributes related to the various limits for the solution process.

7.18.1 Limits library functions

7.18.2 Limits controls

`HEURDIVEITERLIMIT`	Branch and Bound: Simplex iteration limit for reoptimizing during the diving heuristic. [Heuristics, Limits]	p. 573
`IOTIMEOUT`	The maximum number of seconds to wait for an I/O operation before it is cancelled. [File IO, Limits]	p. 584
`LNPITERLIMIT`	Number of iterations to perform in improving each lift-and-project cut. [Cuts, Limits]	p. 586
`LPITERLIMIT`	The maximum number of iterations that will be performed by primal simplex or dual simplex before the optimization process terminates. [Limits, LP]	p. 588
`MAXCHECKSONMAXCUTTIME`	This control is intended for use where optimization runs that are terminated using the `MAXCUTTIME` control are required to be reproduced exactly. [Limits]	p. 590
`MAXCHECKSONMAXTIME`	This control is intended for use where optimization runs that are terminated using the `TIMELIMIT` (or the deprecated `MAXTIME`) control are required to be reproduced exactly. [Limits]	p. 590
`MAXCUTTIME`	The maximum amount of time allowed for generation of cutting planes and reoptimization. [Cuts, Limits]	p. 591
`MAXIIS`	This function controls the number of Irreducible Infeasible Sets to be found using the `XPRSiisall` (`IIS -a`). [Infeasibility, Limits]	p. 592
`MAXMEMORYHARD`	This control sets the maximum amount of memory in megabytes the optimizer should allocate. [Limits, Memory]	p. 593
`MAXMEMORYSOFT`	When `RESOURCESTRATEGY` is enabled, this control sets the maximum amount of memory in megabytes the optimizer targets to allocate. [Limits, Memory]	p. 593
`MAXMIPSOL`	Branch and Bound: This specifies a limit on the number of integer solutions to be found by the Optimizer. [Limits, Solution]	p. 594
`MAXMIPTASKS`	Branch-and-Bound: The maximum number of tasks to run in parallel during a MIP solve. [Determinism, Limits, Parallel]	p. 594
`MAXNODE`	Branch and Bound: The maximum number of nodes that will be explored. [Branch and Bound Search, Limits]	p. 595
`MAXSTALLTIME`	The maximum time in seconds that the Optimizer will continue to search for improving solution after finding a new incumbent. [Limits]	p. 595

MAXTREEFILESIZE	The maximum size, in megabytes, to which the tree file may grow, or 0 for no limit.	[File IO, Limits]	p. 591
MIPABSCUTOFF	Branch and Bound: If the user knows that they are interested only in values of the objective function which are better than some value, this can be assigned to MIPABSCUTOFF.	[Limits]	p. 597
MIPABSSTOP	Branch and Bound: The absolute tolerance determining whether the tree search will continue or not.	[Limits]	p. 598
MIPADDCUTOFF	Branch and Bound: The amount to add to the objective function of the best integer solution found to give the new CURRMIPCUTOFF.	[Branch and Bound Search, Limits]	p. 599
MIPCONCURRENTNODES	Sets the node limit for when a winning solve is selected when concurrent MIP solves are enabled.	[Branch and Bound Search, Limits]	p. 600
MIPREFINEITERLIMIT	This defines an effort limit expressed as simplex iterations for the MIP solution refiner.	[Limits, Solution Refinement]	p. 604
MIPRELCUTOFF	Branch and Bound: Percentage of the incumbent value to be added to the value of the objective function when an integer solution is found, to give the new value of CURRMIPCUTOFF.	[Limits]	p. 604
MIPRELSTOP	Branch and Bound: This determines when the branch and bound tree search will terminate.	[Limits]	p. 605
MIPTERMINATIONMETHOD	Branch and Bound: How a MIP solve should be stopped on early termination when there are still active tasks in the system.	[Branch and Bound Search, Limits]	p. 607
NETSTALLLIMIT	Limit the number of degenerate pivots of the network simplex algorithm, before switching to either primal or dual simplex, depending on ALGAFTERNETWORK.	[Limits, Simplex]	p. 612
PREROOTEFFORT	Dial for the work spent during the Pre-root parallel heuristic phase.	[Heuristics, Limits, Parallel]	p. 625
PREROOTWORKLIMIT	Set an explicit work limit in work units for the Pre-root parallel heuristic phase.	[Heuristics, Limits, Parallel]	p. 626
PRESOLVEMAXGROW	Limit on how much the number of non-zero coefficients is allowed to grow during presolve, specified as a ratio of the number of non-zero coefficients in the original problem.	[Limits, Presolve]	p. 627
PRESOLVEPASSES	Number of reduction rounds to be performed in presolve	[Limits, Presolve]	p. 628

`RELAXTREEMEMORYLIMIT`	When the memory used by the branch and bound search tree exceeds the target specified by the `TREEMEMORYLIMIT` control, the optimizer will try to reduce this by writing nodes to the tree file. [Limits, Memory] p. 635
`REPAIRINFEASTIMELIMIT`	Overall time limit for the repairinfeas tool [Infeasibility, Limits] p. 637
`SBITERLIMIT`	Number of dual iterations to perform the strong branching for each entity. [Branching, Limits] p. 639
`SOLTIMELIMIT`	The maximum time in seconds that the Optimizer will run a MIP solve before it terminates, given that a solution has been found. [Limits] p. 643
`TIMELIMIT`	The maximum time in seconds that the Optimizer will run before it terminates, including the problem setup time and solution time. [Limits] p. 644
`TREEMEMORYLIMIT`	A soft limit, in megabytes, for the amount of memory to use in storing the branch and bound search tree. [Limits, Memory] p. 647
`TUNERMAXTIME`	Tuner: The maximum time in seconds that the tuner will run before it terminates. [Limits, Tuner] p. 649
`WORKLIMIT`	The maximum work (measured in work units) that the Optimizer will run before it terminates. [Determinism, Limits] p. 655

7.18.3 Limits attributes

`CHECKSONMAXCUTTIME`	This attribute is used to set the value of the `MAXCHECKSONMAXCUTTIME` control. [Limits] p. 662
`CHECKSONMAXTIME`	This attribute is used to set the value of the `MAXCHECKSONMAXTIME` control. [Limits] p. 662

7.19 Linear Algebra

Reference section for functions, controls, and attributes related to Linear Algebra routines.

7.19.1 Linear Algebra library functions

`XPRSbasisstability`	Calculates various measures for the stability of the current basis, including the basis condition number. [Linear Algebra, LP] p. 239
`XPRSbtran`	Post-multiplies a (row) vector provided by the user by the inverse of the current basis. [Linear Algebra, LP] p. 241

7.19.2 Linear Algebra console functions

7.20 Logging

Reference section for functions, controls, and attributes related to Logging.

7.20.1 Logging library functions

`XPRSiisprint`	Prints a given Irreducible Infeasible Set (IIS) in the log. [Infeasibility, Logging]	p. 386
`XPRSremovecbcutlog`	Removes a cut log callback function previously added by `XPRSaddcbcutlog`. [Callback, Logging]	p. 462
`XPRSremovecblplog`	Removes a simplex log callback function previously added by `XPRSaddcblplog`. [Callback, Logging]	p. 471
`XPRSremovecbmessage`	Removes a message callback function previously added by `XPRSaddcbmessage`. [Callback, Logging]	p. 472
`XPRSremovecbmiplog`	Removes a MIP log callback function previously added by `XPRSaddcbmiplog`. [Callback, Logging]	p. 468
`XPRSsetlogfile`	This directs all Optimizer output to a log file. [File IO, Logging]	p. 497

7.20.2 Logging controls

`BAROUTPUT`	Newton barrier and hybrid gradient: This specifies the level of solution output provided. [Barrier, Logging, Primal Dual Hybrid Gradient]	p. 538
`IISLOG`	Selects how much information should be printed during the IIS procedure. [Infeasibility, Logging]	p. 581
`LPLOG`	Simplex: The frequency at which the simplex log is printed. [Logging, LP]	p. 588
`LPLOGDELAY`	Time interval between two LP log lines. [Logging, LP]	p. 588
`LPLOGSTYLE`	Simplex: The style of the simplex log. [Logging, LP]	p. 589
`MAXPAGELINES`	Number of lines between page breaks in printable output. [Logging]	p. 595
`MIPLOG`	MIP log print control. [Logging]	p. 603
`MULTIOBJLOG`	Log level for multi-objective optimization. [Logging, Multiobjective]	p. 611
`OUTPUTCONTROLS`	This control toggles the printing of all control settings at the beginning of the search. [Logging]	p. 614
`OUTPUTLOG`	This controls the level of output produced by the Optimizer during optimization. [Logging]	p. 614
`TRACE`	Display the infeasibility diagnosis during presolve. [Logging, Presolve]	p. 645

`TREEDIAGNOSTICS`	A bit-vector (see Section 9.2) providing control over how various tree-management-related messages get printed in the tree log file during the branch-and-bound search. [Bit-vector, Branch and Bound Search, Logging]	p. 646
`TREEFILELOGINTERVAL`	This control sets the interval between progress messages output while writing tree data to the tree file, in seconds. [Branch and Bound Search, Logging]	p. 569
`TUNERVERBOSE`	Tuner: whether the tuner should prints detailed information for each run. [Logging, Tuner]	p. 653

7.21 LP

Reference section for functions, controls, and attributes related to the LP algorithms of the Optimizer. Note that many of them will affect both standalone LP optimization problems, as well as the various LP relaxations that are solved during a Branch and Bound Search.

7.21.1 LP library functions

`XPRSaddcblplog`	Declares a simplex log callback function which is called after every `LPLOG` iterations of the simplex algorithm. [Callback, LP]	p. 206
`XPRSbasisstability`	Calculates various measures for the stability of the current basis, including the basis condition number. [Linear Algebra, LP]	p. 239
`XPRSbndsa`	Returns upper and lower sensitivity ranges for specified variables' lower and upper bounds. [LP, Sensitivity Analysis]	p. 240
`XPRSbtran`	Post-multiplies a (row) vector provided by the user by the inverse of the current basis. [Linear Algebra, LP]	p. 241
`XPRScrossoverlpsol`	Provides a basic optimal solution for a given solution of an LP problem. [Barrier, LP]	p. 267
`XPRSestimaterowdualranges`	Performs a dual side range sensitivity analysis, i.e. calculates estimates for the possible ranges for dual values. [LP, Sensitivity Analysis]	p. 281
`XPRSftran`	Pre-multiplies a (column) vector provided by the user by the inverse of the current matrix. [Linear Algebra, LP]	p. 287
`XPRSgetbasis`	Returns the current basis into the user's data arrays. [Linear Algebra, LP]	p. 290
`XPRSgetbasisval`	Returns the current basis status for a specific column or row. [Linear Algebra, LP]	p. 291
`XPRSgetdualray`	Retrieves a dual ray (dual unbounded direction) for the current problem, if the problem is found to be infeasible. [LP, Solution]	p. 315

`XPRSlpoptimize` — This function begins a search for the optimal continuous (LP) solution.
[LP, Solution Process] p. 424

`XPRSobjsa` — Returns upper and lower sensitivity ranges for specified objective function coefficients.
[LP, Sensitivity Analysis] p. 429

`XPRSreadbasis` — Instructs the Optimizer to read in a previously saved basis from a file.
[File IO, LP] p. 443

`XPRSrhssa` — Returns upper and lower sensitivity ranges for specified right hand side (RHS) function coefficients.
[LP, Sensitivity Analysis] p. 488

`XPRSwritebasis` — Writes the current basis to a file for later input into the Optimizer.
[File IO, LP] p. 518

7.21.2 LP controls

`ALGAFTERCROSSOVER` — The algorithm to be used for the final clean up step after the crossover.
[Barrier, LP] p. 527

`ALGAFTERNETWORK` — The algorithm to be used for the clean up step after the network simplex solver.
[LP, Simplex] p. 527

`AUTOPERTURB` — Simplex: This indicates whether automatic perturbation is performed.
[LP] p. 529

`BARALG` — This control determines which barrier algorithm is used to solve the problem.
[Barrier, LP, Primal Dual Hybrid Gradient] p. 531

`BIGM` — The infeasibility penalty used if the ”Big M” method is implemented.
[LP, Numerics] p. 543

`BIGMMETHOD` — Simplex: This specifies whether to use the ”Big M” method, or the standard phase I (achieving feasibility) and phase II (achieving optimality).
[LP, Numerics] p. 543

`CONCURRENTTHREADS` — Determines the number of threads used by the concurrent solver.
[LP, Parallel] p. 551

`CRASH` — Simplex: This determines the type of crash used when the algorithm begins.
[Bit-vector, LP, Simplex] p. 554

`DEFAULTALG` — This selects the algorithm that will be used to solve LPs, standalone or during MIP optimization.
[Branch and Bound Search, LP] p. 558

`DENSECOLLIMIT` — Newton barrier: Columns with more than `DENSECOLLIMIT` elements are considered to be dense.
[Barrier, LP] p. 559

DUALGRADIENT	Simplex: This specifies the dual simplex pricing method. [LP, Simplex]	p. 560
DUALIZE	For a linear problem or the initial linear relaxation of a MIP, determines whether to form and solve the dual problem. [LP, Problem Transformation]	p. 560
DUALIZEOPS	Bit-vector control (see Section 9.2) for adjusting the behavior when a problem is dualized. [Bit-vector, LP]	p. 561
DUALPERTURB	The factor by which the problem will be perturbed prior to optimization by dual simplex. [LP, Simplex]	p. 561
DUALSTRATEGY	This bit-vector control (see Section 9.2) specifies the dual simplex strategy. [Bit-vector, LP, Simplex]	p. 561
DUALTHREADS	Determines the maximum number of threads that dual simplex is allowed to use. [LP, Parallel]	p. 562
FORCEPARALLELDUAL	Dual simplex: specifies whether the dual simplex solver should always use the parallel simplex algorithm. [LP, Simplex]	p. 568
KEEPBASIS	Simplex: This determines whether the basis should be kept when reoptimizing a problem. [LP, Simplex]	p. 585
LPFLAGS	A bit-vector control (see Section 9.2) which defines the algorithm for solving an LP problem or the initial LP relaxation of a MIP problem. [Bit-vector, LP, Root]	p. 587
LPFOLDING	Simplex and barrier: whether to fold an LP problem before solving it. [LP]	p. 587
LPITERLIMIT	The maximum number of iterations that will be performed by primal simplex or dual simplex before the optimization process terminates. [Limits, LP]	p. 588
LPLOG	Simplex: The frequency at which the simplex log is printed. [Logging, LP]	p. 588
LPLOGDELAY	Time interval between two LP log lines. [Logging, LP]	p. 588
LPLOGSTYLE	Simplex: The style of the simplex log. [Logging, LP]	p. 589
LPREFINEITERLIMIT	This specifies the simplex iteration limit the solution refiner can spend in attempting to increase the accuracy of an LP solution. [LP, Solution Refinement]	p. 589
MARKOWITZTOL	The Markowitz tolerance used for the factorization of the basis matrix. [LP, Simplex]	p. 589

7.21.3 LP attributes

`MAXRELDUALINFEAS`	Maximum calculated relative dual infeasibility in the unscaled original problem. [LP, Numerics]	p. 673
`MAXRELPRIMALINFEAS`	Maximum calculated relative primal infeasibility in the unscaled original problem. [LP, Numerics]	p. 673
`PRIMALINFEAS`	Number of primal infeasibilities. [LP]	p. 688
`SIMPLEXITER`	Number of simplex iterations performed. [LP, Simplex]	p. 691
`SUMPRIMALINF`	Scaled sum of primal infeasibilities. [LP, Numerics]	p. 695

7.22 Memory

Reference section for functions, controls, and attributes related to memory handling and usage.

7.22.1 Memory controls

`MAXMEMORYHARD`	This control sets the maximum amount of memory in megabytes the optimizer should allocate. [Limits, Memory]	p. 593
`MAXMEMORYSOFT`	When `RESOURCESTRATEGY` is enabled, this control sets the maximum amount of memory in megabytes the optimizer targets to allocate. [Limits, Memory]	p. 593
`PREROOTTHREADS`	Specifies an explicit number of threads that should be used for the Pre-root parallel heuristic phase. [Heuristics, Memory, Parallel]	p. 626
`RELAXTREEMEMORYLIMIT`	When the memory used by the branch and bound search tree exceeds the target specified by the `TREEMEMORYLIMIT` control, the optimizer will try to reduce this by writing nodes to the tree file. [Limits, Memory]	p. 635
`RESOURCESTRATEGY`	Controls whether the optimizer is allowed to make nondeterministic decisions if memory is running low in an effort to preserve memory and finish the solve. [Determinism, Memory]	p. 637
`TREECOMPRESSION`	When writing nodes to the gloal file, the optimizer can try to use data-compression techniques to reduce the size of the tree file on disk. [Branch and Bound Search, File IO, Memory]	p. 645
`TREEMEMORYLIMIT`	A soft limit, in megabytes, for the amount of memory to use in storing the branch and bound search tree. [Limits, Memory]	p. 647

TREEMEMORYSAVINGTARGET	When the memory used by the branch-and-bound search tree exceeds the limit specified by the TREEMEMORYLIMIT control, the optimizer will try to save memory by writing lower-rated sections of the tree to the tree file. [Branch and Bound Search, Memory]	p. 648

7.22.2 Memory attributes

AVAILABLEMEMORY	The amount of heap memory detected by Xpress as free. [Memory]	p. 657
CURRENTMEMORY	The amount of dynamically allocated heap memory by the problem being solved. [Memory]	p. 666
MEMORYLIMITDETECTED	The detected amount of memory accessible to the solver process, in megabytes. [Memory]	p. 674
PEAKMEMORY	An estimate of the peak amount of dynamically allocated heap memory by the problem. [Memory]	p. 685
PEAKTOTALTREEMEMORYUSAGE	The peak size, in megabytes, that the branch-and-bound search tree reached during the solve. [Branch and Bound Search, Memory]	p. 685
SYSTEMMEMORY	The amount of non problem specific memory used by the solver. [Memory, System]	p. 695
TOTALMEMORY	The amount of dynamically allocated heap memory by the optimizer, including all problems currently exsisting. [Memory]	p. 696
TREEFILESIZE	The allocated size of the tree file, in megabytes. [Branch and Bound Search, Memory]	p. 696
TREEFILEUSAGE	The number of megabytes of data from the branch-and-bound tree that have been saved to the tree file. [Branch and Bound Search, Memory]	p. 697
TREEMEMORYUSAGE	The amount of physical memory, in megabytes, currently being used to store the branch-and-bound search tree. [Branch and Bound Search, Memory]	p. 697

7.23 MIP Entities

Reference section for functions, controls, and attributes related to MIP Entities, such as binary and integer variables. Further MIP entity column types comprise semi-continuous, semi-integer, partial integer variables. Special Ordered Sets of type I and II are also MIP entities.

7.23.1 MIP Entities library functions

XPRSaddsets, XPRSaddsets64	Allows sets to be added to the problem after passing it to the Optimizer using the input routines. [MIP Entities, Problem Creation]	p. 235

XPRSchgcoltype	Used to change the type of a specified set of columns in the matrix. [MIP Entities, Problem Modification]	p. 250
XPRSchgglblimit	Used to change semi-continuous or semi-integer lower bounds, or upper limits on partial integers. [MIP Entities, Problem Modification]	p. 251
XPRSdelindicators	Delete indicator constraints. [MIP Entities, Problem Creation]	p. 272
XPRSdelsets	Delete sets from a problem. [MIP Entities, Problem Creation]	p. 277
XPRSfixmipentities	Fixes all the MIP entities to the values of the last found MIP solution. [MIP Entities, Problem Modification]	p. 284
XPRSgetcoltype	Returns the column types for the columns in a given range. [MIP Entities, Problem Information]	p. 303
XPRSgetmipentities, XPRSgetmipentities64	Retrieves integr and entity information about a problem. [MIP Entities, Problem Information]	p. 319
XPRSloadmip, XPRSloadmip64	Used to load a MIP problem into the Optimizer data structures. [MIP Entities, Problem Input]	p. 397
XPRSloadmipsol	Loads a starting MIP solution for the problem into the Optimizer. [Data Input, MIP Entities, Solution]	p. 404
XPRSmipoptimize	This function begins a tree search for the optimal MIP solution. [MIP Entities, Solution Process]	p. 428

7.23.2 MIP Entities attributes

INDICATORS	Number of indicator constrains in the problem. [MIP Entities, Problem Information]	p. 670
MAXMIPINFEAS	Maximum integer fractionality in the solution. [MIP Entities, Numerics]	p. 673
MIPENTS	Number of MIP entities (i.e. binary, integer, semi-continuous, partial integer, and semi-continuous integer variables) but excluding the number of special ordered sets. [MIP Entities, Problem Information]	p. 674
MIPINFEAS	Number of integer infeasibilities, including violations of special ordered sets, at the current node. [Branch and Bound Search, MIP Entities]	p. 675
ORIGINALINDICATORS	Number of indicator constraints in the original matrix before presolving. [MIP Entities, Problem Information]	p. 682
ORIGINALMIPENTS	Number of MIP entities (i.e. binary, integer, semi-continuous, partial integer, and semi-continuous integer variables) but excluding the number of special ordered sets in the original matrix before presolving. [MIP Entities, Problem Information]	p. 682

7.24 Misc

Reference section for further miscellaneous functionality of the Optimizer.

7.24.1 Misc library functions

7.24.2 Misc console functions

7.24.3 Misc controls

7.24.4 Misc attributes

7.25 Multiobjective

Reference section for functions, controls, and attributes for solving multiobjective problems.

7.25.1 Multiobjective library functions

`XPRSsetobjdblcontrol`	Sets the value of a given double control parameter associated with an objective. [Controls and Attributes, Multiobjective]	p. 499
`XPRSsetobjintcontrol`	Sets the value of a given integer control parameter associated with an objective. [Controls and Attributes, Multiobjective]	p. 500

7.25.2 Multiobjective controls

`MULTIOBJLOG`	Log level for multi-objective optimization. [Logging, Multiobjective]	p. 611
`MULTIOBJOPS`	Modifies the behaviour of the optimizer when solving multi-objective problems. [Bit-vector, Multiobjective]	p. 610

7.25.3 Multiobjective attributes

`OBJECTIVES`	Number of objectives in the problem. [Multiobjective]	p. 677
`OBJSTOSOLVE`	Number of objectives that will be solved during the current multi-objective solve. [Multiobjective]	p. 679
`SOLVEDOBJS`	Number of objectives that have been solved so far during a multi-objective solve. [Multiobjective]	p. 692

7.26 Names Manager

Reference section for functions, controls, and attributes related to the Names Manager.

7.26.1 Names Manager library functions

`XPRSgetnamelist`	Returns the names for the rows, columns, sets, piecewise linear constraints, general constraints or objectives in a given range. [Names Manager]	p. 344

7.27 Numerics

Reference section for functions, controls, and attributes related to Numerics.

7.27.1 Numerics library functions

`XPRSgetscale`	Returns the the current scaling of the matrix. [Numerics]	p. 369
`XPRSgetscaledinfeas`	Returns a list of scaled infeasible primal and dual variables for the original problem. [LP, Numerics]	p. 370
`XPRSscale`	Re-scales the current matrix. [Numerics]	p. 490

7.27.2 Numerics controls

AUTOSCALING	Whether the Optimizer should automatically select between different scaling algorithms. [Numerics, Problem Transformation]	p. 528
BAROBJSCALE	Defines how the barrier scales the objective. [Barrier, Numerics]	p. 537
BARPERTURB	Newton barrier: In numerically challenging cases it is often advantageous to apply perturbations on the KKT system to improve its numerical properties. [Barrier, Numerics]	p. 539
BARRHSSCALE	Defines how the barrier scales the right hand side. [Barrier, Numerics]	p. 541
BIGM	The infeasibility penalty used if the "Big M" method is implemented. [LP, Numerics]	p. 543
BIGMMETHOD	Simplex: This specifies whether to use the "Big M" method, or the standard phase I (achieving feasibility) and phase II (achieving optimality). [LP, Numerics]	p. 543
CLAMPING	This bit-vector control (see Section 9.2) allows for the adjustment of returned solution values such that they are always within bounds. [Bit-vector, Numerics, Solution]	p. 549
MAXIMPLIEDBOUND	Presolve: When tighter bounds are calculated during MIP preprocessing, only bounds whose absolute value are smaller than MAXIMPLIEDBOUND will be applied to the problem. [Numerics, Presolve]	p. 592
MAXSCALEFACTOR	This determines the maximum scaling factor that can be applied during scaling. [Numerics, Presolve]	p. 595
MIPKAPPAFREQ	Branch and Bound: Specifies how frequently the basis condition number (also known as kappa) should be calculated during the branch-and-bound search. [Branch and Bound Search, Numerics]	p. 602
NUMERICALEMPHASIS	How much emphasis to place on numerical stability instead of solve speed. [Numerics]	p. 613
OBJSCALEFACTOR	Custom objective scaling factor, expressed as a power of 2. [Numerics]	p. 613
SCALING	This bit-vector control (see Section 9.2) determines how the Optimizer will rescale a model internally before optimization. [Bit-vector, Numerics, Problem Transformation]	p. 640

7.27.3 Numerics attributes

`ATTENTIONLEVEL`	A measure between 0 and 1 for how numerically unstable the problem is. [Branch and Bound Search, Numerics]	p. 657
`MAXABSDUALINFEAS`	Maximum calculated absolute dual infeasibility in the unscaled original problem. [LP, Numerics]	p. 672
`MAXABSPRIMALINFEAS`	Maximum calculated absolute primal infeasibility in the unscaled original problem. [LP, Numerics]	p. 672
`MAXKAPPA`	Largest basis condition number (also known as kappa) calculated through all nodes sampled by `MIPKAPPAFREQ`. [Branch and Bound Search, Numerics]	p. 672
`MAXMIPINFEAS`	Maximum integer fractionality in the solution. [MIP Entities, Numerics]	p. 673
`MAXRELDUALINFEAS`	Maximum calculated relative dual infeasibility in the unscaled original problem. [LP, Numerics]	p. 673
`MAXRELPRIMALINFEAS`	Maximum calculated relative primal infeasibility in the unscaled original problem. [LP, Numerics]	p. 673
`PREDICTEDATTLEVEL`	A measure between 0 and 1 to predict how numerically unstable the current MIP solve can be expected to be. [Numerics]	p. 686
`SUMPRIMALINF`	Scaled sum of primal infeasibilities. [LP, Numerics]	p. 695

7.28 Parallel

Reference section for functionality how the Optimizer takes advantage of modern multi-core CPUs. By default, the Optimizer will detect how many cores are available in the system and try to use all of them. The controls in this section affect to which extent the Optimizer uses the parallel hardware.

7.28.1 Parallel library functions

`XPRSaddcbdestroymt`	Declares a destroy MIP thread callback function, called every time a MIP thread is destroyed by the parallel MIP code. [Callback, Parallel]	p. 198
`XPRSaddcbmipthread`	Declares a MIP thread callback function, called every time a MIP worker problem is created by the parallel MIP code. [Callback, Parallel]	p. 209
`XPRSremovecbdestroymt`	Removes a slave thread destruction callback function previously added by `XPRSaddcbdestroymt`. [Callback, Parallel]	p. 465
`XPRSremovecbmipthread`	Removes a callback function previously added by `XPRSaddcbmipthread`. [Callback, Parallel]	p. 473

7.28.2 Parallel controls

7.29 Piecewise Linear and General Constraints

Reference section for functions, controls, and attributes related to Piecewise Linear and General Constraints.

7.29.1 Piecewise Linear and General Constraints library functions

`XPRSaddgencons, XPRSaddgencons64` Adds one or more general constraints to the problem.
[Piecewise Linear and General Constraints, Problem Creation]
p. 223

`XPRSaddpwlcons, XPRSaddpwlcons64` Adds one or more piecewise linear constraints to the problem.
[Piecewise Linear and General Constraints, Problem Creation]
p. 230

`XPRSdelgencons` Delete general constraints from a problem.
[Piecewise Linear and General Constraints, Problem Creation]
p. 271

`XPRSdelpwlcons` Delete piecewise linear constraints from a problem.
[Piecewise Linear and General Constraints, Problem Creation]
p. 274

`XPRSgetgencons, XPRSgetgencons64` Returns the general constraints `y = f(x1, ..., xn, c1, ..., cm)` in a given range.
[Piecewise Linear and General Constraints, Problem Information]
p. 316

`XPRSgetpwlcons, XPRSgetpwlcons64` Returns the piecewise linear constraints `y = f(x)` in a given range.
[Piecewise Linear and General Constraints, Problem Information]
p. 357

7.29.2 Piecewise Linear and General Constraints controls

`GENCONSABSTRANSFORMATION` This control specifies the reformulation method for absolute value general constraints at the beginning of the search.
[Piecewise Linear and General Constraints, Presolve] p. 568

`GENCONSDUALREDUCTIONS` This parameter specifies whether dual reductions should be applied to reduce the number of columns and rows added when transforming general constraints to MIP structs.
[Piecewise Linear and General Constraints, Presolve] p. 568

`PWLDUALREDUCTIONS` This parameter specifies whether dual reductions should be applied to reduce the number of columns, rows and SOS-constraints added when transforming piecewise linear objectives and constraints to MIP structs.
[Piecewise Linear and General Constraints, Presolve, Problem Transformation] p. 631

`PWLNONCONVEXTRANSFORMATION` This control specifies the reformulation method for piecewise linear constraints at the beginning of the search.
[Piecewise Linear and General Constraints, Presolve, Problem Transformation] p. 631

7.29.3 Piecewise Linear and General Constraints attributes

`GENCONCOLS`	Number of input variables in general constraints (i.e. MIN/MAX/AND/OR/ABS constraints) in the problem. [Piecewise Linear and General Constraints, Problem Information] p. 668
`GENCONS`	The number of general constraints (i.e. MIN/MAX/AND/OR/ABS constraints) in the problem. [Piecewise Linear and General Constraints, Problem Information] p. 668
`GENCONVALS`	Number of constant values in general constraints (MIN/MAX constraints) in the problem. [Piecewise Linear and General Constraints, Problem Information] p. 668
`ORIGINALGENCONCOLS`	Number of input variables in general constraints in the original problem before presolving. [Piecewise Linear and General Constraints, Problem Information] p. 681
`ORIGINALGENCONS`	Number of general constraints in the original problem before presolving. [Piecewise Linear and General Constraints, Problem Information] p. 681
`ORIGINALGENCONVALS`	Number of constant values in general constraints in the original problem before presolving. [Piecewise Linear and General Constraints, Problem Information] p. 681
`ORIGINALPWLPOINTS`	Number of breakpoints of piecewise linear constraints in the original problem before presolving. [Piecewise Linear and General Constraints, Problem Information] p. 683
`ORIGINALPWLS`	Number of piecewise linear constraints in the original problem before presolving. [Piecewise Linear and General Constraints, Problem Information] p. 682
`PWLCONS`	Number of piecewise linear constraints in the problem. [Piecewise Linear and General Constraints, Problem Information] p. 688
`PWLPOINTS`	Number of breakpoints of piecewise linear constraints in the problem. [Piecewise Linear and General Constraints, Problem Information] p. 688

7.30 Presolve

Reference section for functions, controls, and attributes related to Presolve. Presolve is a collection of techniques to transform the input problem into an equivalent, but smaller problem by fixing or eliminating columns and rows.

7.30.1 Presolve library functions

`XPRSaddcbpresolve`	Declares a callback to be called after presolve has been performed. [Callback, Presolve]	p. 189
`XPRSgetpresolvemap`	Returns the mapping of the row and column numbers from the presolve problem back to the original problem. [Presolve, Problem Information]	p. 353
`XPRSgetpresolvesol`	Returns the solution for the presolved problem from memory. [Presolve, Solution]	p. 354
`XPRSloadsecurevecs`	Allows the user to mark rows and columns in order to prevent the presolve removing these rows and columns from the matrix. [Data Input, Presolve]	p. 423
`XPRSpostsolve`	Postsolve the current matrix when it is in a presolved state. [Presolve, Solution Process]	p. 437
`XPRSpostsolvesol`	Postsolves a primal solution formulated in the presolved space into the corresponding solution formulated in the original space. [Presolve, Solution]	p. 438
`XPRSpresolverow`	Presolves a row formulated in terms of the original variables such that it can be added to a presolved matrix. [Misc, Presolve]	p. 439
`XPRSremovecbpresolve`	Removes a presolve callback function previously added by `XPRSaddcbpresolve`. [Callback, Presolve]	p. 453

7.30.2 Presolve controls

`BARPRESOLVEOPS`	Newton barrier: This bit-vector (see Section 9.2) controls the Newton-Barrier specific presolve operations. [Barrier, Bit-vector, Presolve]	p. 539
`ELIMFILLIN`	Amount of fill-in allowed when performing an elimination in presolve . [Presolve]	p. 562
`ELIMTOL`	The Markowitz tolerance for the elimination phase of the presolve. [Presolve, Tolerances]	p. 563
`GENCONSABSTRANSFORMATION`	This control specifies the reformulation method for absolute value general constraints at the beginning of the search. [Piecewise Linear and General Constraints, Presolve]	p. 568
`GENCONSDUALREDUCTIONS`	This parameter specifies whether dual reductions should be applied to reduce the number of columns and rows added when transforming general constraints to MIP structs. [Piecewise Linear and General Constraints, Presolve]	p. 568
`IFCHECKCONVEXITY`	Determines if the convexity of the problem is checked before optimization. [Presolve, Quadratic]	p. 581

INDLINBIGM	During presolve, indicator constraints will be linearized using a BigM coefficient whenever that BigM coefficient is small enough. [Presolve, Problem Transformation]	p. 582
INDPRELINBIGM	During presolve, indicator constraints will be linearized using a BigM coefficient whenever that BigM coefficient is small enough. [Presolve, Problem Transformation]	p. 583
MAXIMPLIEDBOUND	Presolve: When tighter bounds are calculated during MIP preprocessing, only bounds whose absolute value are smaller than MAXIMPLIEDBOUND will be applied to the problem. [Numerics, Presolve]	p. 592
MAXSCALEFACTOR	This determines the maximum scaling factor that can be applied during scaling. [Numerics, Presolve]	p. 595
MIPDUALREDUCTIONS	Branch and Bound: Limits operations that can reduce the MIP solution space. [Presolve]	p. 601
MIPPRESOLVE	Branch and Bound: Type of integer processing to be performed. [Bit-vector, Branch and Bound Search, Presolve]	p. 603
NODEPROBINGEFFORT	Adjusts the overall level of node probing. [Branch and Bound Search, Presolve]	p. 612
PENALTY	Minimum absolute penalty variable coefficient. [Presolve]	p. 615
PREANALYTICCENTER	Determines if analytic centers should be computed and used for variable fixing and the generation of alternative reduced costs (-1: Auto 0: Off, 1: Fixing, 2: Redcost, 3: Both) [Barrier, Presolve]	p. 616
PREBASISRED	Determines if a lattice basis reduction algorithm should be attempted as part of presolve [Presolve]	p. 617
PREBNDREDCONE	Determines if second order cone constraints should be used for inferring bound reductions on variables when solving a MIP. [Presolve]	p. 617
PREBNDREDQUAD	Determines if convex quadratic constraints should be used for inferring bound reductions on variables when solving a MIP. [Presolve, Quadratic]	p. 617
PRECLIQUESTRATEGY	Determines how much effort to spend on clique covers in presolve. [Presolve]	p. 618
PRECOEFELIM	Presolve: Specifies whether the optimizer should attempt to recombine constraints in order to reduce the number of non zero coefficients when presolving a mixed integer problem. [Presolve]	p. 618
PRECOMPONENTS	Presolve: determines whether small independent components should be detected and solved as individual subproblems during root node processing. [Presolve]	p. 618

PREPERMUTESEED	This control sets the seed for the pseudo-random number generator for permuting the problem when starting the presolve. [Presolve, Problem Transformation]	p. 624
PREPROBING	Presolve: Amount of probing to perform on binary variables during presolve. [Presolve]	p. 624
PREPROTECTDUAL	Presolve: specifies whether the presolver should protect a given dual solution by maintaining the same level of dual feasibility. [LP, Presolve]	p. 625
PRESOLVE	This control determines whether presolving should be performed prior to starting the main algorithm. [Presolve]	p. 627
PRESOLVEMAXGROW	Limit on how much the number of non-zero coefficients is allowed to grow during presolve, specified as a ratio of the number of non-zero coefficients in the original problem. [Limits, Presolve]	p. 627
PRESOLVEOPS	This bit-vector control (see Section 9.2) specifies the operations which are performed during the presolve. [Bit-vector, Presolve]	p. 628
PRESOLVEPASSES	Number of reduction rounds to be performed in presolve [Limits, Presolve]	p. 628
PRESORT	This bit-vector control (see Section 9.2) specifies whether to sort rows, columns and MIP entities by their names when starting the presolve. [Bit-vector, Presolve, Problem Transformation]	p. 629
PWLDUALREDUCTIONS	This parameter specifies whether dual reductions should be applied to reduce the number of columns, rows and SOS-constraints added when transforming piecewise linear objectives and constraints to MIP structs. [Piecewise Linear and General Constraints, Presolve, Problem Transformation]	p. 631
PWLNONCONVEXTRANSFORMATION	This control specifies the reformulation method for piecewise linear constraints at the beginning of the search. [Piecewise Linear and General Constraints, Presolve, Problem Transformation]	p. 631
SYMMETRY	Adjusts the overall amount of effort for symmetry detection. [Presolve]	p. 643
SYMSELECT	Adjusts the overall amount of effort for symmetry detection. [Presolve]	p. 644
TRACE	Display the infeasibility diagnosis during presolve. [Logging, Presolve]	p. 645

7.30.3 Presolve attributes

PRESOLVEINDEX	Presolve: The row or column index on which presolve detected a problem to be infeasible or unbounded. [Presolve]	p. 687

7.31 Primal Dual Hybrid Gradient

Reference section for functions, controls, and attributes related to the Primal Dual Hybrid Gradient solver, an iterative solver for very large LP relaxations.

7.31.1 Primal Dual Hybrid Gradient controls

`CROSSOVER`
Newton barrier and hybrid gradient: This control determines whether the barrier method will cross over to the simplex method when at optimal solution has been found, to provide an end basis (see `XPRSgetbasis`, `XPRSwritebasis`) and advanced sensitivity analysis information (see `XPRSobjsa`, `XPRSrhssa`, `XPRSbndsa`).
[Barrier, Primal Dual Hybrid Gradient] p. 554

`CROSSOVERACCURACYTOL`
Newton barrier: This control determines how crossover adjusts the default relative pivot tolerance.
[Barrier, Primal Dual Hybrid Gradient] p. 555

`CROSSOVERITERLIMIT`
Newton barrier and hybrid gradient: The maximum number of iterations that will be performed in the crossover procedure before the optimization process terminates.
[Barrier, Primal Dual Hybrid Gradient] p. 555

`CROSSOVEROPS`
Newton barrier and hybrid gradient: a bit-vector (see Section 9.2) for adjusting the behavior of the crossover procedure.
[Barrier, Bit-vector, Primal Dual Hybrid Gradient] p. 555

`CROSSOVERTHREADS`
Determines the maximum number of threads that parallel crossover is allowed to use.
[Barrier, Parallel, Primal Dual Hybrid Gradient] p. 556

7.32 Problem Creation

Reference section for functions, controls, and attributes related to problem creation.

7.32.1 Problem Creation library functions

`XPRSaddcols, XPRSaddcols64`
Adds columns to the optimizer matrix.
[Problem Creation] p. 219

`XPRSaddgencons, XPRSaddgencons64`
Adds one or more general constraints to the problem.
[Piecewise Linear and General Constraints, Problem Creation] p. 223

`XPRSaddnames`
When a model is loaded, the rows, columns, sets, piecewise linear and general constraints of the model may not have names associated with them.
[Problem Creation] p. 229

`XPRSaddobj`
Appends an objective function with the given coefficients to a multi-objective problem.
[Multiobjective, Problem Creation] p. 228

`XPRSaddpwlcons, XPRSaddpwlcons64`
Adds one or more piecewise linear constraints to the problem.
[Piecewise Linear and General Constraints, Problem Creation] p. 230

`XPRSaddqmatrix, XPRSaddqmatrix64`
Adds a new quadratic matrix into a row defined by triplets.
[Problem Creation, Quadratic] p. 232

`XPRSaddrows, XPRSaddrows64`
Adds rows to the optimizer matrix.
[Problem Creation] p. 233

`XPRSaddsets, XPRSaddsets64`	Allows sets to be added to the problem after passing it to the Optimizer using the input routines. [MIP Entities, Problem Creation]	p. 235
`XPRScopyprob`	Copies information defined for one problem to another. [Problem Creation]	p. 265
`XPRScreateprob`	Sets up a new problem within the Optimizer. [Problem Creation]	p. 266
`XPRSdelcols`	Delete columns from a matrix. [Problem Creation]	p. 268
`XPRSdelgencons`	Delete general constraints from a problem. [Piecewise Linear and General Constraints, Problem Creation]	p. 271
`XPRSdelindicators`	Delete indicator constraints. [MIP Entities, Problem Creation]	p. 272
`XPRSdelobj`	Removes an objective function from a multi-objective problem. [Problem Creation]	p. 273
`XPRSdelpwlcons`	Delete piecewise linear constraints from a problem. [Piecewise Linear and General Constraints, Problem Creation]	p. 274
`XPRSdelqmatrix`	Deletes the quadratic part of a row or of the objective function. [Problem Creation, Quadratic]	p. 275
`XPRSdelrows`	Delete rows from a matrix. [Problem Creation]	p. 276
`XPRSdelsets`	Delete sets from a problem. [MIP Entities, Problem Creation]	p. 277
`XPRSdestroyprob`	Removes a given problem and frees any memory associated with it following manipulation and optimization. [Problem Creation]	p. 278
`XPRSreadprob`	Reads an (X)MPS or LP format matrix from file. [File IO, Problem Creation]	p. 447
`XPRSsetindicators`	Specifies that a set of rows in the matrix will be treated as indicator constraints during a tree search. [Problem Creation]	p. 495

7.32.2 Problem Creation controls

`CHECKINPUTDATA`	Check input arrays for bad data. [File IO, Problem Creation]	p. 547
`EXTRACOLS`	The initial number of extra columns to allow for in the matrix. [Problem Creation]	p. 563
`EXTRAELEMS`	The initial number of extra matrix elements to allow for in the matrix, including coefficients for cuts. [Problem Creation]	p. 564
`EXTRAMIPENTS`	The initial number of extra MIP entities to allow for. [Problem Creation]	p. 564

7.33 Problem Information

Reference section for functions, controls, and attributes related to querying information about a problem from the Optimizer.

7.33.1 Problem Information library functions

7.33.2 Problem Information attributes

ORIGINALSETMEMBERS	Number of variables within special ordered sets (set members) in the original matrix before presolving. [MIP Entities, Problem Information]	p. 684
ORIGINALSETS	Number of special ordered sets in the original matrix before presolving. [MIP Entities, Problem Information]	p. 685
PWLCONS	Number of piecewise linear constraints in the problem. [Piecewise Linear and General Constraints, Problem Information]	p. 688
PWLPOINTS	Number of breakpoints of piecewise linear constraints in the problem. [Piecewise Linear and General Constraints, Problem Information]	p. 688
QCELEMS	Number of quadratic row coefficients in the matrix. [Problem Information, Quadratic]	p. 689
QCONSTRAINTS	Number of rows with quadratic coefficients in the matrix. [Problem Information, Quadratic]	p. 689
QELEMS	Number of quadratic non-zeros in the objective. [Problem Information, Quadratic]	p. 689
ROWS	Number of rows (i.e. constraints) in the matrix. [Problem Information]	p. 690
SETMEMBERS	Number of variables within special ordered sets (set members) in the matrix. [MIP Entities, Problem Information]	p. 691
SETS	Number of special ordered sets in the matrix. [MIP Entities, Problem Information]	p. 691
SPARECOLS	Number of spare columns in the matrix. [Problem Information]	p. 693
SPAREELEMS	Number of spare matrix elements in the matrix. [Problem Information]	p. 693
SPAREMIPENTS	Number of spare MIP entities in the matrix. [MIP Entities, Problem Information]	p. 693
SPAREROWS	Number of spare rows in the matrix. [Problem Information]	p. 694
SPARESETELEMS	Number of spare set elements in the matrix. [MIP Entities, Problem Information]	p. 694
SPARESETS	Number of spare sets in the matrix. [MIP Entities, Problem Information]	p. 694

7.34 Problem Input

Reference section for functions, controls, and attributes related to problem input. The load functions in this section can be used to load an entire optimization problem into the Optimizer, which can be considerably faster than incrementally creating a problem.

7.34.1 Problem Input library functions

7.35 Problem Modification

Reference section for functions, controls, and attributes related to Problem Modification. Make adjustments to the current problem in the Optimizer.

7.35.1 Problem Modification library functions

`XPRSchgobj`	Used to change the objective function coefficients. [Problem Modification]	p. 255
`XPRSchgobjn`	Modifies one or more coefficients of an objective function in a multi-objective problem. [Multiobjective, Problem Modification]	p. 254
`XPRSchgobjsense`	Changes the problem's objective function sense to minimize or maximize. [Problem Modification]	p. 256
`XPRSchgqobj`	Used to change a single quadratic coefficient in the objective function corresponding to the variable pair `(objqcol1,objqcol2)` of the Hessian matrix. [Problem Modification, Quadratic]	p. 257
`XPRSchgqrowcoeff`	Changes a single quadratic coefficient in a row. [Problem Modification, Quadratic]	p. 258
`XPRSchgrhs`	Used to change right–hand side values of the problem. [Problem Modification]	p. 259
`XPRSchgrhsrange`	Used to change the range for a row of the problem matrix. [Problem Modification]	p. 260
`XPRSchgrowtype`	Used to change the type of a row in the matrix. [Problem Modification]	p. 261
`XPRSclearrowflags`	Clears extra information attached to a range of rows. [Problem Modification]	p. 263
`XPRSfixmipentities`	Fixes all the MIP entities to the values of the last found MIP solution. [MIP Entities, Problem Modification]	p. 284

7.36 Problem Transformation

Reference section for functions, controls, and attributes related to transformation for specific types of constraints before the actual solution process begins and problem permutation.

7.36.1 Problem Transformation controls

`AUTOSCALING`	Whether the Optimizer should automatically select between different scaling algorithms. [Numerics, Problem Transformation]	p. 528
`DUALIZE`	For a linear problem or the initial linear relaxation of a MIP, determines whether to form and solve the dual problem. [LP, Problem Transformation]	p. 560
`INDLINBIGM`	During presolve, indicator constraints will be linearized using a BigM coefficient whenever that BigM coefficient is small enough. [Presolve, Problem Transformation]	p. 582
`INDPRELINBIGM`	During presolve, indicator constraints will be linearized using a BigM coefficient whenever that BigM coefficient is small enough. [Presolve, Problem Transformation]	p. 583

`PRECONVERTOBJTOCONS`	Presolve: convert a linear or quadratic objective function into an objective transfer constraint [Presolve, Problem Transformation] p. 620
`PRECONVERTSEPARABLE`	Presolve: reformulate problems with a non-diagonal quadratic objective and/or constraints as diagonal quadratic or second-order conic constraints. [Presolve, Problem Transformation] p. 620
`PREPERMUTE`	This bit-vector control (see Section 9.2) specifies whether to randomly permute rows, columns and MIP entities when starting the presolve. [Bit-vector, Presolve, Problem Transformation] p. 624
`PREPERMUTESEED`	This control sets the seed for the pseudo-random number generator for permuting the problem when starting the presolve. [Presolve, Problem Transformation] p. 624
`PRESORT`	This bit-vector control (see Section 9.2) specifies whether to sort rows, columns and MIP entities by their names when starting the presolve. [Bit-vector, Presolve, Problem Transformation] p. 629
`PWLDUALREDUCTIONS`	This parameter specifies whether dual reductions should be applied to reduce the number of columns, rows and SOS-constraints added when transforming piecewise linear objectives and constraints to MIP structs. [Piecewise Linear and General Constraints, Presolve, Problem Transformation] p. 631
`PWLNONCONVEXTRANSFORMATION`	This control specifies the reformulation method for piecewise linear constraints at the beginning of the search. [Piecewise Linear and General Constraints, Presolve, Problem Transformation] p. 631
`SCALING`	This bit-vector control (see Section 9.2) determines how the Optimizer will rescale a model internally before optimization. [Bit-vector, Numerics, Problem Transformation] p. 640
`TUNERPERMUTE`	Tuner: Defines the number of permutations to solve for each control setting. [Problem Transformation, Tuner] p. 651

7.37 Propagation

Reference section for functions, controls, and attributes related to Propagation, which denotes local logical deductions on column bounds during the Branch and Bound Search. Propagation routines are more lightweight than presolving and are hence applicable at every node of the search tree.

7.37.1 Propagation controls

`ALTERNATIVEREDCOSTS`	Controls aggressiveness of searching for alternative reduced cost [Propagation] p. 527
`GLOBALSPATIALBRANCHPROPAGATIONEFFORT`	Limits the effort that is spent on propagation during spatial branching. [Global, Propagation] p. 571

7.38 Quadratic

Reference section for functions, controls, and attributes related to optimization problems with Quadratic terms in the objective function (called 'QP' or 'MIQP'), or in the rows of the problem (called 'QCP' or 'MIQCP'), or both in the objective and the rows.

7.38.1 Quadratic library functions

`XPRSloadmiqp, XPRSloadmiqp64` Used to load a MIQP problem, hence a MIP with quadratic objective coefficients, into the Optimizer data structures.
[Problem Input, Quadratic] p. 417

`XPRSloadqcqp, XPRSloadqcqp64` Used to load a quadratic problem with quadratic side constraints into the Optimizer data structure.
[Problem Input, Quadratic] p. 408

`XPRSloadqp, XPRSloadqp64` Used to load a quadratic problem into the Optimizer data structure.
[Problem Input, Quadratic] p. 420

7.38.2 Quadratic console functions

`CHECKCONVEXITY` Convexity checker.
[Linear Algebra, Quadratic] p. 247

7.38.3 Quadratic controls

`EIGENVALUETOL` A quadratic matrix is considered not to be positive semi-definite, if its smallest eigenvalue is smaller than the negative of this value.
[Quadratic, Tolerances] p. 562

`IFCHECKCONVEXITY` Determines if the convexity of the problem is checked before optimization.
[Presolve, Quadratic] p. 581

`MIQCPALG` This control determines which algorithm is to be used to solve mixed integer quadratic constrained and mixed integer second order cone problems.
[Branch and Bound Search, Quadratic] p. 608

`PREBNDREDQUAD` Determines if convex quadratic constraints should be used for inferring bound reductions on variables when solving a MIP.
[Presolve, Quadratic] p. 617

`PRECONEDECOMP` Presolve: decompose regular and rotated cones with more than two elements and apply Outer Approximation on the resulting components.
[Presolve, Quadratic] p. 619

`PREELIMQUAD` Presolve: Allows for elimination of quadratic variables via doubleton rows.
[Presolve, Quadratic] p. 622

`QCCUTS` Branch and Bound: Limit on the number of rounds of outer approximation cuts generated for the root node, when solving a mixed integer quadratic constrained or mixed integer second order conic problem with outer approximation.
[Cuts, Quadratic] p. 632

`QCROOTALG` This control determines which algorithm is to be used to solve the root of a mixed integer quadratic constrained or mixed integer second order cone problem, when outer approximation is used.
[Quadratic, Root] p. 632

QSIMPLEXOPS	Controls the behavior of the quadratic simplex solvers via a bit-vector (see Section 9.2). [Bit-vector, Quadratic, Simplex]	p. 633
QUADRATICUNSHIFT	Determines whether an extra solution purification step is called after a solution found by the quadratic simplex (either primal or dual). [Quadratic, Simplex]	p. 633
REPAIRINDEFINITEQ	Controls if the optimizer should make indefinite quadratic matrices positive definite when it is possible. [Quadratic]	p. 636
RLTCUTS	Determines whether RLT cuts should be separated in the Xpress Global Solver. [Cuts, Quadratic]	p. 637
TREEQCCUTS	Branch and Bound: Limit on the number of rounds of outer approximation cuts generated for nodes other than the root node, when solving a mixed integer quadratic constrained or mixed integer second order conic problem with outer approximation. [Branch and Bound Search, Quadratic]	p. 648

7.38.4 Quadratic attributes

CONEELEMS	Number of second order cone coefficients in the problem. [Problem Information, Quadratic]	p. 663
CONES	Number of second order and rotated second order cones in the problem. [Problem Information, Quadratic]	p. 664
ORIGINALQCELEMS	Number of quadratic row coefficients in the original matrix before presolving. [Problem Information, Quadratic]	p. 683
ORIGINALQCONSTRAINTS	Number of rows with quadratic coefficients in the original matrix before presolving. [Problem Information, Quadratic]	p. 683
ORIGINALQELEMS	Number of quadratic non-zeros in the original objective before presolving. [Problem Information, Quadratic]	p. 684
QCELEMS	Number of quadratic row coefficients in the matrix. [Problem Information, Quadratic]	p. 689
QCONSTRAINTS	Number of rows with quadratic coefficients in the matrix. [Problem Information, Quadratic]	p. 689
QELEMS	Number of quadratic non-zeros in the objective. [Problem Information, Quadratic]	p. 689

7.39 Root

Reference section for functions, controls, and attributes related to the root node of the Branch and Bound Search. During root node processing, all derived information such as cutting planes or domain reductions are still valid for the whole problem before the search enters the tree search. For this reason, the root node deserves special attention. The optimizer will typically spend a larger amount of time at the root node than at any other node deeper in the tree.

7.39.1 Root controls

`QCROOTALG`	This control determines which algorithm is to be used to solve the root of a mixed integer quadratic constrained or mixed integer second order cone problem, when outer approximation is used. [Quadratic, Root] p. 632
`ROOTPRESOLVE`	Determines if presolving should be performed on the problem after the tree search has finished with root cutting and heuristics. [Root] p. 638

7.40 Save Restore

Reference section for functions, controls, and attributes related to Save and Restore functionality of the Optimizer.

7.40.1 Save Restore library functions

`XPRSrestore`	Restores the Optimizer's data structures from a file created by `XPRSsave` (`SAVE`). [File IO, Save Restore] p. 487
`XPRSsave, XPRSsaveas`	Saves the current data structures to file and terminates the run [File IO, Save Restore] p. 489

7.41 Sensitivity Analysis

Reference section for functions, controls, and attributes related to Sensitivity Analysis.

7.41.1 Sensitivity Analysis library functions

`XPRSbndsa`	Returns upper and lower sensitivity ranges for specified variables' lower and upper bounds. [LP, Sensitivity Analysis] p. 240
`XPRSestimaterowdualranges`	Performs a dual side range sensitivity analysis, i.e. calculates estimates for the possible ranges for dual values. [LP, Sensitivity Analysis] p. 281
`XPRSobjsa`	Returns upper and lower sensitivity ranges for specified objective function coefficients. [LP, Sensitivity Analysis] p. 429
`XPRSrhssa`	Returns upper and lower sensitivity ranges for specified right hand side (RHS) function coefficients. [LP, Sensitivity Analysis] p. 488

7.42 Simplex

Reference section for functions, controls, and attributes related to the Simplex algorithms (primal, dual, network) that the Optimizer uses to solve linear programming problems (LPs) as standalone optimization problems or as relaxations during the Branch and Bound Search.

7.42.1 Simplex library functions

7.42.2 Simplex controls

INVERTMIN	Simplex: The minimum number of iterations between full inversions of the basis matrix. [Simplex]	p. 584
KEEPBASIS	Simplex: This determines whether the basis should be kept when reoptimizing a problem. [LP, Simplex]	p. 585
MARKOWITZTOL	The Markowitz tolerance used for the factorization of the basis matrix. [LP, Simplex]	p. 589
NETSTALLLIMIT	Limit the number of degenerate pivots of the network simplex algorithm, before switching to either primal or dual simplex, depending on ALGAFTERNETWORK. [Limits, Simplex]	p. 612
PIVOTTOL	Simplex: The zero tolerance for matrix elements. [Simplex, Tolerances]	p. 616
PPFACTOR	The partial pricing candidate list sizing parameter. [Simplex]	p. 616
PRICINGALG	Simplex: This determines the primal simplex pricing method. [LP, Simplex]	p. 629
PRIMALOPS	Primal simplex: allows fine tuning the variable selection in the primal simplex solver. [Bit-vector, LP, Simplex]	p. 630
PRIMALPERTURB	The factor by which the problem will be perturbed prior to optimization by primal simplex. [LP, Simplex]	p. 630
PRIMALUNSHIFT	Determines whether primal is allowed to call dual to unshift. [LP, Simplex]	p. 630
QSIMPLEXOPS	Controls the behavior of the quadratic simplex solvers via a bit-vector (see Section 9.2). [Bit-vector, Quadratic, Simplex]	p. 633
QUADRATICUNSHIFT	Determines whether an extra solution purification step is called after a solution found by the quadratic simplex (either primal or dual). [Quadratic, Simplex]	p. 633
REFACTOR	Indicates whether the optimization should restart using the current representation of the factorization in memory. [Branch and Bound Search, Simplex]	p. 634
RELPIVOTTOL	Simplex: At each iteration a pivot element is chosen within a given column of the matrix. [Simplex, Tolerances]	p. 636
SIFTING	Determines whether to enable sifting algorithm with the dual simplex method. [LP, Simplex]	p. 641
SIFTPASSES	Determines how quickly we allow to grow the worker problems during the sifting algorithm. [LP, Simplex]	p. 641

`SIFTPRESOLVEOPS`	Determines the presolve operations for solving the subproblems during the sifting algorithm. [LP, Simplex]	p. 642
`SIFTSWITCH`	Determines which algorithm to use for solving the subproblems during sifting. [LP, Simplex]	p. 642

7.42.3 Simplex attributes

`SIMPLEXITER`	Number of simplex iterations performed. [LP, Simplex]	p. 691

7.43 Solution Process

Reference section for functions, controls, and attributes related to the invocation of a solution algorithm of the Optimizer, and the most fundamental status information after the Optimizer returns.

7.43.1 Solution Process library functions

`XPRSinterrupt`	Interrupts the Optimizer algorithms. [Solution Process]	p. 390
`XPRSlpoptimize`	This function begins a search for the optimal continuous (LP) solution. [LP, Solution Process]	p. 424
`XPRSmipoptimize`	This function begins a tree search for the optimal MIP solution. [MIP Entities, Solution Process]	p. 428
`XPRSoptimize`	This function begins a search for the optimal solution of the problem. [Solution Process]	p. 430
`XPRSpostsolve`	Postsolve the current matrix when it is in a presolved state. [Presolve, Solution Process]	p. 437

7.43.2 Solution Process attributes

`OBJVAL`	Value of the objective function of the last problem solved with XPRSoptimize. [Problem Information, Solution Process]	p. 678
`OPTIMIZETYPEUSED`	The type of solver used in the last call to `XPRSoptimize`, `XPRSmipoptimize`, `XPRSlpoptimize` or `XPRSnlpoptimize`. [Solution Process]	p. 680
`SOLSTATUS`	Status of the solution of the last problem solved with XPRSoptimize. [Solution, Solution Process]	p. 692
`SOLVESTATUS`	Status of the solve of the last problem solved with XPRSoptimize. [Solution Process]	p. 693
`STOPSTATUS`	Status of the optimization process. [Solution Process]	p. 694

`TIME` Time spent solving the problem as measured by the optimizer.
[Solution Process] p. 696

`WORK` Amount of deterministic algorithmic ”work” spent since the invocation of the search process.
[Determinism, Solution Process] p. 698

7.44 Solution Refinement

Reference section for functions, controls, and attributes related to Solution Refinement, an iterative procedure that helps reducing numerical artifacts in a solution.

7.44.1 Solution Refinement controls

`FEASTOLTARGET` This specifies the target feasibility tolerance for the solution refiner.
[Solution Refinement, Tolerances] p. 567

`LPREFINEITERLIMIT` This specifies the simplex iteration limit the solution refiner can spend in attempting to increase the accuracy of an LP solution.
[LP, Solution Refinement] p. 589

`MIPREFINEITERLIMIT` This defines an effort limit expressed as simplex iterations for the MIP solution refiner.
[Limits, Solution Refinement] p. 604

`MIPTOLTARGET` Target MIPTOL value used by the automatic MIP solution refiner as defined by `REFINEOPS`.
[Solution Refinement, Tolerances] p. 608

`OPTIMALITYTOLTARGET` This specifies the target optimality tolerance for the solution refiner.
[Solution Refinement, Tolerances] p. 614

`REFINEOPS` This specifies when the solution refiner should be executed to reduce solution infeasibilities.
[Bit-vector, Solution Refinement] p. 634

7.45 Solution

Reference section for functions, controls, and attributes related to the handling of optimal or intermediate solutions.

7.45.1 Solution library functions

`XPRSaddcbintsol` Declares a user integer solution callback function, called every time an integer solution is found by heuristics or during the Branch and Bound search.
[Callback, Solution] p. 204

`XPRSaddcbpreintsol` Declares a user integer solution callback function, called when an integer solution is found by heuristics or during the branch and bound search, but before it is accepted by the Optimizer.
[Callback, Solution] p. 214

7.45.2 Solution console functions

7.45.3 Solution controls

CLAMPING	This bit-vector control (see Section 9.2) allows for the adjustment of returned solution values such that they are always within bounds. [Bit-vector, Numerics, Solution]	p. 549
MAXMIPSOL	Branch and Bound: This specifies a limit on the number of integer solutions to be found by the Optimizer. [Limits, Solution]	p. 594

7.45.4 Solution attributes

MIPOBJVAL	Objective function value of the last integer solution found. [Branch and Bound Search, Solution]	p. 675
MIPSOLNODE	Node at which the last integer feasible solution was found. [Branch and Bound Search, Solution]	p. 675
MIPSOLS	Number of integer solutions that have been found. [Branch and Bound Search, Solution]	p. 675
MIPSOLTIME	Time at which the last integer feasible solution was found. [Branch and Bound Search, Solution]	p. 676
SOLSTATUS	Status of the solution of the last problem solved with XPRSoptimize. [Solution, Solution Process]	p. 692

7.46 System

Reference section for functions, controls, and attributes related to Computer hardware and system specifications.

7.46.1 System controls

CORESPERCPU	Used to override the detected value of the number of cores on a CPU. [System]	p. 552
CPUPLATFORM	Newton Barrier: Selects the AMD, Intel x86 or ARM vectorization instruction set that Barrier should run optimized code for. [System]	p. 553
CPUTIME	How time should be measured when timings are reported in the log and when checking against time limits [System]	p. 553

7.46.2 System attributes

CORESDETECTED	Number of logical cores detected by the optimizer, which is the total number of threads the hardware can execute across all CPUs. [System]	p. 664

`CORESPERCPUDETECTED`	Number of logical cores per CPU unit detected by the optimizer, which is the number of threads each CPU can execute. [System]	p. 664
`CPUSDETECTED`	Number of CPU units detected by the optimizer. [System]	p. 665
`PHYSICALCORESDETECTED`	The total number of physical cores across all CPUs detected by the optimizer. [System]	p. 686
`PHYSICALCORESPERCPUDETECTED`	The number of physical cores per CPU detected by the optimizer. [System]	p. 686
`SYSTEMMEMORY`	The amount of non problem specific memory used by the solver. [Memory, System]	p. 695

7.47 Tolerances

Reference section for functions, controls, and attributes related to feasibility and optimality tolerances.

7.47.1 Tolerances controls

`BARDUALSTOP`	Newton barrier and hybrid gradient: This is a convergence parameter, representing the tolerance for dual infeasibilities. [Barrier, Primal Dual Hybrid Gradient, Tolerances]	p. 532
`BARPRIMALSTOP`	Newton barrier and hybrid gradient: This is a convergence parameter, indicating the tolerance for primal infeasibilities. [Barrier, Primal Dual Hybrid Gradient, Tolerances]	p. 539
`EIGENVALUETOL`	A quadratic matrix is considered not to be positive semi-definite, if its smallest eigenvalue is smaller than the negative of this value. [Quadratic, Tolerances]	p. 562
`ELIMTOL`	The Markowitz tolerance for the elimination phase of the presolve. [Presolve, Tolerances]	p. 563
`ETATOL`	Tolerance on eta elements. [Simplex, Tolerances]	p. 563
`FEASTOL`	This tolerance determines when a solution is treated as feasible. [Tolerances]	p. 566
`FEASTOLPERTURB`	This tolerance determines how much a feasible primal basic solution is allowed to be perturbed when performing basis changes. [Simplex, Tolerances]	p. 566
`FEASTOLTARGET`	This specifies the target feasibility tolerance for the solution refiner. [Solution Refinement, Tolerances]	p. 567
`INPUTTOL`	The tolerance on input values elements. [File IO, Problem Creation, Tolerances]	p. 583
`MATRIXTOL`	The zero tolerance on matrix elements. [Tolerances]	p. 590

`MIPTOL` — Branch and Bound: This is the tolerance within which a decision variable's value is considered to be integral.
[Tolerances] p. 607

`MIPTOLTARGET` — Target MIPTOL value used by the automatic MIP solution refiner as defined by `REFINEOPS`.
[Solution Refinement, Tolerances] p. 608

`OPTIMALITYTOL` — Simplex: This is the zero tolerance for reduced costs.
[Tolerances] p. 613

`OPTIMALITYTOLTARGET` — This specifies the target optimality tolerance for the solution refiner.
[Solution Refinement, Tolerances] p. 614

`OUTPUTTOL` — Zero tolerance on print values.
[File IO, Tolerances] p. 615

`PIVOTTOL` — Simplex: The zero tolerance for matrix elements.
[Simplex, Tolerances] p. 616

`RELPIVOTTOL` — Simplex: At each iteration a pivot element is chosen within a given column of the matrix.
[Simplex, Tolerances] p. 636

`SOSREFTOL` — The minimum relative gap between the ordering values of elements in a special ordered set.
[Tolerances] p. 643

7.48 Tuner

Reference section for functions, controls, and attributes related to the algorithmic tuning of control values to customize the Optimizer.

7.48.1 Tuner library functions

`XPRStune` — This function begins a tuner session for the current problem.
[Tuner] p. 513

`XPRStuneprobsetfile` — This function begins a tuner session for a set of problems.
[Tuner] p. 514

`XPRStunerreadmethod` — This function loads a user defined tuner method from the given file.
[File IO, Tuner] p. 515

`XPRStunerwritemethod` — This function writes the current tuner method to a given file or prints it to the console.
[File IO, Tuner] p. 516

7.48.2 Tuner console functions

`TUNE` — Console Tuner command.
[Tuner] p. 511

7.48.3 Tuner controls

`TUNERHISTORY`	Tuner: Whether to reuse and append to previous tuner results of the same problem. [Tuner]	p. 648
`TUNERMAXTIME`	Tuner: The maximum time in seconds that the tuner will run before it terminates. [Limits, Tuner]	p. 649
`TUNERMETHOD`	Tuner: Selects a factory tuner method. [Tuner]	p. 649
`TUNERMETHODFILE`	Tuner: Defines a file from which the tuner can read user-defined tuner method. [File IO, Tuner]	p. 650
`TUNERMODE`	Tuner: Whether to always enable the tuner or disable it. [Tuner]	p. 650
`TUNEROUTPUT`	Tuner: Whether to output tuner results and logs to the file system. [File IO, Tuner]	p. 651
`TUNEROUTPUTPATH`	Tuner: Defines a root path to which the tuner writes the result file and logs. [File IO, Tuner]	p. 651
`TUNERPERMUTE`	Tuner: Defines the number of permutations to solve for each control setting. [Problem Transformation, Tuner]	p. 651
`TUNERSESSIONNAME`	Tuner: Defines a session name for the tuner. [Tuner]	p. 652
`TUNERTARGET`	Tuner: Defines the tuner target – what should be evaluated when comparing two runs with different control settings. [Tuner]	p. 652
`TUNERTHREADS`	Tuner: the number of threads used by the tuner. [Parallel, Tuner]	p. 653
`TUNERVERBOSE`	Tuner: whether the tuner should prints detailed information for each run. [Logging, Tuner]	p. 653

Chapter 8

Console and Library Functions

A large number of routines are available for both Console and Library users of the FICO Xpress Optimizer, ranging from simple routines for the input and solution of problems from matrix files to sophisticated callback functions and greater control over the solution process. Of these, the core functionality is available to both sets of users and comprises the 'Console Mode'. Library users additionally have access to a set of more 'advanced' functions, which extend the functionality provided by the Console Mode, providing more control over their program's interaction with the Optimizer and catering for more complicated problem development.

8.1 Console Mode Functions

With both the Console and Advanced Mode functions described side-by-side in this chapter, library users can use this as a quick reference for the full capabilities of the Optimizer library. For users of Console Optimizer, only the following functions will be of relevance:

Command	Description	Page
CHECKCONVEXITY	Convexity checker. [Linear Algebra, Quadratic]	p. 247
EXIT	Terminate the Console Optimizer. [Misc]	p. 280
HELP	Quick reference help for the Optimizer console. [Misc]	p. 378
IIS	Console IIS command. [Infeasibility]	p. 379
PRINTSOL	Write the current solution to screen. [Solution]	p. 441
QUIT	Terminate the Console Optimizer. [Misc]	p. 442
STOP	Terminate the Console Optimizer. [Misc]	p. 506
TUNE	Console Tuner command. [Tuner]	p. 511
SETARCHCONSISTENCY	Sets whether to force the same execution path on various CPU architecture extensions, in particular (pre-)AVX and AVX2. [Global Environment]	p. 171
BASISSTABILITY	Calculates various measures for the stability of the current basis, including the basis condition number. [Linear Algebra, LP]	p. 239
CHGOBJSENSE	Changes the problem's objective function sense to minimize or maximize. [Problem Modification]	p. 256
DUMPCONTROLS	Displays the list of controls and their current value for those controls that have been set to a non default value. [Controls and Attributes]	p. 279

For a list of functions by task, refer to 2.8.

8.2 Layout for Function Descriptions

All functions mentioned in this chapter are described under the following set of headings:

Function Name

The description of each routine starts on a new page. The library name for a function is on the left and the Console Optimizer command name, if one exists, is on the right.

Purpose

A short description of the routine and its purpose begins the information section.

Synopsis

A synopsis of the syntax for usage of the routine is provided. "Optional" arguments and flags may be specified as `NULL` if not required. Where this possibility exists, it will be described alongside the argument, or in the Further Information at the end of the routine's description. If the function is available in the Console, the library syntax is described first, followed by the Console Optimizer syntax.

Arguments

A list of arguments to the routine with a description of possible values for them follows.

Error Values

Optimizer return codes are described in Chapter 11.

Likely error values returned by this for each function are listed in the Error Values section. A description of the error may be obtained using the `XPRSgetlasterror` function. If no attention need be drawn to particular error values, this section will be omitted.

Associated Controls

Controls which affect a given routine are listed next, separated into lists by type. The control name given here should have `XPRS_` prefixed by library users, in a similar way to the `XPRSgetintattrib` example in the Error Values section above. Console Xpress users should use the controls without this prefix, as described in FICO Xpress Getting Started manual. These controls must be set before the routine is called if they are to have any effect.

Examples

One or two examples are provided which explain certain aspects of the routine's use.

Further Information

Additional information not contained elsewhere in the routine's description is provided at the end.

Related Topics

Finally a list of related routines and topics is provided for comparison and reference.

XPRS_bo_addbounds

Purpose

Adds new bounds to a branch of a user branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_addbounds(XPRSbranchobject bo, int branch, int nbounds,
        const char bndtype[], const int colind[], const double bndval[]);
```

Arguments

`bo`	The user branching object to modify.
`branch`	The number of the branch to add the new bounds for. This branch must already have been created using `XPRS_bo_addbranches`. Branches are indexed starting from zero.
`nbounds`	Number of new bounds to add.
`bndtype`	Character array of length `nbounds` indicating the type of bounds to add: `L` Lower bound. `U` Upper bound.
`colind`	Integer array of length `nbounds` containing the column indices for the new bounds.
`bndval`	Double array of length `nbounds` giving the bound values.

Example

See `XPRS_bo_create` for an example using `XPRS_bo_addbounds`.

See also example `mostviolated.c`.

Related topics

`XPRS_bo_create`.

XPRS_bo_addbranches

Purpose

Adds new, empty branches to a user defined branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_addbranches(XPRSbranchobject bo, int nbranches);
```

Arguments

`bo`	The user branching object to modify.
`nbranches`	Number of new branches to create.

Example

See `XPRS_bo_create` for an example using `XPRS_bo_addbranches`.

See also example `mostviolated.c`.

Related topics

`XPRS_bo_create`.

XPRS_bo_addcuts

Purpose

Adds stored user cuts as new constraints to a branch of a user branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_addcuts(XPRSbranchobject bo, int branch, int ncuts,
       const XPRScut cutind[]);
```

Arguments

`bo`	The user branching object to modify.
`branch`	The number of the branch to add the cuts for. This branch must already have been created using `XPRS_bo_addbranches`. Branches are indexed starting from zero.
`ncuts`	Number of cuts to add.
`cutind`	Array of length `ncuts` containing the pointers to user cuts that should be added to the branch.

Related topics

`XPRS_bo_create`, `XPRS_bo_addrows`.

XPRS_bo_addrows

Purpose

Adds new constraints to a branch of a user branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_addrows(XPRSbranchobject bo, int branch, int nrows, int
      ncoefs, const char rowtype[], const double rhs[], const int start[],
      const int colind[], const double rowcoef[]);
```

Arguments

`bo`	The user branching object to modify.
`branch`	The number of the branch to add the new constraints for. This branch must already have been created using `XPRS_bo_addbranches`. Branches are indexed starting from zero.
`nrows`	Number of new constraints to add.
`ncoefs`	Number of non-zero coefficients in all new constraints.
`rowtype`	Character array of length `nrows` indicating the type of constraints to add: `L` Less than type. `G` Greater than type. `E` Equality type.
`rhs`	Double array of length `nrows` containing the right hand side values.
`start`	Integer array of length `nrows` containing the offsets of the `colind` and `rowcoef` arrays of the start of the non zero coefficients in the new constraints.
`colind`	Integer array of length `ncoefs` containing the column indices for the non zero coefficients.
`rowcoef`	Double array of length `ncoefs` containing the non zero coefficient values.

Example

The following function will create a branching object that branches on constraints $x_1 + x_2 \geqslant 1$ or $x_1 + x_2 \leqslant 0$:

```
XPRSbranchobject CreateConstraintBranch(XPRSprob xp_mip, int icol)
{
  char   cRowType;
  double dRowRHS;
  int    mRowBeg;
  int    mElemCol[2];
  double dElemVal[2];

  XPRSbranchobject bo = NULL;
  int isoriginal = 1;

  /* Create the new object with two empty branches. */
  XPRS_bo_create(&bo, xp_mip, isoriginal);
  XPRS_bo_addbranches(bo, 2);

  /* Add the constraint x1 + x2 >= 1. */
  cRowType = 'G';
  dRowRHS  = 1.0;
  mRowBeg  = 0;
  mElemCol[0] = 0; mElemCol[1] = 1;
  dElemVal[0] = 1.0; dElemVal[1] = 1.0;
  XPRS_bo_addrows
    (bo, 0, 1, 2, &cRowType, &dRowRHS, &mRowBeg, mElemCol, dElemVal);
```

```
  /* Add the constraint x1 + x2 <= 0. */
  cRowType = 'L';
  dRowRHS  = 0.0;
  XPRS_bo_addrows
    (bo, 1, 1, 2, &cRowType, &dRowRHS, &mRowBeg, mElemCol, dElemVal);

  /* Set a low priority value so our branch object is picked up */
  /* before the default branch candidates. */
  XPRS_bo_setpriority(bo, 100);

  return bo;
}
```

Related topics

XPRS_bo_create.

XPRS_bo_create

Purpose

Creates a new user defined branching object for the Optimizer to branch on. This function should be called only from within one of the callback functions set by `XPRSaddcboptnode`, `XPRSaddcbchgbranchobject`, or `XPRSaddcbpreintsol` (only if the `soltype` is 0).

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_create(XPRSbranchobject* p_bo, XPRSprob prob, int
       isoriginal);
```

Arguments

`p_bo`	Pointer to where the new object should be returned.
`prob`	The problem structure that the branching object should be created for.
`isoriginal`	If the branching object will be set up for the original matrix, which determines how column indices are interpreted when adding bounds and rows to the object: `0` Column indices should refer to the current (presolved) node problem. `1` Column indices should refer to the original matrix.

Further information

1. In addition to the standard MIP entities supported by the Optimizer, the Optimizer also allows the user to define their own entities for branching, using branching objects.

2. A branching object of type `XPRSbranchobject` should provide a linear description of how to branch on the current node for a user's entities. Any number of branches is allowed and each branch description can contain any combination of columns bounds and new constraints.

3. Branching objects must always contain at least one branch and all branches of the object must contain at least one bound or constraint.

4. If multiple lower or multiple upper bounds on the same variable are given, the strictest one will be applied.

5. When the Optimizer branches the current node on a user's branching object, a new child node will be created for each branch defined in the object. The child nodes will inherit the bounds and constraint of the current node, plus any new bounds or constraints defined for that branch in the object.

6. Inside the callback function set by `XPRSaddcboptnode`, a user can define any number of branching objects and pass them to the Optimizer. These objects are added to the set of infeasible entities for the current node and the Optimizer will select a best candidate from this extended set using all of its normal evaluation methods.

7. The callback function set by `XPRSaddcbchgbranchobject` can be used to override the Optimizer's selected branching candidate with the user's own object. This can for example be used to modify how to branch on the MIP entity selected by the Optimizer.

8. The following functions are available to set up a new user branching object:

`XPRS_bo_create`	Creates a new, empty branching object with no branches.
`XPRS_bo_addbranches`	Adds new, empty branches to the object. Branches must be created before column bounds or rows can be added to a branch.
`XPRS_bo_addbounds`	Adds new column bounds to a given branch of the object.
`XPRS_bo_addrows`	Adds new constraints to a given branch of the object.
`XPRS_bo_setpriority`	Sets the priority value for the object. These are equivalent to the priority values for regular MIP entities that can be set through directives (see also Appendix A.5).
`XPRS_bo_setpreferredbranch`	Specifies which of the child nodes corresponding to the branches of the object should be explored first.
`XPRS_bo_store`	Adds the created object to the candidate list for branching.

Example

The following function will create a branching object equivalent to a standard binary branch on a column:

```
XPRSbranchobject CreateBinaryBranchObject(XPRSprob xp_mip, int icol)
{
  char   cBndType;
  double dBndValue;
  int isoriginal = 1;

  XPRSbranchobject bo = NULL;

  /* Create the new object with two empty branches. */
  XPRS_bo_create(&bo, xp_mip, isoriginal);
  XPRS_bo_addbranches(bo, 2);

  /* Add bounds to branch the column to either zero or one. */
  cBndType = 'U';
  dBndValue = 0.0;
```

```
  XPRS_bo_addbounds(bo, 0, 1, &cBndType, &icol, &dBndValue);
  cBndType = 'L';
  dBndValue = 1.0;
  XPRS_bo_addbounds(bo, 1, 1, &cBndType, &icol, &dBndValue);

  /* Set a low priority value so our branch object is picked up */
  /* before the default branch candidates. */
  XPRS_bo_setpriority(bo, 100);

  return bo;
}
```

See also example `mostviolated.c`.

Related topics

`XPRSaddcboptnode`, `XPRSaddcbchgbranchobject`.

XPRS_bo_destroy

Purpose

Frees all memory for a user branching object, when the object was not stored with the Optimizer.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_destroy(XPRSbranchobject bo);
```

Argument

`bo`	The user branching object to free.

Related topics

`XPRS_bo_create`, `XPRS_bo_store`.

Example

See also example `mostviolated.c`.

XPRS_bo_getbounds

Purpose

Returns the bounds for a branch of a user branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_getbounds(XPRSbranchobject bo, int branch, int*
        p_nbounds, int maxbounds, char bndtype[], int colind[], double
        bndval[]);
```

Arguments

`bo`	The branching object to inspect.
`branch`	The number of the branch to get the bounds for.
`p_nbounds`	Location where the number of bounds for the given branch should be returned.
`maxbounds`	Maximum number of bounds to return.
`bndtype`	Character array of length `maxbounds` where the types of bounds will be returned: `L` Lower bound. `U` Upper bound. Allowed to be NULL.
`colind`	Integer array of length `maxbounds` where the column indices will be returned. Allowed to be NULL.
`bndval`	Double array of length `maxbounds` where the bound values will be returned. Allowed to be NULL.

Related topics

`XPRS_bo_create`, `XPRS_bo_addbounds`.

XPRS_bo_getbranches

Purpose

Returns the number of branches of a branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_getbranches(XPRSbranchobject bo, int* p_nbranches);
```

Arguments

`bo`	The user branching object to inspect.
`p_nbranches`	Memory where the number of branches should be returned.

Related topics

`XPRS_bo_create`, `XPRS_bo_addbranches`.

XPRS_bo_getid

Purpose

Returns the unique identifier assigned to a branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_getid(XPRSbranchobject bo, int* p_id);
```

Arguments

`bo`	A branching object.
`p_id`	Pointer to an integer where the identifier should be returned.

Further information

1. Branching objects associated with existing column entities (binaries, integers, semi-continuous and partial integers), are given an identifier from 1 to `MIPENTS`.
2. Branching objects associated with existing Special Ordered Sets, are given an identifier from `MIPENTS+1` to `MIPENTS+SETS`.
3. User created branching objects will always have a negative identifier.

Related topics

`XPRS_bo_create`.

XPRS_bo_getlasterror

Purpose

Returns the last error encountered during a call to the given branch object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_getlasterror(XPRSbranchobject bo, int* p_msgcode, char*
       msg, int maxbytes, int* p_nbytes);
```

Arguments

bo	The branch object.
p_msgcode	Location where the error code will be returned. Can be NULL if not required.
msg	A character buffer of size maxbytes in which the last error message relating to the given branching object will be returned.
maxbytes	The size of the character buffer msg.
p_nbytes	The size of the required character buffer to fully return the error string.

Example

The following shows how this function might be used in error checking:

```
XPRSbranchobject bo;
...
char* cbuf;
int cbuflen;
if (XPRS_bo_setpreferredbranch(bo,3)) {
  XPRS_bo_getlasterror(bo,NULL,NULL,0,&cbuflen);
  cbuf = malloc(cbuflen);
  XPRS_bo_getlasterror(bo,NULL, cbuf, cbuflen, NULL);
  printf("ERROR when setting preferred branch: %s\n", cbuf);
}
```

Related topics

XPRS_ge_addcbmsghandler.

XPRS_bo_getrows

Purpose

Returns the constraints for a branch of a user branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_getrows(XPRSbranchobject bo, int branch, int* p_nrows,
      int maxrows, int* p_ncoefs, int maxcoefs, char rowtype[], double
      rhs[], int start[], int colind[], double rowcoef[]);
```

Arguments

`bo`	The user branching object to inspect.
`branch`	The number of the branch to get the constraints from.
`p_nrows`	Memory location where the number of rows should be returned.
`maxrows`	Maximum number of rows to return.
`p_ncoefs`	Memory location where the number of non zero coefficients in the constraints should be returned.
`maxcoefs`	Maximum number of non zero coefficients to return.
`rowtype`	Character array of length `maxrows` where the types of the rows will be returned: `L` Less than type. `G` Greater than type. `E` Equality type.
`rhs`	Double array of length `maxrows` where the right hand side values will be returned.
`start`	Integer array of length `maxrows` which will be filled with the offsets of the `colind` and `rowcoef` arrays of the start of the non zero coefficients in the returned constraints.
`colind`	Integer array of length `maxcoefs` which will be filled with the column indices for the non zero coefficients.
`rowcoef`	Double array of length `maxcoefs` which will be filled with the non zero coefficient values.

Related topics

`XPRS_bo_create`, `XPRS_bo_addrows`.

XPRS_bo_setpreferredbranch

Purpose

Specifies which of the child nodes corresponding to the branches of the object should be explored first.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_setpreferredbranch(XPRSbranchobject bo, int branch);
```

Arguments

`bo`	The user branching object.
`branch`	The number of the branch to mark as preferred.

Related topics

XPRS_bo_create.

XPRS_bo_setpriority

Purpose

Sets the priority value of a user branching object.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_setpriority(XPRSbranchobject bo, int priority);
```

Arguments

`bo`	The user branching object.
`priority`	The new priority value to assign to the branching object, which must be a number from 0 to 1000. User branching objects are created with a default priority value of 500.

Further information

1. A candidate branching object with lowest priority number will always be selected for branching before an object with a higher number.
2. Priority values must be an integer from 0 to 1000. User branching objects and MIP entities are by default assigned a priority value of 500. Special branching objects, such as those arising from structural branches or split disjunctions are assigned a priority value of 400.

Related topics

`XPRS_bo_create`, Section A.5.

XPRS_bo_store

Purpose

Adds a new user branching object to the Optimizer's list of candidates for branching. This function is available only through the callback functions set by `XPRSaddcboptnode` or `XPRSaddcbpreintsol` (if the `soltype` is 0).

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_store(XPRSbranchobject bo, int* p_status);
```

Arguments

`bo`	The new user branching object to store. After this call the `bo` object is no longer valid and should not be referred to again.
`p_status`	The returned status from checking the provided branching object: `0` The object was accepted successfully. `1` Failed to presolve the object due to dual reductions in presolve. `2` Failed to presolve the object due to duplicate column reductions in presolve. `3` The object contains an empty branch. The object was not added to the candidate list if a non zero status is returned.

Further information

1. To ensure that a user branching object expressed in terms of the original matrix columns can be applied to the presolved problem, it might be necessary to turn off certain presolve operations.
2. If any of the original matrix columns referred to in the object are unbounded, dual reductions might prevent the corresponding bound or constraint from being presolved. To avoid this, dual reductions should be turned off in presolve, by clearing bit 3 of the integer control `PRESOLVEOPS`.
3. If one or more of the original matrix columns of the object are duplicates in the original matrix, but not in the branching object, it might not be possible to presolve the object due to duplicate column eliminations in presolve. To avoid this, duplicate column eliminations should be turned off in presolve, by clearing bit 5 of `PRESOLVEOPS`.
4. As an alternative to turning off the above mentioned presolve features, it is possible to protect individual columns of a the problem from being modified by presolve. Use the `XPRSloadsecurevecs` function to mark any columns that might be branched on using branching objects.
5. The branching candidate is only for the current node. Once Xpress branched on that node, all branching candidates are deleted.
6. The branching candidates stored by this function are evaluated together with Xpress's internal branching candidates. A branching object's priority can be used to give preference to objects, see `XPRS_bo_setpriority`.

Related topics

`XPRS_bo_create`, `XPRS_bo_validate`.

XPRS_bo_validate

Purpose

Verifies that a given branching object is valid for branching on the current branch-and-bound node of a MIP solve. The function will check that all branches are non-empty, and if required, verify that the branching object can be presolved.

Topic area

Branching Object

Synopsis

```
int XPRS_CC XPRS_bo_validate(XPRSbranchobject bo, int* p_status);
```

Arguments

`bo`	A branching object.
`p_status`	The returned status from checking the provided branching object: `0` The object is acceptable. `1` Failed to presolve the object due to dual reductions in presolve. `2` Failed to presolve the object due to duplicate column reductions in presolve. `3` The object contains an empty branch.

Related topics

`XPRS_bo_create`.

XPRS_ge_addcbmsghandler

Purpose

Declares an output callback function in the global environment, called every time a line of message text is output by any data in the library. This callback function will be called in addition to any output callbacks already added by `XPRS_ge_addcbmsghandler`.

Topic areas

Global Environment, Callback

Synopsis

```
int XPRS_CC XPRS_ge_addcbmsghandler(int (XPRS_CC *msghandler)(XPRSobject
       xprsobj, void * cbdata, void * thread, const char * msg, int msgtype,
       int msgcode), void *data, int priority);
```

Arguments

`msghandler`	The callback function which takes six arguments, `xprsobj`, `cbdata`, `thread`, `msg`, `msgtype` and `msgcode`. Use a NULL value to cancel a callback function.
`xprsobj`	The data sending the message. Use `XPRSgetobjecttypename` to get the name of the data type.
`cbdata`	The user-defined data passed to the callback function.
`thread`	The system id of the thread sending the message cast to a `void *`.
`msg`	A null terminated character array (string) containing the message, which may simply be a new line. When the callback is called for the first time `msg` will be a NULL pointer.
`msgtype`	Indicates the type of output message: `1` information messages; `2` (not used); `3` warning messages; `4` error messages. When the callback is called for the first time `msgtype` will be a negative value.
`msgcode`	The number associated with the message. If the message is an error or a warning then you can look up the number in the section Optimizer Error and Warning Messages for advice on what it means and how to resolve the associated issue.
`data`	A user-defined data to be passed to the callback function.
`priority`	An integer that determines the order in which multiple message handler callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

1. To send all messages to a log file the built in message handler `XPRSlogfilehandler` can be used. This can be done with:

```
XPRS_ge_addcbmsghandler(XPRSlogfilehandler, "log.txt", 0);
```

2. The return value of the callback function `msghandler` is currently ignored.

Related topics

`XPRS_ge_removecbmsghandler`, `XPRSgetobjecttypename`.

XPRS_ge_getcomputeallowed

Purpose

Query whether the current application is allowed to use the Insight Compute interface.

Topic areas

Global Environment, Compute Interface

Synopsis

```
int XPRS_CC XPRS_ge_getcomputeallowed(int* p_allow);
```

Argument

`p_allow`	Memory location in which the value will be returned. Value will equal one of the following constants.	
	`XPRS_ALLOW_COMPUTE_ALWAYS`	Always allow solves to be sent to Compute.
	`XPRS_ALLOW_COMPUTE_NEVER`	Never allow solves to be sent to Compute.
	`XPRS_ALLOW_COMPUTE_DEFAULT`	Allow solves to be sent to Compute only from non-OEM applications.

Example

The following shows how this function might be used:

```
int isComputeAllowed;
if (XPRS_ge_getcomputeallowed(&isComputeAllowed)!=0) {
  switch (isComputeAllowed) {
    case XPRS_ALLOW_COMPUTE_ALWAYS:
      printf("Compute integration is always allowed.\n");
      break;
    case XPRS_ALLOW_COMPUTE_NEVER:
      printf("Compute integration is never allowed.\n");
      break;
    case XPRS_ALLOW_COMPUTE_DEFAULT:
      printf("Compute integration is allowed if "
        "this is not an OEM application.\n");
      break;
  }
}
```

Further information

This value supplied by this function describes whether this process would be allowed to use Insight 5 Compute Interface - this is not affected by whether or not the user has tried to enable use of the Insight 5 Compute interface.

Related topics

`XPRS_ge_setcomputeallowed`.

XPRS_ge_getlasterror

Purpose

Returns the last error encountered during a call to the Xpress global environment.

Topic area

Global Environment

Synopsis

```
int XPRS_CC XPRS_ge_getlasterror(int* p_msgcode, char* msg, int maxbytes,
        int* p_nbytes);
```

Arguments

`p_msgcode`	Memory location in which the error code will be returned. Can be NULL if not required.
`msg`	A character buffer of size `maxbytes` in which the last error message relating to the global environment will be returned. If the message is longer than `maxbytes` then it will be truncated. The message will always be terminated by a `NUL` character (provided `maxbytes` is bigger than 0).
`maxbytes`	The size of the character buffer `msg`.
`p_nbytes`	Memory location in which the minimum required size of the buffer to hold the full error string will be returned. This minimum size includes the terminating `NUL` character. Can be NULL if not required.

Example

The following shows how this function might be used in error checking:

```
char* cbuf;
int cbuflen;
if (XPRS_ge_addcbmsghandler(myfunc,NULL,0)!=0) {
  XPRS_ge_getlasterror(NULL,NULL,0,&cbuflen);
  cbuf = malloc(cbuflen);
  XPRS_ge_getlasterror(NULL, cbuf, cbuflen, NULL);
  printf("ERROR from Xpress global environment: %s\n", cbuf);
}
```

Related topics

`XPRS_ge_addcbmsghandler`.

XPRS_ge_removecbmsghandler

Purpose

Removes a message callback function previously added by `XPRS_ge_addcbmsghandler`. The specified callback function will no longer be called after it has been removed.

Topic areas

Global Environment, Callback

Synopsis

```
int XPRS_CC XPRS_ge_removecbmsghandler(int (XPRS_CC *msghandler)
       (XPRSobject object, void * cbdata, void * thread, const char * msg,
       int msgtype, int msgcode), void * data);
```

Arguments

`msghandler`	The callback function to remove. If NULL then all message callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all message callbacks with the function pointer `msghandler` will be removed.

Related topics

`XPRS_ge_addcbmsghandler`

XPRS_ge_setarchconsistency — SETARCHCONSISTENCY

Purpose

Sets whether to force the same execution path on various CPU architecture extensions, in particular (pre-)AVX and AVX2.

Topic area

Global Environment

Synopsis

```
int XPRS_CC XPRS_ge_setarchconsistency(int consistent);
SETARCHCONSISTENCY consistent
```

Argument

`consistent`	Whether to force the same execution path: `0` Do not force the same execution path (default behavior); `1` Force the same execution path.

Further information

Note that, using this general environment API function is different from setting the `CPUPLATFORM` control. Setting `CPUPLATFORM` selects a vectorization instruction set for the barrier method.

XPRS_ge_setcomputeallowed

Purpose

Set whether the current application is allowed to use the Insight Compute interface.

Topic areas

Global Environment, Compute Interface

Synopsis

```
int XPRS_CC XPRS_ge_setcomputeallowed(int allow);
```

Argument

`allow`	Whether the Insight Compute interface may be used; must be one of the following constants:	
	`XPRS_ALLOW_COMPUTE_ALWAYS`	Always allow solves to be sent to Compute.
	`XPRS_ALLOW_COMPUTE_NEVER`	Never allow solves to be sent to Compute.
	`XPRS_ALLOW_COMPUTE_DEFAULT`	Allow solves to be sent to Compute only from non-OEM applications.

Example

The following shows how this function might be used:

```
XPRS_ge_setcomputeallowed(XPRS_ALLOW_COMPUTE_NEVER);
```

Further information

1. This function controls whether this process would be allowed to use the Insight Compute Interface if the user tries to enable it.
2. If the user enables the Insight Compute Interface but the value specified through this function does not allow the Insight Compute Interface to be used, any solves will terminate with an immediate error. This function can be used to prevent solves being sent to Insight Compute but cannot be used to force solves to be performed locally. The purpose of this function is to allow an application to prevent the optimization model being sent to the Insight Compute Interface.

Related topics

`XPRS_ge_getcomputeallowed`.

XPRS_nml_addnames

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

The XPRS_nml_* functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user. Use the XPRS_nml_addnames to add names to a name list, or modify existing names on a namelist.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_addnames(XPRSnamelist nml, const char names[], int
       first, int last);
```

Arguments

nml	The name list to which you want to add names. Must be an object previously returned by XPRS_nml_create, as XPRSnamelist objects returned by other functions are immutable and cannot be changed.
names	Character buffer containing the null-terminated string names. The length of an individual name (including the terminating NUL) must not exceed 1MB.
first	The index of the first name to add/replace. Name indices in a namelist always start from 0.
last	The index of the last name to add/replace.

Example

```
char mynames[0] = "fred\0jim\0sheila"
...
XPRS_nml_addnames(nml,mynames,0,2);
```

Related topics

XPRS_nml_create, XPRS_nml_removenames, XPRS_nml_copynames, XPRSaddnames.

XPRS_nml_copynames

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

The `XPRS_nml_*` functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user.

`XPRS_nml_copynames` allows you to copy all the names from one name list to another. As name lists representing row/column names cannot be modified, `XPRS_nml_copynames` will be most often used to copy such names to a namelist where they can be modified, for some later use.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_copynames(XPRSnamelist dest, XPRSnamelist src);
```

Arguments

`dest`	The namelist object to copy names to. Any names already in this name list will be removed. Must be an object previously returned by `XPRS_nml_create`.
`src`	The namelist object from which all the names should be copied.

Example

```
XPRSprob prob;
XPRSnamelist rnames, rnames_on_prob;
...
/* Create a namelist */
XPRS_nml_create(&rnames);
/* Get a namelist through which we can access the row names */
XPRSgetnamelistobject(prob,1,&rnames_on_prob);
/* Now copy these names from the immutable 'XPRSprob' namelist
   to another one */
XPRS_nml_copynames(rnames,rnames_on_prob);
/* The names in the list can now be modified then put to some
   other use */
```

Related topics

`XPRS_nml_create`, `XPRS_nml_addnames`, `XPRSgetnamelistobject`.

XPRS_nml_create

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

The `XPRS_nml_*` functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user.

`XPRS_nml_create` will create a new namelist to which the user can add, remove and otherwise modify names.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_create(XPRSnamelist* p_nml);
```

Argument

`p_nml`	Pointer to location where the new namelist will be returned.

Example

```
XPRSnamelist mylist;
XPRS_nml_create(&mylist);
```

Related topics

`XPRSgetnamelistobject`, `XPRS_nml_destroy`.

XPRS_nml_destroy

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

Destroys a namelist and frees any memory associated with it. Note you need only destroy namelists created by `XPRS_nml_destroy` - those returned by `XPRSgetnamelistobject` are automatically destroyed when you destroy the problem object.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_destroy(XPRSnamelist nml);
```

Argument

`nml`	The namelist to be destroyed.

Example

```
XPRSnamelist mylist;
XPRS_nml_create(&mylist);
...
XPRS_nml_destroy(&mylist);
```

Related topics

`XPRS_nml_create`, `XPRSgetnamelistobject`, `XPRSdestroyprob`.

XPRS_nml_findname

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

The `XPRS_nml_*` functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user.

`XPRS_nml_findname` returns the index of the given name in the given name list.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_findname(XPRSnamelist nml, const char* name, int*
        p_index);
```

Arguments

`nml`	The namelist in which to look for the name.
`name`	Null-terminated string containing the name for which to search.
`p_index`	Pointer to variable in which the index of the name is returned, or in which is returned -1 if the name is not found in the namelist.

Example

```
XPRSnamelist mylist;
int idx;
...
XPRS_nml_findname(mylist, "profit_after_work", &idx);
if (idx==-1)
  printf("'profit_after_work' was not found in the namelist");
else
  printf("'profit_after_work' was found at position %d", idx);
```

Related topics

`XPRS_nml_addnames`, `XPRS_nml_getnames`.

XPRS_nml_getlasterror

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

Returns the last error encountered during a call to a namelist object.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_getlasterror(XPRSnamelist nml, int* p_msgcode, char*
       msg, int maxbytes, int* p_nbytes);
```

Arguments

`nml`	The namelist object.
`p_msgcode`	Variable in which the error code will be returned. Can be NULL if not required.
`msg`	A character buffer of size `iStringBufferBytes` in which the last error message relating to this namelist will be returned.
`maxbytes`	The size of the character buffer `msg`.
`p_nbytes`	Memory location in which the minimum required size of the buffer to hold the full error string will be returned. Can be NULL if not required.

Example

```
XPRSnamelist nml;
char* cbuf;
int cbuflen;
...
if (XPRS_nml_removenames(nml,2,35)) {
  XPRS_nml_getlasterror(nml, NULL, NULL, 0, &cbuflen);
  cbuf = malloc(cbuflen);
  XPRS_nml_getlasterror(nml, NULL, cbuf, cbuflen, NULL);
  printf("ERROR removing names: %s\n", cbuf);
}
```

Related topics

None.

XPRS_nml_getmaxnamelen

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

The `XPRS_nml_*` functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user. `XPRS_nml_getmaxnamelen` returns the length of the longest name in the namelist.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_getmaxnamelen(XPRSnamelist nml, int* p_namelen);
```

Arguments

`nml`	The namelist object.
`p_namelen`	Pointer to a variable into which shall be written the length of the longest name.

Related topics

None.

XPRS_nml_getnamecount

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

The `XPRS_nml_*` functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user.

`XPRS_nlm_getnamecount` returns the number of names in the namelist.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_getnamecount(XPRSnamelist nml, int* p_count);
```

Arguments

`nml`	The namelist object.
`p_count`	Pointer to a variable into which shall be written the number of names.

Example

```
XPRSnamelist mylist;
int count;
...
XPRS_nml_getnamecount(mylist,&count);
printf("There are %d names", count);
```

Related topics

None.

XPRS_nml_getnames

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

The `XPRS_nml_*` functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user. The `XPRS_nml_getnames` function returns some of the names held in the name list. The names shall be returned in a character buffer, and with each name being separated by a `NULL` character.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_getnames(XPRSnamelist nml, int pad, char buffer[], int
       maxbytes, int* p_nbytes, int first, int last);
```

Arguments

`nml`	The namelist object.
`pad`	The minimum length of each name. If >0 then names shorter than `pad` will be concatenated with whitespace to make them this length.
`buffer`	Buffer of length `maxbytes` into which the names shall be returned.
`maxbytes`	The maximum number of bytes that may be written to the character buffer `buffer`.
`p_nbytes`	A pointer to a variable into which will be written the number of bytes required to contain the names. May be NULL if not required.
`first`	The index of the first name in the namelist to return. Note name list indexes always start from 0.
`last`	The index of the last name in the namelist to return.

Example

```
XPRSnamelist mylist;
char* cbuf;
int o, i, cbuflen;
...
/* Find out how much space we'll require for these names */
XPRS_nml_getnames(mylist, 0, NULL, 0, &cbuflen, 0, 5 );
/* Allocate a buffer large enough to hold the names */
cbuf = malloc(cbuflen);
/* Retrieve the names */
XPRS_nml_getnames(mylist, 0, cbuf, cbuflen, NULL, 0, 5);
/* Display the names */
o=0;
for (i=0;i<6;i++) {
  printf("Name #%d = %s\n", i, cbuf+o);
  o += strlen(cbuf)+1;
}
```

Related topics

None.

XPRS_nml_removenames

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.
The `XPRS_nml_*` functions provide a simple, generic interface to lists of names, which may be names of rows/columns on a problem or may be a list of arbitrary names provided by the user.
`XPRS_nml_removenames` will remove the specified names from the name list. Any subsequent names will be moved down to replace the removed names.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRS_nml_removenames(XPRSnamelist nml, int first, int last);
```

Arguments

`nml`	The name list from which you want to remove names. Must be an object previously returned by `XPRS_nml_create`, as `XPRSnamelist` objects returned by other functions are immutable and cannot be changed.
`first`	The index of the first name to remove. Note that indices for names in a name list always start from 0.
`last`	The index of the last name to remove.

Example

```
XPRS_nml_removenames(mylist, 3, 5);
```

Related topics

`XPRS_nml_addnames`.

XPRSaddcbbariteration

Purpose

Declares a barrier iteration callback function, called after each iteration during the interior point algorithm, with the ability to access the current barrier solution/slack/duals or reduced cost values, and to ask barrier to stop. This callback function will be called in addition to any callbacks already added by XPRSaddcbbariteration.

This functionality is not available for the hybrid gradient method.

Topic areas

Callback, Barrier

Synopsis

```
int XPRS_CC XPRSaddcbbariteration (XPRSprob prob, void (XPRS_CC
       *bariteration)( XPRSprob cbprob, void *cbdata, int *p_action), void
       *data, int priority);
```

Arguments

`prob`	The current problem.
`bariteration`	The callback function itself. This takes three arguments, `cbprob`, `cbdata`, and `p_action` serving as an integer return value. This function is called at every barrier iteration.
`cbprob`	The problem passed to the callback function, `bariteration`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbbariteration`.
`p_action`	Defines a return value controlling barrier: `<0` continue with the next iteration; `=0` let barrier decide (use default stopping criteria); `1` barrier stops with status not defined; `2` barrier stops with optimal status; `3` barrier stops with dual infeasible status; `4` barrier stops wih primal infeasible status.
`data`	A user-defined data to be passed to the callback function, `bariteration`.
`priority`	An integer that determines the order in which callbacks of this type will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

This simple example demonstrates how the solution might be retrieved for each barrier iteration.

```
// Barrier iteration callback
void XPRS_CC BarrierIterCallback(XPRSprob cbprob,
                  void *cbdata, int *p_action) {
  int currentIteration, nCols, nRows;
  double primalObj, dualObj, gap, primalInf, dualInf,
    complementaryGap;

  my_object_s *my = (my_object_s *) cbdata;

  XPRSgetintattrib(cbprob, XPRS_BARITER, &currentIteration);
  XPRSgetintattrib(cbprob, XPRS_INPUTCOLS, &nCols);
  XPRSgetintattrib(cbprob, XPRS_INPUTROWS, &nRows);

  // try to get all the solution values
  XPRSgetcallbacksolution(cbprob, NULL, my->x, 0, nCols - 1);
  XPRSgetcallbackslacks(cbprob, NULL, my->slacks, 0, nRows - 1);
```

```
  XPRSgetcallbackduals(cbprob, NULL, my->y, 0, nRows - 1);
  XPRSgetcallbackredcosts(cbprob, NULL, my->dj, 0, nCols - 1);

  XPRSgetdblattrib(cbprob, XPRS_BARPRIMALOBJ, &primalObj);
  XPRSgetdblattrib(cbprob, XPRS_BARdUALOBJ, &dualObj);
  gap = dualObj - primalObj;
  XPRSgetdblattrib(cbprob, XPRS_BARPRIMALINF, &primalInf);
  XPRSgetdblattrib(cbprob, XPRS_BARDUALINF, &dualInf);
  XPRSgetdblattrib(cbprob, XPRS_BARCGAP, &complementaryGap);

  // decide if stop or continue
  *p_action = BARRIER_CHECKSTOPPING;
  if (currentIteration >= 50
      || gap <= 0.1*max(fabs(primalObj),fabs(dualObj))) {
    *p_action = BARRIER_OPTIMAL;
  }
}

// To set callback:
XPRSaddcbbariteration(xprob, BarrierIterCallback, (void *) &my, 0);
```

Further information

1. Only the functions for retrieving the solution and the attribute/control value retrieving and setting routines are expected to be called from the callback.
2. General barrier iteration values are available by using `XPRSgetdblattrib` to retrieve:
   - `BARPRIMALOBJ` - current primal objective
   - `BARDUALOBJ` - current dual objective
   - `BARPRIMALINF` - current primal infeasibility
   - `BARDUALINF` - current dual infeasibility
   - `BARCGAP` - current complementary gap
3. Please note that these values refer to the scaled and presolved problem used by barrier, and may differ from the ones calculated from the postsolved solution returned by `XPRSgetcallbacksolution`.

Related topics

`XPRSremovecbbariteration`.

XPRSaddcbbarlog

Purpose

Declares a barrier log callback function, called at each iteration during the interior point algorithm. This callback function will be called in addition to any barrier log callbacks already added by XPRSaddcbbarlog.

This functionality is not available for the hybrid gradient method.

Topic areas

Callback, Barrier

Synopsis

```
int XPRS_CC XPRSaddcbbarlog (XPRSprob prob, int (XPRS_CC *barlog)(XPRSprob
       cbprob, void *cbdata), void *data, int priority);
```

Arguments

`prob`	The current problem.
`barlog`	The callback function itself. This takes two arguments, `cbprob` and `cbdata`, and has an integer return value. If the value returned by `barlog` is nonzero, the solution process will be interrupted. This function is called at every barrier iteration.
`cbprob`	The problem passed to the callback function, `barlog`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbbarlog`.
`data`	A user-defined data to be passed to the callback function, `barlog`.
`priority`	An integer that determines the order in which multiple barrier log callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

This simple example prints a line to the screen for each iteration of the algorithm.

```
XPRSaddcbbarlog(prob,barLog,NULL,0);
XPRSlpoptimize(prob,"b");
```

The callback function might resemble:

```
int XPRS_CC barLog(XPRSprob prob, void *data)
{
  printf("Next barrier iteration\n");
}
```

Further information

If the callback function returns a nonzero value, the Optimizer run will be interrupted.

Related topics

`XPRSremovecbbarlog`, `XPRSaddcbmiplog`, `XPRSaddcblplog`, `XPRSaddcbmessage`.

XPRSaddcbafterobjective

Purpose

Declares a callback which will be called after each objective in a multi-objective problem is solved.

Topic areas

Callback, Multiobjective

Synopsis

```
int XPRS_CC XPRSaddcbafterobjective (XPRSprob prob, void (XPRS_CC
      *afterobjective)(XPRSprob cbprob, void *cbdata), void *data, int
      priority);
```

Arguments

`prob`	The current problem.
`afterobjective`	The callback function itself. This takes two arguments, `cbprob` and `cbdata`. This function is called after every solve.
`cbprob`	The problem passed to the callback function, `afterobjective`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbafterobjective`.
`data`	A user-defined data to be passed to the callback function, `afterobjective`.
`priority`	An integer that determines the order in which multiple after-objective callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

This example logs a message after each objective is solved:

```
XPRSaddcbafterobjective(prob,afterobjective,NULL,0);
XPRSoptimize(prob,"",NULL,NULL);
```

The callback function might resemble:

```
int XPRS_CC afterobjective(XPRSprob prob, void *data)
{
  int solvedobjs;
  XPRSgetintattribute(prob, XPRS_SOLVEDOBJS, &solvedobjs);
  printf("Completed solve %d\n", solvedObjs);
}
```

Further information

If the solve does not stop early, this callback will be invoked `OBJSTOSOLVE` times. Note that this can be less than than `OBJECTIVES` times, since all objectives with the same priority are combined into a single objective and solved at once.

Related topics

`XPRSremovecbafterobjective`, `XPRSaddcbbeforeobjective`.

XPRSaddcbbeforeobjective

Purpose

Declares a callback which will be called before each objective in a multi-objective problem is solved.

Topic areas

Callback, Multiobjective

Synopsis

```
int XPRS_CC XPRSaddcbbeforeobjective (XPRSprob prob, void (XPRS_CC
       *beforeobjective)(XPRSprob cbprob, void *cbdata), void *data, int
       priority);
```

Arguments

`prob`	The current problem.
`beforeobjective`	The callback function itself. This takes two arguments, `cbprob` and `cbdata`. This function is called before every solve.
`cbprob`	The problem passed to the callback function, `beforeobjective`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbbeforeobjective`.
`data`	A user-defined data to be passed to the callback function, `beforeobjective`.
`priority`	An integer that determines the order in which multiple before-objective callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

This example sets a node limit for the second multi-objective solve:

```
XPRSaddcbbeforeobjective(prob,beforeobjective,NULL,0);
XPRSoptimize(prob,"",NULL,NULL);
```

The callback function might resemble:

```
int XPRS_CC beforeobjective(XPRSprob prob, void *data)
{
  int solvedobjs;
  XPRSgetintattribute(prob, XPRS_SOLVEDOBJS, &solvedobjs);
  if (solvedobjs == 1) {
    XPRSsetintcontrol(prob, XPRS_MAXNODE, 100);
  }
}
```

Further information

1. If the solve does not stop early, this callback will be invoked `OBJSTOSOLVE` times. Note that this can be less than than `OBJECTIVES` times, since all objectives with the same priority are combined into a single objective and solved at once.
2. During this callback, the main objective (as queried with `XPRSgetobj`) will reflect the current objective being solved, which may be a weighted combination of several objectives with the same priority level. But changes made to the problem to fix earlier objectives at their optimal value will not be visible in the problem during this callback.
3. Controls set during this callback will retain their values when solving all subsequent objectives, except for controls which set resource limits: `TIMELIMIT`, `SOLTIMELIMIT`, `MAXTIME`, `MAXNODE`, `LPITERLIMIT`, `NLPMAXTIME` and `SLPITERLIMIT`. These are recomputed before each call to `beforeobjective`, based on their original values before the solve and the resources consumed when solving all previous objectives. Changes to resource limit controls will therefore only be retained while solving the current objective.

Related topics

`XPRSremovecbbeforeobjective`, `XPRSaddcbafterobjective`.

XPRSaddcbcomputerestart

Purpose

Declares a callback to be called when a solve executed in compute mode needs to be restarted.

Topic areas

Callback, Compute Interface

Synopsis

```
int XPRS_CC XPRSaddcbcomputerestart (XPRSprob prob, void (XPRS_CC
      *computerestart)(XPRSprob cbprob, void *cbdata), void *data, int
      priority);
```

Arguments

`prob`	The current problem.
`computerestart`	The callback function itself. This takes two arguments, `cbprob` and `cbdata`, and has no return value. This function is called when a solve had to be restarted in compute mode.
`cbprob`	The problem passed to the callback function, `computerestart`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbcomputerestart`.
`data`	A user-defined data to be passed to the callback function, `computerestart`.
`priority`	An integer that determines the order in which multiple computerestart callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related topics

`XPRSremovecbcomputerestart`

XPRSaddcbpresolve

Purpose

Declares a callback to be called after presolve has been performed.

Topic areas

Callback, Presolve

Synopsis

```
int XPRS_CC XPRSaddcbpresolve (XPRSprob prob, void (XPRS_CC
        *presolve)(XPRSprob cbprob, void *cbdata), void *data, int priority);
```

Arguments

`prob`	The current problem.
`presolve`	The callback function itself. This takes two arguments, `cbprob` and `cbdata`, and has no return value. This function is called after presolve is complete.
`cbprob`	The problem passed to the callback function.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbpresolve`.
`data`	A user-defined data to be passed to the callback function, `presolve`.
`priority`	An integer that determines the order in which multiple presolve callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related topics

`XPRSremovecbpresolve`

XPRSaddcbchecktime

Purpose

Declares a callback function which is called every time the Optimizer checks if the time limit has been reached. This callback function will be called in addition to any callbacks already added by XPRSaddcbchecktime.

Topic areas

Callback, Limits

Synopsis

```
int XPRS_CC XPRSaddcbchecktime(XPRSprob prob, int (XPRS_CC
      *checktime)(XPRSprob cbprob, void* cbdata), void* data, int
      priority);
```

Arguments

`prob`	The current problem.
`checktime`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has an integer return value. This function is called every time the Optimizer checks against the time limit.
`cbprob`	The problem passed to the callback function, `checktime`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbchecktime`.
`data`	A user-defined data to be passed to the callback function, `checktime`.
`priority`	An integer that determines the order in which multiple checktime callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

If the callback function returns a nonzero value the solution process will be interrupted. If multiple callbacks of this type are registered, then all of them will be invoked and the solution process will be interrupted if any of them returned a nonzero value.

Related topics

`XPRSremovecbchecktime` `MAXTIME` `CHECKSONMAXTIME` `MAXCHECKSONMAXTIME`

XPRSaddcbchgbranch

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSaddcbchgbranchobject` instead.
Declares a branching variable callback function, called every time a new branching variable is set or selected during the branch and bound search. This callback function will be called in addition to any change branch callbacks already added by XPRSaddcbchgbranch.

Topic areas

Callback, Branching

Synopsis

```
int XPRS_CC XPRSaddcbchgbranch(XPRSprob prob, void (XPRS_CC
      *chgbranch)(XPRSprob cbprob, void *cbdata, int *p_entity, int *p_up,
      double *p_estdeg), void *data, int priority);
```

Arguments

`prob`	The current problem.
`chgbranch`	The callback function, which takes five arguments, `cbprob`, `cbdata`, `p_entity`, `p_up` and `p_estdeg`, and has no return value. This function is called every time a new branching variable or set is selected.
`cbprob`	The problem passed to the callback function, `chgbranch`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbchgbranch`.
`p_entity`	A pointer to the variable or set on which to branch. Integer variables are identified by their column index, i.e. `0`, `1`,...(`COLS- 1`) and sets by their set index, i.e. `0`, `1`,...,(`SETS- 1`).
`p_up`	If `p_entity` is a variable, this is `1` if the upward branch is to be made first, or `0` otherwise. If `p_entity` is a set, this is `3` if the upward branch is to be made first, or `2` otherwise.
`p_estdeg`	This value is obsolete. It will be set to zero and any returned value is ignored.
`data`	A user-defined data to be passed to the callback function, `chgbranch`.
`priority`	An integer that determines the order in which multiple branching variable callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

The arguments initially contain the default values of the branching variable, branching variable, branching direction and estimated degradation. If they are changed then the Optimizer will use the new values, if they are not changed then the default values will be used.

Related topics

`XPRSremovecbchgbranch`, `XPRSaddcbchgnode`, `XPRSaddcboptnode`, `XPRSaddcbinfnode`, `XPRSaddcbintsol`, `XPRSaddcbnodecutoff`, `XPRSaddcbprenode`.

XPRSaddcbchgbranchobject

Purpose

Declares a callback function that will be called after the selection of a MIP entity to branch on. This callback allows the user to inspect and replace the Optimizer's choice of how to branch the current node. This callback will also be called in the case when there are no candidates to branch on, that is, when all MIP entities are already satisfied. This callback function will be called in addition to any callbacks already added by XPRSaddcbchgbranchobject.

Topic areas

Callback, Branching

Synopsis

```
int XPRS_CC XPRSaddcbchgbranchobject(XPRSprob prob, void (XPRS_CC
       *chgbranchobject)(XPRSprob cbprob, void* cbdata, XPRSbranchobject
       branch, XPRSbranchobject* p_newbranch), void* data, int priority);
```

Arguments

`prob`	The current problem.
`chgbranchobject`	The callback function, which takes four arguments: `cbprob`, `cbdata`, `branch` and `p_newbranch`. This function is called every time the Optimizer has selected a candidate entity for branching.
`cbprob`	The problem passed to the callback function, `chgbranchobject`.
`cbdata`	The user defined data passed as `data` when setting up the callback with `XPRSaddcbchgbranchobject`.
`branch`	The candidate branching data selected by the Optimizer. Will be `NULL` if no candidates exist.
`p_newbranch`	Optional new branching data to replace the Optimizer's selection. If `branch` or `NULL` is passed back, no change will be applied.
`data`	A user-defined data to be passed to the callback function, `chgbranchobject`.
`priority`	An integer that determines the order in which multiple callbacks of this type will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

1. The branching data given by the Optimizer provides a linear description of how the Optimizer intends to branch on the selected candidate. This will often be one of standard MIP entities of the current problem, but can also be e.g. a split disjunction or a structural branch, if those features are turned on.

2. The functions `XPRS_bo_getbranches`, `XPRS_bo_getbounds` and `XPRS_bo_getrows` can be used to inspect the given branching data.

3. Refer to `XPRS_bo_create` on how to create a new branching data to replace the Optimizer's selection. Note that the new branching data should be created with a priority value no higher than the current data to guarantee it will be used for branching.

Related topics

`XPRSremovecbchgbranchobject`, `XPRS_bo_create`.

Example

See also example `mostviolated.c`.

XPRSaddcbchgnode

Purpose

This subroutine is deprecated and will be removed in a future release.
Declares a callback that is fired every time the code performs a full backtrack to select a node to start a dive on. Note that it is no longer possible to change the selected node but the callback may still be used to keep track of dives. This callback function will be called in addition to any callbacks already added by XPRSaddcbchgnode.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSaddcbchgnode(XPRSprob prob, void (XPRS_CC
      *chgnode)(XPRSprob cbprob, void *cbdata, int *p_node), void *data,
      int priority);
```

Arguments

`prob`	The current problem.
`chgnode`	The callback function which takes three arguments, `cbprob`, `cbdata` and `p_node`, and has no return value. This function is called every time a new node is selected.
`cbprob`	The problem passed to the callback function, `chgnode`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbchgnode`.
`p_node`	A pointer to the number of the node selected by the Optimizer. This value cannot be changed.
`data`	A user-defined data to be passed to the callback function, `chgnode`.
`priority`	An integer that determines the order in which multiple node selection callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related topics

`XPRSremovecbchgnode`

XPRSaddcbcutlog

Purpose

Declares a cut log callback function, called each time the cut log is printed. This callback function will be called in addition to any callbacks already added by XPRSaddcbcutlog.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSaddcbcutlog(XPRSprob prob, int (XPRS_CC *cutlog)(XPRSprob
        cbprob, void *cbdata), void *data, int priority);
```

Arguments

`prob`	The current problem.
`cutlog`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has an integer return value.
`cbprob`	The problem passed to the callback function, `cutlog`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbcutlog`.
`data`	A user-defined data to be passed to the callback function, `cutlog`.
`priority`	An integer that determines the order in which multiple cut log callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

Return a non-zero value from `cutlog` to stop cutting on the current node.

Related topics

`XPRSremovecbcutlog`, `XPRSaddcbcutmgr`.

XPRSaddcbcutmgr

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSaddcboptnode` *instead.*

Declares a user-defined cut manager routine, called at each node of the branch and bound search. This callback function will be called in addition to any callbacks already added by XPRSaddcbcutmgr.

Topic areas

Callback, Cuts

Synopsis

```
int XPRS_CC XPRSaddcbcutmgr(XPRSprob prob, int (XPRS_CC *cutmgr)(XPRSprob
       cbprob, void *cbdata), void *data, int priority);
```

Arguments

`prob`	The current problem
`cutmgr`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has an integer return value. This function is called at each node in the Branch and Bound search.
`cbprob`	The problem passed to the callback function, `cutmgr`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbcutmgr`.
`data`	A user-defined data to be passed to the callback function, `cutmgr`.
`priority`	An integer that determines the order in which multiple callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

1. When returning from the user function `f_cutlog`, the Optimizer will automatically reoptimize the LP relaxation of the node problem. If a non-zero value is returned from `f_cutlog`, the function will be called again afterwards, unless the LP relaxation has become infeasible or was cut off due to the objective function value. Return `0` from `f_cutlog` to prevent the function from being called again for the same branch and bound node.

2. `f_cutlog` is called for a branch-and-bound node problem after the Optimizer has already applied any internal cuts and heuristics, but before determining if the node problem should be branched or if the node LP relaxation solution is MIP feasible.

3. The Optimizer ensures that cuts added to a node are automatically restored at descendant nodes. To do this, all cuts are stored in a cut pool and the Optimizer keeps track of which cuts from the cut pool must be restored at each node.

Related topics

`XPRSremovecbcutmgr`, `XPRSaddcbcutlog`, `CALLBACKCOUNT_CUTMGR`.

XPRSaddcbcutround

Purpose

Declares a callback function that is called when the Optimizer could separate cutting planes during the branch and bound search. The callback allows a user to add cuts managed by the Optimizer, through `XPRSaddmanagedcuts`. It also allows the user some degree of control over when and how many rounds of cuts the Optimizer should apply for the local node of the branch-and-bound search. This callback function will be called in addition to any callbacks already added by XPRSaddcbcutmgr.

Topic areas

Callback, Cuts

Synopsis

```
int XPRS_CC XPRSaddcbcutround(XPRSprob prob, void (XPRS_CC
       *cutround)(XPRSprob cbprob, void *cbdata, int ifxpresscuts, int
       *p_action), void *data, int priority);
```

Arguments

`prob`	The current problem
`cutround`	The callback function which takes four arguments, `cbprob`, `cbdata`, `ifxpresscuts` and `p_action` and has no return value. This function is called before each potential round of cut separation in the branch and bound search.
`cbprob`	The problem passed to the callback function, `cutround`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbcutround`.
`ifxpresscuts`	An integer set to 1 if the Optimizer will apply a round of cuts after this callback. 0 otherwise.
`p_action`	An integer return value that specifies the action the Optimizer should take: `-1` Continue unchanged. The default action. `0` No further rounds of cuts should be applied on this node. `1` The Optimizer should apply one more round of cutting, regardless of the value of `ifxpresscuts` `2` The Optimizer should process any changes applied during this callback and fire the callback again, but skip any Optimizer cutting.
`data`	A user-defined data to be passed to the callback function, `cutround`.
`priority`	An integer that determines the order in which multiple callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

1. Optimizer managed cuts can be added using XPRSaddmanagedcuts. The Optimizer will check for any violated cuts, together with any new or violated Optimizer created cuts (if p_action is −1 or +1), to add to the problem.

2. It is possible to use XPRSaddcuts, XPRSloadcuts and XPRSdelcuts to add and remove cutting planes and XPRSchgbounds to tighten bounds, although it is recommended to use XPRSaddmanagedcuts to provide new cutting planes to the Optimizer in this callback. Separation of delayed constraints is best done using either the optnode or preintsol callbacks.

3. It is possible to add new solutions using XPRSaddmipsol. These will be checked immediately after the callback, but adding solutions will not automatically trigger another firing of this callback.

4. If cutting planes are added with XPRSaddmanagedcuts, the callback will be fired again after the reoptimization if and only if a non-zero value is returned in p_action.

5. If cuts are added directly to, or deleted from, the problem, or bounds are tightened, a reoptimization of the node LP relaxation will automatically be triggered open the return from the callback function. If separation of Optimizer cuts is enabled or managed cuts were added during the callback, such cuts will only be separated based on the new solution after reoptimization. The callback will be fired again if and only if a non-zero value is returned in p_action.

6. The attribute CUTROUNDS can be used for determining how many rounds of cuts has been applied on the current node.

Related topics

XPRSremovecbcutround, XPRSaddmanagedcuts, XPRSaddmipsol, XPRSaddcuts, XPRSloadcuts, XPRSdelcuts, XPRSchgbounds, CUTROUNDS.

Example

See also examples els_managedcuts.c, els_usercuts.c .

XPRSaddcbdestroymt

Purpose

Declares a destroy MIP thread callback function, called every time a MIP thread is destroyed by the parallel MIP code. This callback function will be called in addition to any callbacks already added by XPRSaddcbdestroymt.

Topic areas

Callback, Parallel

Synopsis

```
int XPRS_CC XPRSaddcbdestroymt(XPRSprob prob, void (XPRS_CC
       *destroymt)(XPRSprob cbprob, void *cbdata), void *data, int
       priority);
```

Arguments

`prob`	The current thread problem.
`destroymt`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has no return value.
`cbprob`	The thread problem passed to the callback function.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbdestroymt`.
`data`	A user-defined data to be passed to the callback function.
`priority`	An integer that determines the order in which multiple callbacks of this type will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related controls

Integer

`MIPTHREADS`	Number of MIP threads to create.

Further information

This callback is useful for freeing up any user data created in the MIP thread callback.

Related topics

`XPRSremovecbdestroymt`,`XPRSaddcbmipthread`.

XPRSaddcbgapnotify

Purpose

Declares a gap notification callback, to be called when a MIP solve reaches a predefined target, set using the `MIPRELGAPNOTIFY`, `MIPABSGAPNOTIFY`, `MIPABSGAPNOTIFYOBJ` and/or `MIPABSGAPNOTIFYBOUND` controls.

Topic areas

Callback, Limits

Synopsis

```
int XPRS_CC XPRSaddcbgapnotify(XPRSprob prob, void (XPRS_CC
       *gapnotify)(XPRSprob cbprob, void* cbdata, double*
       p_relgapnotifytarget, double* p_absgapnotifytarget, double*
       p_absgapnotifyobjtarget, double* p_absgapnotifyboundtarget), void*
       data, int priority);
```

Arguments

`prob`	The current problem.
`gapnotify`	The callback function.
`cbprob`	The current problem.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbgapnotify`.
`p_relgapnotifytarget`	The value the `MIPRELGAPNOTIFY` control will be set to after this callback. May be modified within the callback in order to set a new notification target.
`p_absgapnotifytarget`	The value the `MIPABSGAPNOTIFY` control will be set to after this callback. May be modified within the callback in order to set a new notification target.
`p_absgapnotifyobjtarget`	The value the `MIPABSGAPNOTIFYOBJ` control will be set to after this callback. May be modified within the callback in order to set a new notification target.
`p_absgapnotifyboundtarget`	The value the `MIPABSGAPNOTIFYBOUND` control will be set to after this callback. May be modified within the callback in order to set a new notification target.
`data`	A user-defined data to be passed to the callback function, `gapnotify`.
`priority`	An integer that determines the order in which multiple estimate callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

The following example prints a message when the gap reaches 10% and 1%

```
void XPRS_CC gapnotify(XPRSprob prob, void* data,
  double* p_relgapnotifytarget, double* p_absgapnotifytarget,
  double* p_absgapnotifyobjtarget, double* p_absgapnotifyboundtarget)
{
  double obj, bound, relgap;
  XPRSgetdblattrib(prob, XPRS_MIPOBJVAL, &obj);
  XPRSgetdblattrib(prob, XPRS_BESTBOUND, &bound);
  if (obj != 0.0 || bound != 0.0)
    relgap = fabs((obj - bound)/ max(fabs(obj), fabs(bound)));
  else
    relgap = 0.0;
  if (relgap<=0.10) {
    printf("Gap reached 10%");
    *p_relgapnotifytarget = 0.1;
  }
  if (relgap<=0.01) {
```

```
    printf("Gap reached 1%");
    *p_relgapnotifytarget = -1; /* Don't call gapnotify again */
  }
}

XPRSsetdblcontrol( prob, XPRS_MIPRELGAPNOTIFY, 0.10 );
XPRSaddcbgapnotify( prob, gapnotify, NULL, 0 );
XPRSmipoptimize(prob, "");
```

Further information

The target values that caused the callback to be triggered will automatically be reset to prevent the same callback from being fired again.

Related topics

`MIPRELGAPNOTIFY`, `MIPABSGAPNOTIFY`, `MIPABSGAPNOTIFYOBJ`, `MIPABSGAPNOTIFYBOUND`, `XPRSremovecbgapnotify`.

XPRSaddcbgloballog

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSaddcbmiplog` *instead.*

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSaddcbgloballog(XPRSprob prob, int (XPRS_CC
       *globallog)(XPRSprob cbprob, void *cbdata), void *data, int
       priority);
```

XPRSaddcbmiplog

Purpose

Declares a MIP log callback function, called each time the MIP log is printed. This callback function will be called in addition to any callbacks already added by XPRSaddcbmiplog.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSaddcbmiplog(XPRSprob prob, int (XPRS_CC *miplog)(XPRSprob
       cbprob, void *cbdata), void *data, int priority);
```

Arguments

`prob`	The current problem.
`miplog`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has an integer return value. This function is called whenever the MIP log is printed as determined by the `MIPLOG` control.
`cbprob`	The problem passed to the callback function, `miplog`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbmiplog`.
`data`	A user-defined data to be passed to the callback function, `miplog`.
`priority`	An integer that determines the order in which multiple MIP log callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related controls

Integer

`MIPLOG`	MIP print flag.

Example

The following example prints at each node of the tree search the node number and its depth:

```
XPRSsetintcontrol(prob, XPRS_MIPLOG, 3);
XPRSaddcbmiplog(prob, mipLog, NULL, 0);
XPRSmipoptimize(prob,"");
```

The callback function may resemble:

```
int XPRS_CC mipLog(XPRSprob prob, void *data)
{
  int node, nodedepth;

  XPRSgetintattrib(prob, XPRS_NODEDEPTH, &nodedepth);
  XPRSgetintattrib(prob, XPRS_CURRENTNODE, &node);
  printf("Node %d with depth %d has just been processed\n",
         node, nodedepth);

  return 0;
}
```

See the example `depthfirst.c` in the `examples/optimizer/c` folder.

Further information

If the callback function returns a nonzero value, the tree search will be interrupted.

Related topics

`XPRSremovecbmiplog`, `XPRSaddcbbarlog`, `XPRSaddcblplog`, `XPRSaddcbmessage`.

XPRSaddcbinfnode

Purpose

Declares a user infeasible node callback function, called after the current node has been found to be infeasible during the Branch and Bound search. The callback is also invoked if the current node gets cut off, i.e., if its objective is proven to exceed the current primal bound. This callback function will be called in addition to any callbacks already added by XPRSaddcbinfnode.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSaddcbinfnode(XPRSprob prob, void (XPRS_CC
       *infnode)(XPRSprob cbprob, void *cbdata), void *data, int priority);
```

Arguments

`prob`	The current problem
`infnode`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has no return value. This function is called after the current node has been found to be infeasible.
`cbprob`	The problem passed to the callback function, `infnode`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbinfnode`.
`data`	A user-defined data to be passed to the callback function, `infnode`.
`priority`	An integer that determines the order in which multiple user infeasible node callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

The following notifies the user whenever an infeasible node is found during the tree search:

```
XPRSaddcbinfnode(prob,nodeInfeasible,NULL,0);
XPRSmipoptimize(prob,"");
```

The callback function may resemble:

```
void XPRS_CC nodeInfeasible(XPRSprob prob, void *data)
{
  int node;
  XPRSgetintattrib(prob, XPRS_CURRENTNODE, &node);
  printf("Node %d infeasible\n", node);
}
```

See the example `depthfirst.c` in the `examples/optimizer/c` folder.

Further information

Nodes that are cut off while *not* being the active node, are reported in the `nodecutoff` callback.

Related topics

`XPRSremovecbinfnode`, `XPRSaddcboptnode`, `XPRSaddcbintsol`, `XPRSaddcbnodecutoff`.

XPRSaddcbintsol

Purpose

Declares a user integer solution callback function, called every time an integer solution is found by heuristics or during the Branch and Bound search. This callback function will be called in addition to any callbacks already added by XPRSaddcbintsol.

Topic areas

Callback, Solution

Synopsis

```
int XPRS_CC XPRSaddcbintsol(XPRSprob prob, void (XPRS_CC *intsol)(XPRSprob
       cbprob, void *cbdata), void *data, int priority);
```

Arguments

`prob`	The current problem.
`intsol`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has no return value. This function is called if the current node is found to have an integer feasible solution, i.e. every time an integer feasible solution is found.
`cbprob`	The problem passed to the callback function, `intsol`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbintsol`.
`data`	A user-defined data to be passed to the callback function, `intsol`.
`priority`	An integer that determines the order in which multiple integer solution callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

The following example prints integer solutions as they are discovered in the tree search:

```
XPRSaddcbintsol(prob,printsol,NULL,0);
XPRSmipoptimize(prob,"");
```

The callback function might resemble:

```
void XPRS_CC printsol(XPRSprob cbprob, void *data)
{
  int i, cols;
  double objval, *x;

  XPRSgetintattrib(cbprob, XPRS_INPUTCOLS, &cols);
  XPRSgetdblattrib(cbprob, XPRS_LPOBJVAL, &objval);
  x = malloc(cols * sizeof(double));
  if (!x) return;
  XPRSgetcallbacksolution(cbprob, NULL, x, 0, cols - 1);

  printf("\nInteger solution found: %f\n", objval);
  for(i=0;i<cols;i++) printf(" x[%d] = %d\n", i, x[i]);
  free(x);
}
```

See also examples `roundint.c`, `savesol.c` .

Further information

1. This callback is useful if the user wants to retrieve the integer solution when it is found.
2. To retrieve the integer solution, use `XPRSgetcallbacksolution`, `XPRSgetcallbackpresolvesolution` or `XPRSgetsolution`.
3. This callback is called after a new integer solution was found by the Optimizer. Use a callback set by `XPRSaddcbpreintsol` in order to be notified *before* a new integer solution is accepted by the Optimizer, which allows for the new solution to be rejected.
4. For this callback attribute `MIPINFEAS` is set to -1 if the solution was found by a heuristic.

Related topics

`XPRSremovecbintsol`, `XPRSaddcbpreintsol`.

XPRSaddcblplog

Purpose

Declares a simplex log callback function which is called after every `LPLOG` iterations of the simplex algorithm. This callback function will be called in addition to any callbacks already added by XPRSaddcblplog.

Topic areas

Callback, LP

Synopsis

```
int XPRS_CC XPRSaddcblplog(XPRSprob prob, int (XPRS_CC *lplog)(XPRSprob
       cbprob, void* cbdata), void* data, int priority);
```

Arguments

`prob`	The current problem.
`lplog`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has an integer return value. This function is called every `LPLOG` simplex iterations including iteration `0` and the final iteration.
`cbprob`	The problem passed to the callback function, `lplog`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcblplog`.
`data`	A user-defined data to be passed to the callback function, `lplog`.
`priority`	An integer that determines the order in which multiple lplog callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related controls

Integer

`LPLOG`	Frequency and type of simplex algorithm log.

Example

The following code sets a callback function, `lpLog`, to be called every `10` iterations of the optimization:

```
XPRSsetintcontrol(prob,XPRS_LPLOG,10);
XPRSsetintcontrol(prob,XPRS_LPLOGSTYLE,0);
XPRSaddcblplog(prob,lpLog,NULL,0);
XPRSreadprob(prob,"problem","");
XPRSmipoptimize(prob,"");
```

The callback function may resemble:

```
int XPRS_CC lpLog(XPRSprob cbprob, void *data)
{
  int iter; double obj;

  XPRSgetintattrib(cbprob, XPRS_SIMPLEXITER, &iter);
  XPRSgetdblattrib(cbprob, XPRS_LPOBJVAL, &obj);
  printf("At iteration %d objval is %g\n", iter, obj);
  return 0;
}
```

Further information

If the callback function returns a nonzero value the solution process will be interrupted.

Related topics

`XPRSremovecblplog`, `XPRSaddcbbarlog`, `XPRSaddcbmiplog`, `XPRSaddcbmessage`.

XPRSaddcbmessage

Purpose

Declares an output callback function, called every time a text line relating to the given XPRSprob is output by the Optimizer. This callback function will be called in addition to any callbacks already added by XPRSaddcbmessage. Note that Optimizer messages passed to the callback do not end with a newline character; the user callback is expected to append such required newline characters itself.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSaddcbmessage(XPRSprob prob, void (XPRS_CC
       *message)(XPRSprob cbprob, void *cbdata, const char *msg, int msglen,
       int msgtype), void *data, int priority);
```

Arguments

`prob`	The current problem.
`message`	The callback function which takes five arguments, `cbprob`, `cbdata`, `msg`, `msglen` and `msgtype`, and has no return value. Use a `NULL` value to cancel a callback function.
`cbprob`	The problem passed to the callback function.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbmessage`.
`msg`	A null terminated character array (string) containing the message, which may simply be a new line. The total number of bytes (including `NUL` terminator) will not exceed `XPRS_MAXMESSAGELENGTH`. If a message needs to be truncated to meet this limit, the last four bytes in `msg` are set to "...\0".
`msglen`	The length of the message string, excluding the null terminator.
`msgtype`	Indicates the type of output message: `1` information messages; `2` (not used); `3` warning messages; `4` error messages. A negative value indicates that the Optimizer is about to finish and the buffers should be flushed at this time if the output is being redirected to a file.
`data`	A user-defined data to be passed to the callback function.
`priority`	An integer that determines the order in which callbacks of this type will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related controls

Integer

`OUTPUTLOG` All messages are disabled if set to zero.

Example

The following example simply sends all output to the screen (`stdout`):

```
XPRSaddcbmessage(prob,Message,NULL,0);
```

The callback function might resemble:

```
void XPRS_CC Message(XPRSprob cbprob, void* data,
                     const char *msg, int msglen, int msgtype)
{
  switch(msgtype)
  {
    case 4:  /* error */
```

```
    case 3:  /* warning */
    case 2:  /* not used */
    case 1:  /* information */
             printf("%s\n", msg);
             break;
    default: /* exiting - buffers need flushing */
             fflush(stdout);
             break;
  }
}
```

See also examples addmipsol.c, ComplexUserFunctions.c, els_managedcuts.c, els_usercuts.c, fixbv.c, globjpar.c, glrhspar.c, goalprog.c, knapsack.c, loadlp.c, mipsolenum.c, mipsolpool.c, mostviolated.c, Polygon_initialvalue.c, Polygon_textformula.c, Polygon_tokens.c, Polygon_userfunc.c, Polygon_userfunc_map.c, Polygon_userfunc_mapdelta.c, Polygon_userfunc_multimap.c, Polygon_userfunc_multimapdelta.c, Polygon_userfunc_vecmap.c, Polygon_userfunc_vecmapdelta.c, repair.c, roundint.c, savesol.c, tableau.c, trimloss.c, tsp.c .

Further information

1. Screen output is automatically created by the Optimizer Console only. To produce output when using the Optimizer library, it is necessary to define this callback function and use it to print the messages to the screen (`stdout`).
2. This function offers one method of handling the messages which describe any warnings and errors that may occur during execution. Other methods are to check the return values of functions and then get the error code using the `ERRORCODE` attribute, obtain the last error message directly using `XPRSgetlasterror`, or send messages direct to a log file using `XPRSsetlogfile`.
3. Visual Basic users must use the alternative function `XPRSaddcbmessageVB` to define the callback; this is required because of the different way VB handles strings.

Related topics

`XPRSremovecbmessage`, `XPRSaddcbbarlog`, `XPRSaddcbmiplog`, `XPRSaddcblplog`, `XPRSsetlogfile`.

XPRSaddcbmipthread

Purpose

Declares a MIP thread callback function, called every time a MIP worker problem is created by the parallel MIP code. This callback function will be called in addition to any callbacks already added by XPRSaddcbmipthread.

Topic areas

Callback, Parallel

Synopsis

```
int XPRS_CC XPRSaddcbmipthread(XPRSprob prob, void (XPRS_CC
       *mipthread)(XPRSprob cbprob, void *cbdata, XPRSprob threadprob), void
       *data, int priority);
```

Arguments

`prob`	The current problem.
`mipthread`	The callback function which takes three arguments, `cbprob`, `cbdata` and `threadprob`, and has no return value.
`cbprob`	The problem passed to the callback function.
`cbdata`	The user-defined data passed to the callback function.
`threadprob`	The problem pointer for the MIP thread
`data`	A user-defined data to be passed to the callback function.
`priority`	An integer that determines the order in which multiple callbacks of this type will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related controls

Integer

`MIPTHREADS`	Number of MIP threads to create.

Example

The following example clears the message callback for each of the MIP threads:

```
XPRSaddcbmipthread(prob,mipthread,NULL, 0);
```

```
void XPRS_CC mipthread(XPRSprob cbprob, void* cbdata,
                        XPRSprob mipthread)
{
  /* clear the message callback*/
  XPRSremovecbmessage(cbprob, mipthread, NULL);

}
```

Further information

This function will be called when a new MIP worker problem is created. Each worker problem receives a unique identifier that can be obtained through the `MIPTHREADID` attribute. Worker problems can be matched with different system threads at different points of a solve, so the system thread that is responsible for executing the callback is not necessarily the same thread used for all subsequent callbacks for the same worker problem. On the other hand, worker problems are always assigned to a single thread at a time and the same nodes are always solved on the same worker problem in repeated runs of a deterministic MIP solve. A worker problem therefore acts as a virtual thread through the node solves.

Related topics

`XPRSremovecbmipthread`,`XPRSaddcbdestroymt`, `MIPTHREADS`, `MAXMIPTASKS`.

XPRSaddcbnewnode

Purpose

Declares a callback function that will be called every time a new node is created during the branch and bound search. This callback function will be called in addition to any callbacks already added by XPRSaddcbnewnode.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSaddcbnewnode(XPRSprob prob, void (XPRS_CC
      *newnode)(XPRSprob cbprob, void* cbdata, int parentnode, int node,
      int branch), void* data, int priority);
```

Arguments

`prob`	The current problem.
`newnode`	The callback function, which takes five arguments: `myprob`, `cbdata`, `parentnode`, `node` and `branch`. This function is called every time a new node is created through branching.
`cbprob`	The problem passed to the callback function, `newnode`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbnewnode`.
`parentnode`	Unique identifier for the parent of the new node.
`node`	Unique identifier assigned to the new node.
`branch`	The sequence number of the new node amongst the child nodes of `parentnode`. For regular branches on a MIP entity this will be either `0` or `1`.
`data`	A user-defined data to be passed to the callback function.
`priority`	An integer that determines the order in which callbacks of this type will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Further information

1. For regular branches on a MIP entity, `branch` will be either zero or one, depending on whether the new node corresponds to branching the MIP entity up or down.
2. When branching on an `XPRSbranchobject`, `branch` refers to the given branch index of the data.

Related topics

`XPRSremovecbnewnode`, `XPRSaddcbchgnode`.

XPRSaddcbnodecutoff

Purpose

Declares a user node cutoff callback function, called every time a node is cut off as a result of an improved integer solution being found during the branch and bound search. This callback function will be called in addition to any callbacks already added by XPRSaddcbnodecutoff.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSaddcbnodecutoff(XPRSprob prob, void (XPRS_CC
       *nodecutoff)(XPRSprob cbprob, void *cbdata, int node), void *data,
       int priority);
```

Arguments

`prob`	The current problem.
`nodecutoff`	The callback function, which takes three arguments, `cbprob`, `cbdata` and `node`, and has no return value. This function is called every time a node is cut off as the result of an improved integer solution being found.
`cbprob`	The problem passed to the callback function, `nodecutoff`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbnodecutoff`.
`node`	The node id of the node that is cut off. This id cannot be queried from the `CURRENTNODE` attribute as for other callbacks since this callback is not invoked in the context of the node being cutoff.
`data`	A user-defined data to be passed to the callback function, `nodecutoff`.
`priority`	An integer that determines the order in which multiple node-optimal callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

The following notifies the user whenever a node is cutoff during the tree search:

```
XPRSaddcbnodecutoff(prob,Cutoff,NULL,0);
XPRSmipoptimize(prob,"");
```

The callback function might resemble:

```
void XPRS_CC Cutoff(XPRSprob prob, void *data, int node)
{
  printf("Node %d cutoff\n", node);
}
```

See the example `depthfirst.c` in the `examples/optimizer/c` folder.

Further information

1. This function allows the user to keep track of the eligible nodes. Note that the LP solution will not be available from this callback.
2. This callback is invoked only for nodes that are pruned from the active node pool. If the *current* node gets cut off then the `infnode` callback is invoked to report this.

Related topics

`XPRSremovecbnodecutoff`, `XPRSaddcboptnode`, `XPRSaddcbinfnode`, `XPRSaddcbintsol`.

XPRSaddcbnodelpsolved

Purpose

Declares a callback function, called during the branch and bound search, after the LP relaxation has been solved for the current node, but before any internal cuts and heuristics have been applied. This callback function will be called in addition to any callbacks already added by XPRSaddcbnodelpsolved.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSaddcbnodelpsolved(XPRSprob prob, void (XPRS_CC
      *nodelpsolved)(XPRSprob cbprob, void *cbdata), void *data, int
      priority);
```

Arguments

`prob`	The current problem.
`nodelpsolved`	The callback function which takes two arguments, `cbprob` and `cbdata`, and has no return value.
`cbprob`	The problem passed to the callback function, `nodelpsolved`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddnodelpsolved`.
`data`	A user-defined data to be passed to the callback function, `nodelpsolved`.
`priority`	An integer that determines the order in which multiple node-optimal callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related topics

`XPRSremovecbnodelpsolved`, `XPRSaddcboptnode`,

XPRSaddcboptnode

Purpose

Declares an optimal node callback function, called during the branch and bound search, after the LP relaxation has been solved for the current node, and after any internal cuts and heuristics have been applied, but before the Optimizer checks if the current node should be branched. This callback function will be called in addition to any callbacks already added by XPRSaddcboptnode.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSaddcboptnode(XPRSprob prob, void (XPRS_CC
      *optnode)(XPRSprob cbprob, void *cbdata, int *p_infeasible), void
      *data, int priority);
```

Arguments

`prob`	The current problem.
`optnode`	The callback function which takes three arguments, `cbprob`, `cbdata` and `p_infeasible`, and has no return value.
`cbprob`	The problem passed to the callback function, `optnode`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcboptnode`.
`p_infeasible`	The feasibility status. If set to a nonzero value by the user, the current node will be declared infeasible.
`data`	A user-defined data to be passed to the callback function, `optnode`.
`priority`	An integer that determines the order in which multiple node-optimal callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

The following prints the optimal objective value of the node LP relaxations:

```
XPRSaddcboptnode(prob,nodeOptimal,NULL,0);
XPRSmipoptimize(prob,"");
```

The callback function might resemble:

```
void XPRS_CC nodeOptimal(XPRSprob prob, void *data, int *p_infeasible)
{
  int node;
  double objval;

  XPRSgetintattrib(prob, XPRS_CURRENTNODE, &node);
  printf("NodeOptimal: node number %d\n", node);
  XPRSgetdblattrib(prob, XPRS_LPOBJVAL, &objval);
  printf("\tObjective function value = %f\n", objval);
}
```

See the example `depthfirst.c` in the `examples/optimizer/c` folder.

See also examples `addmipsol.c`, `knapsack.c` .

Related topics

`XPRSremovecboptnode`, `XPRSaddcbinfnode`, `XPRSaddcbintsol`, `XPRSaddcbnodecutoff`, `CALLBACKCOUNT_OPTNODE`.

XPRSaddcbpreintsol

Purpose

Declares a user integer solution callback function, called when an integer solution is found by heuristics or during the branch and bound search, but before it is accepted by the Optimizer. This callback allows the user to accept or reject a candidate MIP solution before it is accepted as the new incumbent, and in some cases allows the user to modify the node problem (see notes below). This callback function will be called in addition to any integer solution callbacks already added by XPRSaddcbpreintsol.

Topic areas

Callback, Solution

Synopsis

```
int XPRS_CC XPRSaddcbpreintsol(XPRSprob prob, void (XPRS_CC
       *preintsol)(XPRSprob cbprob, void *cbdata, int soltype, int
       *p_reject, double *p_cutoff), void *data, int priority);
```

Arguments

`prob`	The current problem.
`preintsol`	The callback function which takes five arguments, `cbprob`, `cbdata`, `soltype`, `p_reject` and `p_cutoff`, and has no return value. This function is called when an integer solution is found, but before the solution is accepted by the Optimizer, allowing the user to reject the solution.
`cbprob`	The problem passed to the callback function, `preintsol`.
`cbdata`	The user-defined data passed as data when setting up the callback with `XPRSaddcbpreintsol`.
`soltype`	The type of MIP solution that has been found: `0` The continuous relaxation solution to the current node of the tree search, which has been found to be integer feasible. `1` A MIP solution found by a heuristic. `2` A MIP solution provided by the user.
`p_reject`	Set this to 1 if the solution should be rejected. When `soltype` is zero, this will also drop the node problem.
`p_cutoff`	The new cutoff value that the Optimizer will use if the solution is accepted. If the user changes `p_cutoff`, the new value will be used instead. The cutoff value will not be updated if the solution is rejected.
`data`	A user-defined data to be passed to the callback function, `preintsol`.
`priority`	An integer that determines the order in which callbacks of this type will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Related controls

Integer

`MIPABSCUTOFF`	Branch and Bound: If the user knows that they are interested only in values of the objective function which are better than some value, this can be assigned to `MIPABSCUTOFF`.

Further information

1. To retrieve the integer solution, use XPRSgetcallbacksolution or XPRSgetcallbackpresolvesolution. The matching objective value should be queried through the LPOBJVAL attribute. XPRSgetsolution and MIPOBJVAL will not return the newly found solution and matching objective value, because the incumbent solution has not been updated at this point.
2. When soltype is zero, the current node problem can be modified in this callback, including adding cuts, changing bounds and adding branching objects. If the problem is changed, the node LP relaxation will be resolved and checked for branching candidates again. This might retrigger any optnode (ref. XPRSaddcboptnode) callback and the node will be branched, if required. Making any changes when soltype is not zero will raise an error.
3. When soltype is zero, it is possible to add new solutions using XPRSaddmipsol. If the node problem was not modified during the callback, any added solution will be checked after the current solution has been accepted as a new incumbent. Otherwise, the added solution will be checked after the node relaxation has been resolved.
4. When soltype is zero and the solution is rejected, the node itself will be dropped without further branching, regardless of whether the node problem is modified within the callback.
5. The cutoff value returned in p_cutoff is applied only if the solution is not rejected and the node problem (when soltype is zero) is not modified. To change the cutoff value when rejecting a solution, the control MIPABSCUTOFF should be set instead.
6. For this callback attribute MIPINFEAS is set to -1 if the solution was found by a heuristic.

Related topics

XPRSremovecbpreintsol, XPRSaddcboptnode, XPRSaddcbintsol.

Example

See also example tsp.c.

XPRSaddcbprenode

Purpose

Declares a preprocess node callback function, called before the LP relaxation of a node has been optimized, so the solution at the node will not be available. This callback function will be called in addition to any callbacks already added by XPRSaddcbprenode.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSaddcbprenode(XPRSprob prob, void (XPRS_CC
       *prenode)(XPRSprob cbprob, void *cbdata, int *p_infeasible), void
       *data, int priority);
```

Arguments

`prob`	The current problem.
`prenode`	The callback function, which takes three arguments, `cbprob`, `cbdata` and `p_infeasible`, and has no return value. This function is called before a node is reoptimized and the node may be made infeasible by setting `*p_infeasible` to `1`.
`cbprob`	The problem passed to the callback function, `prenode`.
`cbdata`	The user-defined data passed as `data` when setting up the callback with `XPRSaddcbprenode`.
`p_infeasible`	The feasibility status. If set to a nonzero value by the user, the current node will be declared infeasible by the Optimizer.
`data`	A user-defined data to be passed to the callback function, `prenode`.
`priority`	An integer that determines the order in which multiple preprocess node callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to 0 if not required.

Example

The following example notifies the user before each node is processed:

```
XPRSaddcbprenode(prob, preNode, NULL, 0);
XPRSmipoptimize(prob,"");
```

The callback function might resemble:

```
void XPRS_CC preNode(XPRSprob prob, void* data, int *p_infeasible)
{
 *p_infeasible = 0; /* set to 1 if node is infeasible */
}
```

Related topics

`XPRSremovecbprenode`, `XPRSaddcbchgnode`, `XPRSaddcbinfnode`, `XPRSaddcbintsol`, `XPRSaddcbnodecutoff`, `XPRSaddcboptnode`.

XPRSaddcbusersolnotify

Purpose

Declares a callback function to be called each time a solution added by `XPRSaddmipsol` has been processed. This callback function will be called in addition to any callbacks already added by `XPRSaddcbusersolnotify`.

Topic areas

Callback, Solution

Synopsis

```
int XPRS_CC XPRSaddcbusersolnotify(XPRSprob prob, void (XPRS_CC
      *usersolnotify)(XPRSprob cbprob, void* cbdata, const char* solname,
      int status), void* data, int priority);
```

Arguments

`prob`	The current problem.
`usersolnotify`	The callback function which takes four arguments, `cbprob`, `cbdata`, `solname` and `status` and has no return value.
`cbprob`	The problem passed to the callback function, `usersolnotify`.
`cbdata`	The user-defined data passed as data when setting up the callback with `XPRSaddcbusersolnotify`.
`solname`	The string name assigned to the solution when it was loaded into the Optimizer using `XPRSaddmipsol`.
`status`	One of the following status values: 0 An error occurred while processing the solution. (`XPRS_USERSOLSTATUS_ERROR`) 1 Solution is feasible. (`XPRS_USERSOLSTATUS_ACCEPTED_FEASIBLE`) 2 Solution is feasible after reoptimizing with fixed MIP entities. (`XPRS_USERSOLSTATUS_ACCEPTED_OPTIMIZED`) 3 A local search heuristic was applied and a feasible solution discovered. (`XPRS_USERSOLSTATUS_SEARCHED_SOL`) 4 A local search heuristic was applied but a feasible solution was not found. (`XPRS_USERSOLSTATUS_SEARCHED_NOSOL`) 5 Solution is infeasible and a local search could not be applied. (`XPRS_USERSOLSTATUS_REJECTED_INFEAS_NOSEARCH`) 6 Solution is partial and a local search could not be applied. (`XPRS_USERSOLSTATUS_REJECTED_PARTIAL_NOSEARCH`) 7 Failed to reoptimize the problem with MIP entities fixed to the provided solution. Likely because a time or iteration limit was reached. (`XPRS_USERSOLSTATUS_REJECTED_FAILED_OPTIMIZE`) 8 Solution is dropped. This can happen if the MIP problem is changed or solved to completion before the solution could be processed. (`XPRS_USERSOLSTATUS_DROPPED`) 9 The solution is worse than the current MIP cutoff value. (`XPRS_USERSOLSTATUS_REJECTED_CUTOFF`)
`data`	A user-defined data to be passed to the callback function, `usersolnotify`.
`priority`	An integer that determines the order in which multiple callbacks will be invoked. The callback added with a higher priority will be called before a callback with a lower priority. Set to `0` if not required.

Further information

If presolve is turned on, any solution added with `XPRSaddmipsol` will first be presolved before it can be checked. The value returned in `status` refers to the presolved solution, which might have had values adjusted due to bound changes, fixing of variables, etc.

Related topics

`XPRSremovecbusersolnotify`, `XPRSaddmipsol`.

XPRSaddcols, XPRSaddcols64

Purpose

Adds columns to the optimizer matrix.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSaddcols(XPRSprob prob, int ncols, int ncoefs, const double
      objcoef[], const int start[], const int rowind[], const double
      rowcoef[], const double lb[], const double ub[]);

int XPRS_CC XPRSaddcols64(XPRSprob prob, int ncols, XPRSint64 ncoefs, const
      double objcoef[], const XPRSint64 start[], const int rowind[], const
      double rowcoef[], const double lb[], const double ub[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of new columns.
`ncoefs`	Number of new nonzeros in the added columns.
`objcoef`	Double array of length `ncols` containing the objective function coefficients of the new columns. As of version 43.00 this can be `NULL` to indicate that all newly created variables have an objective coefficient of 0 (zero).
`start`	Integer array of length `ncols` containing the offsets in the `rowind` and `rowcoef` arrays of the start of the elements for each column. This can be `NULL` if the `ncols` or the `ncoefs` argument is 0 (zero).
`rowind`	Integer array of length `ncoefs` containing the row indices for the elements in each column.
`rowcoef`	Double array of length `ncoefs` containing the element values.
`lb`	Double array of length `ncols` containing the lower bounds on the added columns. This can be `NULL` in which case all variables will have a lower bound of 0.
`ub`	Double array of length `ncols` containing the upper bounds on the added columns. This can be `NULL` in which hcase all variables will have an infinite upper bound.

Related controls

Integer

`EXTRACOLS`	Number of extra columns to be allowed for.
`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAMIPENTS`	Number of extra MIP entities to be allowed for.

Double

`MATRIXTOL`	Tolerance on matrix elements.

Example

In this example, we consider the two problems:

(a)	maximize:	$2x + y$			(b)	maximize:	$2x + y + 3z$		
	subject to:	$x + 4y$	$\leqslant$	24		subject to:	$x + 4y + 2z$	$\leqslant$	24
		y	$\leqslant$	5			$y + z$	$\leqslant$	5
		$3x + y$	$\leqslant$	20			$3x + y$	$\leqslant$	20
		$x + y$	$\leqslant$	9			$x + y + 3z$	$\leqslant$	9
							z	$\leqslant$	12

Using XPRSaddcols, the following transforms (a) into (b) and then names the new variable using XPRSaddnames:

```
obj[0] = 3;
start[] = {0};
rowind[] = {0, 1, 3};
matval[] = {2.0, 1.0, 3.0};
lb[0] = XPRS_MINUSINFINITY; ub[0] = 12.0;
...
XPRSaddcols(prob,1,3,obj,start,rowind,matval,lb,ub);
XPRSaddnames(prob,XPRS_NAMES_COLUMN,"z",2,2);
```

See also examples els_managedcuts.c, els_usercuts.c, goalprog.c, trimloss.c, tsp.c .

Further information

1. The double constants XPRS_PLUSINFINITY and XPRS_MINUSINFINITY defined in the library header file can be used to represent plus and minus infinity respectively in the bound arrays.
2. If the columns are added to a MIP problem then they will be continuous variables. Use XPRSchgcoltype to impose integrality conditions on such new columns.
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use XPRSpostsolve to restore the problem to a non-solving state.

Related topics

XPRSaddnames, XPRSaddrows, XPRSdelcols, XPRSchgcoltype.

XPRSaddcuts, XPRSaddcuts64

Purpose

Adds cuts directly to the matrix at the current node. The cuts will automatically be added to the cut pool. Cuts added to a node will automatically be inherited on any descendant node, unless explicitly deleted with a call to `XPRSdelcuts`.

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSaddcuts(XPRSprob prob, int ncuts, const int cuttype[],
       const char rowtype[], const double rhs[], const int start[], const
       int colind[], const double cutcoef[]);

int XPRS_CC XPRSaddcuts64(XPRSprob prob, int ncuts, const int cuttype[],
       const char rowtype[], const double rhs[], const XPRSint64 start[],
       const int colind[], const double cutcoef[]);
```

Arguments

`prob`	The current problem.
`ncuts`	Number of cuts to add.
`cuttype`	Integer array of length `ncuts` containing the user assigned cut types. The cut types can be any integer chosen by the user, and are used to identify the cuts in other cut manager routines using user supplied parameters. The cut type can be interpreted as an integer or a bitmap - see `XPRSdelcuts`.
`rowtype`	Character array of length `ncuts` containing the row types: `L` indicates a $\leqslant$ row; `G` indicates a $\geqslant$ row; `E` indicates an = row.
`rhs`	Double array of length `ncuts` containing the right hand side elements for the cuts.
`start`	Integer array containing offset into the `colind` and `cutcoef` arrays indicating the start of each cut. This array is of length `ncuts+1` with the last element, `start[ncuts]`, being where cut `ncuts+1` would start.
`colind`	Integer array of length `start[ncuts]` containing the column indices in the cuts.
`cutcoef`	Double array of length `start[ncuts]` containing the matrix values for the cuts.

Further information

1. The columns and elements of the cuts must be stored contiguously in the `colind` and `cutcoef` arrays passed to `XPRSaddcuts`. The starting point of each cut must be stored in the `start` array. To determine the length of the final cut, the `start` array must be of length `ncuts+1` with the last element of this array containing the position in `colind` and `cutcoef` where the cut `ncuts+1` would start. `start[ncuts]` denotes the number of nonzeros in the added cuts.

2. The cuts added to the matrix are always added at the end of the matrix and the number of rows is always set to the original number rows plust the number of cuts added. If `ncuts` have been added, then the rows `0`,...,`ROWS-ncuts-1` are the original rows, whilst the rows `ROWS-ncuts`,...,`ROWS-1` are the added cuts. The number of cuts can be found by consulting the `CUTS` problem attribute.

3. This function should be called only from within callback functions set by either `XPRSaddcbnodelpsolved`, `XPRSaddcboptnode`, `XPRSaddcbpreintsol` or `XPRSaddcbprenode`.

Related topics

`XPRSaddrows`, `XPRSdelcpcuts`, `XPRSdelcuts`, `XPRSgetcpcutlist`, `XPRSgetcutlist`, `XPRSloadcuts`, `XPRSstorecuts`, Section 5.9.

Example

See also examples `els_usercuts.c`, `tsp.c` .

XPRSaddgencons, XPRSaddgencons64

Purpose

Adds one or more general constraints to the problem. Each general constraint `y = f(x1, ..., xn, c1, ..., cn)` consists of one or more (input) columns xi, zero or more constant values ci and a resultant (output column) y, different from all xi. General constraints include `maximum` and `minimum` (arbitrary number of input columns of any type and arbitrary number of input values, at least one total), `and` and `or` (at least one binary input column, no constant values, binary resultant) and `absolute value` (exactly one input column of arbitrary type, no constant values).

Topic areas

Problem Creation, Piecewise Linear and General Constraints

Synopsis

```
int XPRS_CC XPRSaddgencons(XPRSprob prob, int ncons, int ncols, int nvals,
       const int contype[], const int resultant[], const int colstart[],
       const int colind[], const int valstart[], const double val[]);

int XPRS_CC XPRSaddgencons64(XPRSprob prob, int ncons, XPRSint64 ncols,
       XPRSint64 nvals, const int contype[], const int resultant[], const
       XPRSint64 colstart[], const int colind[], const XPRSint64 valstart[],
       const double val[]);
```

Arguments

`prob`	The current problem.
`ncons`	The number of general constraints to add.
`ncols`	The total number of input variables in general constraints that should be added.
`nvals`	The total number of constant values in general constraints that should be added.
`contype`	Integer array of length `ncons` containing the types of the general constraints: `XPRS_GENCONS_MAX (0)` indicates a `maximum` constraint; `XPRS_GENCONS_MIN (1)` indicates a `minimum` constraint; `XPRS_GENCONS_AND (2)` indicates an `and` constraint. `XPRS_GENCONS_OR (3)` indicates an `or` constraint; `XPRS_GENCONS_ABS (4)` indicates an `absolute value` constraint.
`resultant`	Integer array of length `ncons` containing the indices of the output variables of the general constraints.
`colstart`	Integer array of length `ncons` containing the start index of each general constraint in the `colind` array.
`colind`	Integer array of length `ncols` containing the input variables in all general constraints.
`valstart`	Integer array of length `ncons` containing the start index of each general constraint in the `val` array (may be `NULL` if `ncoefs = 0`).
`val`	Double array of length `nvals` containing the constant values in all general constraints (may be `NULL` if `ncoefs = 0`).

Example

This adds two new general constraints `x2 = max(x0, x1, 5)` and `x3 = |x1|`:

```
int contype[] = {XPRS_GENCONS_MAX, XPRS_GENCONS_ABS};
int resultant[] = {2, 3};
int colstart[] = {0, 2};
int colind[] = {0, 1, 1};
int valstart[] = {0, 1};
double val[] = {5.0};
```

```
...
XPRSaddgencons(prob, 2, 3, 1, contype, resultant, colstart, colind, valstart, val);
XPRSmipoptimize(prob,"");
```

Further information

1. General constraints must be set up before solving the problem. They are converted to additional binary variables, indicator and linear constraints with the exact formulation and number of added entities depending on the performed presolving.
2. Note that using non-binary variables in `and`/`or` constraints or adding constant values to them or `absolute value` constraints will give an error at solve time.
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related controls

Integer

`GENCONSDUALREDUCTIONS` Controls whether dual reductions may be applied to reduce the number of added variables and constraints.

Related topics

`XPRSgetgencons`, `XPRSdelgencons`.

XPRSaddmanagedcuts, XPRSaddmanagedcuts64

Purpose

Adds cuts to the Optimizer's internal cut pool from within the *cutround* callback set by `XPRSaddcbcutround`. The cuts will be added to an internal pool of cuts managed by the Optimizer. The Optimizer will use internal priorities to dynamically load violated cuts from this pool into branch-and-bound node problems and remove inactive cuts. Cuts can be either *local* or *global*. Cuts flagged as *local* are assumed to be valid only for the the current node of the branch-and-bound search or any of its descendants. Global cuts are assumed to be valid for the whole problem and might be used on any node of the branch-and-bound search tree. The cuts should be formulated in the original space of variables and will automatically be presolved.

Topic areas

Callback, Cuts

Synopsis

```
int XPRS_CC XPRSaddmanagedcuts(XPRSprob prob, int globalvalid, int ncuts,
       const char rowtype[], const double rhs[], const int start[], const
       int colind[], const double cutcoef[]);

int XPRS_CC XPRSaddcuts64(XPRSprob prob, int globalvalid, int ncuts, const
       char rowtype[], const double rhs[], const XPRSint64 start[], const
       int colind[], const double cutcoef[]);
```

Arguments

`prob`	The current problem.
`globalvalid`	Nonzero if the cuts should be assumed to be valid for the whole problem. If zero, cuts will be assumed to be valid only for the current node and its descendants.
`ncuts`	Number of cuts to add.
`rowtype`	Character array of length `ncuts` containing the row types: `L` indicates a $\leqslant$ row; `G` indicates a $\geqslant$ row; `E` indicates an = row.
`rhs`	Double array of length `ncuts` containing the right hand side elements for the cuts.
`start`	Integer array containing offset into the `colind` and `cutcoef` arrays indicating the start of each cut. This array is of length `ncuts+1` with the last element, `start[ncuts]`, being where cut `ncuts+1` would start.
`colind`	Integer array of length `start[ncuts]` containing the column indices in the cuts.
`cutcoef`	Double array of length `start[ncuts]` containing the matrix values for the cuts.

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. The columns and elements of the cuts must be stored contiguously in the `colind` and `cutcoef` arrays passed to `XPRSaddmanagedcuts`. The starting point of each cut must be stored in the `start` array. To determine the length of the final cut, the `start` array must be of length `ncuts+1` with the last element of this array containing the position in `colind` and `cutcoef` where the cut `ncuts+1` would start. `start[ncuts]` denotes the number of nonzeros in the added cuts.
3. Cuts will not be loaded into the problem if `CUTSTRATEGY` is set to zero.
4. Cuts added by `XPRSaddmanagedcuts` cannot be retrieved by `XPRSgetcutlist` or `XPRSgetcpcutlist`.
5. This function can only be called from a callback function set by `XPRSaddcbcutround`.

Related topics

`XPRSaddcbcutround`, `XPRSaddcuts`, `CUTSTRATEGY`, Section 5.9.

Example

See also example `els_managedcuts.c`.

XPRSaddmipsol

Purpose

Adds a new feasible, infeasible or partial MIP solution for the problem to the Optimizer.

Topic area

Solution

Synopsis

```
int XPRS_CC XPRSaddmipsol(XPRSprob prob, int length, const double solval[],
        const int colind[], const char* name);
```

Arguments

`prob`	The current problem.
`length`	Number of columns for which a value is provided.
`solval`	Double array of length `length` containing solution values.
`colind`	Optional integer array of length `length` containing the column indices for the solution values provided in `solval`. Should be `NULL` when `length` is equal to `COLS`, in which case it is assumed that `solval` provides a complete solution vector.
`name`	An optional name to associate with the solution. Can be NULL.

Further information

1. The function returns immediately after passing the solution to the Optimizer. The solution is placed in a pool until the Optimizer is able to analyze the solution during a MIP solve.
2. If the provided solution is found to be infeasible, a limited local search heuristic will be run in an attempt to find a close feasible integer solution.
3. If a partial solution is provided, integer columns will be fixed to any provided values and a limited local search will be run in an attempt to find integer feasible values for the remaining unspecified columns. Values provided for continuous column in partial solutions are currently ignored.
4. The `XPRSaddcbusersolnotify` callback function can be used to discover the outcome of a loaded solution. The optional name provided as `name` will be returned in the callback function.
5. If one or more solutions are loaded during the `XPRSaddcboptnode` callback, the Optimizer will process all loaded solutions and fire the callback again. This will be repeated as long as new solutions are loaded during the callback. You can check the `CALLBACKCOUNT_OPTNODE` attribute if you only want to run once.

Related controls

Integer

`CALLBACKCOUNT_OPTNODE` Counts the number of times the `XPRSaddcboptnode` callback has been called, in particular after rerunning due to XPRSaddmipsol.

`USERSOLHEURISTIC` Controls the local search heuristic for an infeasible or partial solution.

Related topics

`XPRSaddcbusersolnotify`, `XPRSaddcboptnode`.

Example

See also examples `addmipsol.c`, `tsp.c` .

XPRSaddobj

Purpose

Appends an objective function with the given coefficients to a multi-objective problem. The weight and priority of the objective are set to the given values.

Topic areas

Problem Creation, Multiobjective

Synopsis

```
int XPRS_CC XPRSaddobj(XPRSprob prob, int ncols, const int colind[], const
       double objcoef[], int priority, double weight);
```

Arguments

`prob`	The current problem.
`ncols`	Number of objective function coefficient elements to add.
`colind`	Integer array of length `ncols` containing the indices of the columns whose objective coefficients will change. An index of `-1` indicates that the fixed part of the objective function on the right hand side should change.
`objcoef`	Double array of length `ncols` giving the new objective function coefficients.
`priority`	The priority for the objective function. During optimization, objectives with the same priority are combined together in a weighted sum.
`weight`	The weight for the objective function. If the weight is negative, the sense of this objective is reversed.

Example

Adding a second objective function to a problem:

```
colind[0] = 0; colind[1] = 2; colind[2] = 5;
objcoef[0] = 25.0; objcoef[1] = 5.3; objcoef[2] = 0.0;
XPRSaddobj(prob,3,colind,objcoef,1,1);
```

Related topics

`XPRSchgobjn`, `XPRSgetobjn`, `XPRSdelobj`, `XPRSchgobj`.

XPRSaddnames

Purpose

When a model is loaded, the rows, columns, sets, piecewise linear and general constraints of the model may not have names associated with them. This may not be important as the rows, columns, sets, piecewise linear and general constraints can be referred to by their sequence numbers. However, if you wish row, column, set, piecewise linear and general constraint names to appear in the ASCII solutions files, the names for a range of rows/columns/... can be added with `XPRSaddnames`.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSaddnames(XPRSprob prob, int type, const char names[], int
        first, int last);
```

Arguments

`prob`	The current problem.
`type`	`XPRS_NAMES_ROW (=1)` for row names; `XPRS_NAMES_COLUMN (=2)` for column names; `XPRS_NAMES_SET (=3)` for set names; `XPRS_NAMES_PWLCONS (=4)` for piecewise linear constraint names; `XPRS_NAMES_GENCONS (=5)` for general constraint names; `XPRS_NAMES_OBJECTIVE (=6)` for objective names.
`names`	Character buffer containing the null-terminated string names. The length of an individual name (including the terminating `NUL`) must not exceed 1MB.
`first`	Start of the range of rows, columns, sets, piecewise linear constraints, general constraints or objectives.
`last`	End of the range of rows, columns, sets, piecewise linear constraints, general constraints or objectives.

Example

Add variable names (`a` and `b`), and constraint names (`first`, `second`, and `third`) to a problem:

```
char rnames[] = "first\0second\0third"
char cnames[] = "a\0b";
...
XPRSaddnames(prob,XPRS_NAMES_ROW,rnames,0,nrow-1);
XPRSaddnames(prob,XPRS_NAMES_COLUMN,cnames,0,ncol-1);
```

See also examples `ComplexUserFunctions.c`, `els_managedcuts.c`, `els_usercuts.c`, `loadlp.c`, `Polygon_initialvalue.c`, `Polygon_textformula.c`, `Polygon_tokens.c`, `Polygon_userfunc.c`, `Polygon_userfunc_map.c`, `Polygon_userfunc_mapdelta.c`, `Polygon_userfunc_multimap.c`, `Polygon_userfunc_multimapdelta.c`, `Polygon_userfunc_vecmap.c`, `Polygon_userfunc_vecmapdelta.c`, `trimloss.c`, `tsp.c` .

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddcols`, `XPRSaddrows`, `XPRSgetnamelist`.

XPRSaddpwlcons, XPRSaddpwlcons64

Purpose

Adds one or more piecewise linear constraints to the problem. Each piecewise linear constraint `y = f(x)` consists of an (input) column x, a (different) resultant (output column) y and a piecewise linear function f. The piecewise linear function f is described by at least two breakpoints, which are given as combinations of x- and y-values. Discontinuous piecewise linear functions are supported, in this case both the left and right limit at a given point need to be entered as breakpoints. To differentiate between left and right limit, the breakpoints need to be given as a list with non-decreasing x-values.

Topic areas

Problem Creation, Piecewise Linear and General Constraints

Synopsis

```
int XPRS_CC XPRSaddpwlcons(XPRSprob prob, int npwls, int npoints, const int
      colind[], const int resultant[], const int start[], const double
      xval[], const double yval[]);

int XPRS_CC XPRSaddpwlcons64(XPRSprob prob, int npwls, XPRSint64 npoints,
      const int colind[], const int resultant[], const XPRSint64 start[],
      const double xval[], const double yval[]);
```

Arguments

`prob`	The current problem.
`npwls`	The number of piecewise linear constraints to add.
`npoints`	The total number of breakpoints of all piecewise linear constraints that should be added.
`colind`	Integer array of length `npwls` containing the indices of the input variables x of the piecewise linear functions.
`resultant`	Integer array of length `npwls` containing the indices of the output variables y of the piecewise linear functions.
`start`	Integer array of length `npwls` containing the start index of each piecewise linear constraint in the `xval` and `yval` arrays.
`xval`	Double array of length `npoints` containing the x-values of the breakpoints.
`yval`	Double array of length `npoints` containing the y-values of the breakpoints.

Example

The following example adds a new piecewise linear constraint `y = f(x)`, where

`f(x) = -x`	if `x < 0`
`f(x) = 1`	if `0 <= x <= 2`
`f(x) = 2x-3`	if `x > 2`

This function can be defined using the breakpoints `(x = -1, y = 1), (0,0), (0,1), (2,1), (3,3)` (note that the first breakpoint could also be replaced, e.g., by `(x = -2, y = 2)`, similarly for the last):

```
int colind[] = {0};
int resultant[] = {1};
int start[] = {0};
double xval[] = {-1,0,0,2,3};
double yval[] = {1,0,1,1,3};
...
XPRSaddpwlcons(prob,1,5,colind,resultant,start,xval,yval);
XPRSmipoptimize(prob,"");
```

Further information

1. Piecewise linear constraints must be set up before solving the problem. They are converted to additional linear constraints, continuous variables and SOS2 constraints, with the exact formulation and number of added entities depending on the convexity of the piecewise linear function and some presolving steps that are applied.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related controls

Integer

`PWLDUALREDUCTIONS` Controls whether dual reductions may be applied to reduce the number of added variables and constraints.

Related topics

`XPRSgetpwlcons`, `XPRSdelpwlcons`.

XPRSaddqmatrix, XPRSaddqmatrix64

Purpose

Adds a new quadratic matrix into a row defined by triplets.

Topic areas

Problem Creation, Quadratic

Synopsis

```
int XPRS_CC XPRSaddqmatrix(XPRSprob prob, int row, int ncoefs, const int
       rowqcol1[], const int rowqcol2[], const double rowqcoef[]);

int XPRS_CC XPRSaddqmatrix64(XPRSprob prob, int row, XPRSint64 ncoefs,
       const int rowqcol1[], const int rowqcol2[], const double rowqcoef[]);
```

Arguments

`prob`	The current problem.
`row`	Index of the row where the quadratic matrix is to be added.
`ncoefs`	Number of triplets used to define the quadratic matrix. This may be less than the number of coefficients in the quadratic matrix, since off diagonals and their transposed pairs are defined by one triplet.
`rowqcol1`	First index in the triplets.
`rowqcol2`	Second index in the triplets.
`rowqcoef`	Coefficients in the triplets.

Further information

1. The triplets should be filled to define the upper-triangular part of the quadratic expression. This means that to add $[x^2 + 6\ xy]$ the `rowqcoef` array shall contain the coefficients 1 and 3, respectively.

2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSloadqcqp`, `XPRSgetqrowcoeff`, `XPRSchgqrowcoeff`, `XPRSgetqrowqmatrix`, `XPRSgetqrowqmatrixtriplets`, `XPRSgetqrows`, `XPRSchgqobj`, `XPRSchgmqobj`, `XPRSgetqobj`.

XPRSaddrows, XPRSaddrows64

Purpose

Adds rows to the optimizer matrix.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSaddrows(XPRSprob prob, int nrows, int ncoefs, const char
      rowtype[], const double rhs[], const double rng[], const int start[],
      const int colind[], const double rowcoef[]);

int XPRS_CC XPRSaddrows64(XPRSprob prob, int nrows, XPRSint64 ncoefs, const
      char rowtype[], const double rhs[], const double rng[], const
      XPRSint64 start[], const int colind[], const double rowcoef[]);
```

Arguments

`prob`	The current problem.
`nrows`	Number of new rows.
`ncoefs`	Number of new nonzeros in the added rows.
`rowtype`	Character array of length nrows containing the row types: `L` indicates a $\leqslant$ row; `G` indicates $\geqslant$ row; `E` indicates an = row. `R` indicates a range constraint; `N` indicates a nonbinding constraint.
`rhs`	Double array of length `nrows` containing the right hand side elements.
`rng`	Double array of length `nrows` containing the row range elements. This may be `NULL` if there are no ranged constraints. The values in the `rng` array will only be read for `R` type rows. The entries for other type rows will be ignored. The lower bound on a range row is the right hand side value minus the range value. The sign of the range value is ignored - the absolute value is used in all cases.
`start`	Integer array of length `nrows` containing the offsets in the `colind` and `rowcoef` arrays of the start of the elements for each row.
`colind`	Integer array of length `ncoefs` containing the (contiguous) column indices for the elements in each row.
`rowcoef`	Double array of length `ncoefs` containing the (contiguous) element values.

Related controls

Integer

`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAROWS`	Number of extra rows to be allowed for.

Double

`MATRIXTOL`	Tolerance on matrix elements.

Example

Suppose the current problem is:

maximize:	$2x + y + 3z$		
subject to:	$x + 4y + 2z$	$\leqslant$	24
	$y + z$	$\leqslant$	5
	$3x + y$	$\leqslant$	20
	$x + y + 3z$	$\leqslant$	9

Then the following adds the row $8x + 9y + 10z \leqslant 25$ to the problem and names it `NewRow`:

```
rowtype[0] = 'L';
rhs[0] = 25.0;
start[] = {0};
colind[] = {0, 1, 2};
rowcoef[] = {8.0, 9.0, 10.0};
...
XPRSaddrows(prob,1,3,rowtype,rhs,NULL,start,colind, rowcoef);
XPRSaddnames(prob,XPRS_NAMES_ROW,"NewRow",4,4);
```

See also examples `els_managedcuts.c`, `els_usercuts.c`, `goalprog.c`, `loadlp.c`, `trimloss.c`, `tsp.c` .

Further information

1. Range rows are automatically converted to type `L`, with an upper bound in the slack. This must be taken into consideration, when retrieving row type, right-hand side values or range information for rows.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddcols`, `XPRSaddcuts`, `XPRSaddnames`, `XPRSdelrows`.

XPRSaddsets, XPRSaddsets64

Purpose

Allows sets to be added to the problem after passing it to the Optimizer using the input routines.

Topic areas

Problem Creation, MIP Entities

Synopsis

```
int XPRS_CC XPRSaddsets(XPRSprob prob, int nsets, int nelems, const char
      settype[], const int start[], const int colind[], const double
      refval[]);

int XPRS_CC XPRSaddsets64(XPRSprob prob, int nsets, XPRSint64 nelems, const
      char settype[], const XPRSint64 start[], const int colind[], const
      double refval[]);
```

Arguments

`prob`	The current problem.
`nsets`	Number of new sets.
`nelems`	Number of new nonzeros in the added sets.
`settype`	Character array of length nsets containing the set types: `1` indicates a SOS1; `2` indicates a SOS2;
`start`	Integer array of length `nsets` containing the offsets in the `colind` and `refval` arrays of the start of the elements for each set.
`colind`	Integer array of length `nelems` containing the (contiguous) column indices for the elements in each set.
`refval`	Double array of length `nelems` containing the (contiguous) reference values. These define the order for SOS2 constraints and may be used in branching for both types.

Further information

1. Reference values must be sufficiently distinct (see `SOSREFTOL`).
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSdelsets`.

XPRSaddsetnames

Purpose

This subroutine is deprecated and will be removed in a future release. Use `XPRSaddnames` instead.
When a model with MIP entities is loaded, any special ordered sets may not have names associated with them. If you wish names to appear in the ASCII solutions files, the names for a range of sets can be added with this function.

Topic areas

Problem Creation, MIP Entities

Synopsis

```
int XPRS_CC XPRSaddsetnames(XPRSprob prob, const char names[], int first,
        int last);
```

Arguments

`prob`	The current problem.
`names`	Character buffer containing the null-terminated string names.
`first`	Start of the range of sets.
`last`	End of the range of sets.

Example

Add set names (`set1` and `set2`) to a problem:

```
char snames[] = "set1\0set2"
...
XPRSaddsetnames(prob,snames,0,1);
```

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddnames`, `XPRSloadmip`, `XPRSloadmiqp`.

XPRSalter ALTER

Purpose

This subroutine is deprecated and will be removed in a future release. To change matrix coefficients, use `XPRSchgmcoef`. To change column bounds, use `XPRSchgbounds`. To change constraint right-hand side values, use `XPRSchgrhs`. To change constraint types, use `XPRSchgrowtype`. To change objective coefficients, use `XPRSchgobj`.
Alters or changes matrix elements, right hand sides and constraint senses in the current problem.

Topic area

File IO

Synopsis

```
int XPRS_CC XPRSalter(XPRSprob prob, const char *filename);
ALTER [filename]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters specifying the file to be read. If omitted, the default *problem_name* is used with a `.alt` extension.

Related controls

Integer

`EXTRAELEMS`	Number of extra matrix elements to be allowed for.

Double

`MATRIXTOL`	Tolerance on matrix elements.

Example 1 (Library)

Since the following call does not specify a filename, the file *problem_name*`.alt` is read in, from which commands are taken to alter the current matrix.

```
XPRSalter(prob,"");
```

Example 2 (Console)

The following example reads in the file `fred.alt`, from which instructions are taken to alter the current matrix:

```
ALTER fred
```

Further information

1. The file `filename.alt` is read. It is an ASCII file containing matrix revision statements in the format described in Section A.7. The `MODIFY` format of the MPS `REVISE` data is also supported.
2. It is not possible to alter a problem that is in a presolved state. Call `XPRSpostsolve` to bring the problem back to its original state.
3. If the problem was read from an .lp file, the name to use for the right-hand side is the one given by the attribute `RHSNAME` which by default is set to `RHS00001`.

Related topics

Section A.7.

XPRSbasiscondition BASISCONDITION

Purpose

This subroutine is deprecated and will be removed in a future release. Please use the `XPRSbasisstability` function instead.
Calculates the condition number of the current basis after solving the LP relaxation.

Topic areas

Linear Algebra, LP

Synopsis

```
int XPRS_CC XPRSbasiscondition(XPRSprob prob, double *p_cond, double
       *p_scaledcond);
BASISCONDITION
```

Arguments

`prob`	The current problem.
`p_cond`	The returned condition number of the current basis.
`p_scaledcond`	The returned condition number of the current basis for the scaled problem.

Example 1 (Library)

Get the condition number after optimizing a problem.

```
XPRSlpoptimize(prob," ");
XPRSbasiscondition(prob,&cond,&scaledcond);
printf("Condition no's are %g %g\n",cond,scaledcond);
```

Example 2 (Console)

Print the condition number after optimizing a problem.

```
READPROB problem.mps
LPOPTIMIZE
BASISCONDITION
```

Further information

1. The condition number of an invertible matrix is the norm of the matrix multiplied with the norm of its inverse. This number is an indication of how accurate the solution can be calculated and how sensitive it is to small changes in the data. The larger the condition number is, the less accurate the solution is likely to become.
2. When using the `BASISCONDITION` command in the Console Optimizer, the condition number is shown both for the scaled problem and in parenthesis for the original problem.

XPRSbasisstability BASISSTABILITY

Purpose

Calculates various measures for the stability of the current basis, including the basis condition number.

Topic areas

Linear Algebra, LP

Synopsis

```
int XPRS_CC XPRSbasisstability(XPRSprob prob, int type, int norm, int
        scaled, double *p_value);
BASISSTABILITY [-flags]
```

Arguments

prob	The current problem.	
type	0	Condition number of the basis.
	1	Stability measure for the solution relative to the current basis.
	2	Stability measure for the duals relative to the current basis.
	3	Stability measure for the right hand side relative to the current basis.
	4	Stability measure for the basic part of the objective relative to the current basis.
norm	0	Use the infinity norm.
	1	Use the 1 norm.
	2	Use the Euclidian norm for vectors, and the Frobenius norm for matrices.
scaled	If the stability values are to be calculated in the scaled, or the unscaled matrix.	
p_value	Pointer to a double, where the calculated value is to be returned.	
flags	x	Stability measure for the solution and right–hand side values relative to the current basis.
	d	Stability measure for the duals and the basic part of the objective relative to the current basis.
	c	Condition number of the basis (default).
	i	Use the infinity norm (default).
	o	Use the one norm.
	e	Use the Euclidian norm for vectors, and the Frobenius norm for matrices.
	u	Calculate values in the unscaled matrix.

Further information

1. The Console Optimizer command `BASISSTABILITY` uses `0` as the default value for type and norm, and calculates the values in the scaled matrix.

2. The condition number (`type` = `0`) of an invertible matrix is the norm of the matrix multiplied with the norm of its inverse. This number is an indication of how accurate the solution can be calculated and how sensitive it is to small changes in the data. The larger the condition number is, the less accurate the solution is likely to become.

3. The stability measures (`type` = `1...4`) are using the original matrix and the basis to recalculate the various vectors related to the solution and the duals. The returned stability measure is the norm of the difference of the recalculated vector to the original one.

XPRSbndsa

Purpose

Returns upper and lower sensitivity ranges for specified variables' lower and upper bounds. If the bounds are varied within these ranges the current basis remains optimal and feasible.

Topic areas

Sensitivity Analysis, LP

Synopsis

```
int XPRS_CC XPRSbndsa(XPRSprob prob, int ncols, const int colind[], double
       lblower[], double lbupper[], double ublower[], double ubupper[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of variables whose sensitivity is sought.
`colind`	Integer array of length `ncols` containing the indices of the columns whose bounds' ranges are required.
`lblower`	Double array of length `ncols` where the variable lower bound lower range values are to be returned.
`lbupper`	Double array of length `ncols` where the variable lower bound upper range values are to be returned.
`ublower`	Double array of length `ncols` where the variable upper bound lower range values are to be returned.
`ubupper`	Double array of length `ncols` where the variable upper bound upper range values are to be returned.

Further information

`XPRSbndsa` can only be called when an optimal solution to the current LP has been found. It cannot be used when the problem is MIP presolved.

Related topics

`XPRSrhssa`, `XPRSobjsa`.

XPRSbtran

Purpose

Post-multiplies a (row) vector provided by the user by the inverse of the current basis.

Topic areas

Linear Algebra, LP

Synopsis

```
int XPRS_CC XPRSbtran(XPRSprob prob, double vec[]);
```

Arguments

`prob`	The current problem.
`vec`	Double array of length `ROWS` containing the values by which the basis inverse is to be multiplied. The transformed values will also be returned in this array.

Related controls

Double

`ETATOL`	Tolerance on eta elements.

Example

Get the (unscaled) tableau row `z` of constraint number `irow`, assuming that all arrays have been dimensioned.

```
/* Minimum size of arrays:
 * y: nrow + ncol;
 * start: 2;
 * rowind, rowcoef: nrow.
 */

/* set up the unit vector y to pick out row irow */
for(i = 0; i < nrow; i++) y[i] = 0.0;
y[irow] = 1.0;

rc = XPRSbtran(prob,y);        /* y = e*B^{-1} */

/* Form z = y * A */
for(j = 0; j < ncol, j++) {
  rc = XPRSgetcols(prob, start, rowind, rowcoef,
                   nrow, &ncoefs, j, j);
  for(d = 0.0, ielt = 0, ielt < ncoefs; ielt++)
    d += y[rowind[ielt]] * rowcoef[ielt];
  y[nrow + j] = d;
}
```

See also example `tableau.c`.

Further information

If the matrix is in a presolved state, `XPRSbtran` works with the basis for the presolved problem.

Related topics

`XPRSftran`.

XPRScalcobjn

Purpose

Calculates the objective value of the given objective function in a multi-objective problem.

Topic areas

Solution, Multiobjective

Synopsis

```
int XPRS_CC XPRScalcobjn(XPRSprob prob, int objidx, const double
       solution[], double* p_objval);
```

Arguments

`prob`	The current problem.
`objidx`	Index of the objective function to calculate.
`solution`	Double array of length `COLS` that holds the solution, or `NULL` to use the current solution.
`p_objval`	Pointer to a double in which the calculated objective value is returned.

Example

The following obtains the optimal values of all objective functions and displays them to the console:

```
double objval;
int nobjective, i;
XPRSgetintattrib(prob, XPRS_OBJECTIVES, &nobjective);
XPRSlpoptimize(prob,"");
for (i = 0; i < nobjective; i++) {
  XPRScalcobjn(prob,i,NULL,&objval);
  printf("Objective %d: %f\n",i,objval);
}
```

Related topics

`XPRScalcobjective`.

XPRScalcobjective

Purpose

Calculates the objective value of a given solution.

Topic area

Solution

Synopsis

```
int XPRS_CC XPRScalcobjective(XPRSprob prob, const double solution[],
      double* p_objval);
```

Arguments

`prob`	The current problem.
`solution`	Double array of length COLS that holds the solution.
`p_objval`	Pointer to a double in which the calculated objective value is returned.

Further information

The calculations are always carried out in the original problem, even if the problem is currently presolved.

Related topics

`XPRScalcslacks`, `XPRScalcsolinfo`, `XPRScalcreducedcosts`.

XPRScalcreducedcosts

Purpose

Calculates the reduced cost values for a given (row) dual solution.

Topic areas

Linear Algebra, Solution

Synopsis

```
int XPRS_CC XPRScalcreducedcosts(XPRSprob prob, const double duals[], const
       double solution[], double djs[]);
```

Arguments

`prob`	The current problem.
`duals`	Double array of length ROWS that holds the dual solution to calculate the reduced costs for.
`solution`	Optional double array of length COLS that holds the primal solution. This is necessary for quadratic problems.
`djs`	Double array of length COLS in which the calculated reduced costs are returned.

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. The calculations are always carried out in the original problem, even if the problem is currently presolved.
3. If using the function during a solve (e.g. from a callback), use ORIGINALCOLS and ORIGINALROWS to retrieve the non-presolved dimensions of the problem.

Related topics

`XPRScalcslacks`, `XPRScalcsolinfo`, `XPRScalcobjective`.

XPRScalcslacks

Purpose

Calculates the row slack values for a given solution.

Topic area

Solution

Synopsis

```
int XPRS_CC XPRScalcslacks(XPRSprob prob, const double solution[], double
        slacks[]);
```

Arguments

`prob`	The current problem.
`solution`	Double array of length COLS that holds the solution to calculate the slacks for.
`slacks`	Double array of length ROWS in which the calculated row slacks are returned.

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only. For nonlinear, please use `XPRSnlpcalcslacks`.
2. The calculations are always carried out in the original problem, even if the problem is currently presolved.
3. If using the function during a solve (e.g. from a callback), use ORIGINALCOLS and ORIGINALROWS to retrieve the non-presolved dimensions of the problem.

Related topics

`XPRScalcreducedcosts`, `XPRScalcsolinfo`, `XPRScalcobjective`, `XPRSnlpcalcslacks`.

XPRScalcsolinfo

Purpose

Calculates the required property of a solution, like maximum infeasibility of a given primal and dual solution.

Topic area

Solution

Synopsis

```
int XPRS_CC XPRScalcsolinfo(XPRSprob prob, const double solution[], const
       double duals[], int property, double* p_value);
```

Arguments

`prob`	The current problem.
`solution`	Double array of length COLS that holds the solution. May be `NULL` when asking for dual infeasibility.
`duals`	Double array of length ROWS that holds the dual solution. May be `NULL` when asking for primal/MIP infeasibility.
`property`	Defined the property to be calculated. `XPRS_SOLINFO_ABSPRIMALINFEAS` the calculated maximum absolute primal infeasibility is returned. `XPRS_SOLINFO_RELPRIMALINFEAS` the calculated maximum relative primal infeasibility is returned. `XPRS_SOLINFO_ABSDUALINFEAS` the calculated maximum absolute dual infeasibility is returned. `XPRS_SOLINFO_RELDUALINFEAS` the calculated maximum relative dual infeasibility is returned. `XPRS_SOLINFO_MAXMIPFRACTIONAL` the calculated maximum absolute MIP fractionality or SOS infeasibility. `XPRS_SOLINFO_ABSMIPINFEAS` the calculated maximum absolute MIP infeasibility (including delayed rows, indicators, general and piecewise linear constraints) is returned. `XPRS_SOLINFO_RELMIPINFEAS` the calculated maximum relative MIP infeasibility (including delayed rows, indicators, general and piecewise linear constraints) is returned.
`p_value`	Pointer to a double where the calculated value is returned.

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. The calculations are always carried out in the original problem, even if the problem is currently presolved.

Related topics

`XPRScalcslacks`, `XPRScalcobjective`, `XPRScalcreducedcosts`.

CHECKCONVEXITY

Purpose

Checks if the loaded problem is convex. Applies to quadratic, mixed integer quadratic and quadratically constrained problems. Checking convexity takes some time, thus for problems that are known to be convex it might be reasonable to switch the checking off. Returns an error if the problem is not convex.

Topic areas

Linear Algebra, Quadratic

Synopsis

```
CHECKCONVEXITY
```

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. This console function checks the positive semi-definiteness of all quadratic matrices in the problem. Note, that when optimizing a problem, for quadratic programming and mixed integer quadratic problems, the checking of the objective function is performed after presolve, thus it is possible that an otherwise indefinite quadratic matrix will be found positive semi-definite (the indefinite part might have been fixed and dropped by presolve).

Related topics

`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`), `IFCHECKCONVEXITY`, `EIGENVALUETOL`.

XPRSchgbounds

Purpose

Used to change the bounds on columns in the matrix.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgbounds(XPRSprob prob, int nbounds, const int colind[],
      const char bndtype[], const double bndval[]);
```

Arguments

`prob`	The current problem.
`nbounds`	Number of bounds to change.
`colind`	Integer array of size `nbounds` containing the indices of the columns on which the bounds will change.
`bndtype`	Character array of length `nbounds` indicating the type of bound to change: `U` indicates change the upper bound; `L` indicates change the lower bound; `B` indicates change both bounds, i.e. fix the column.
`bndval`	Double array of length `nbounds` giving the new bound values.

Example

The following changes column `0` of the current problem to have an upper bound of `0.5`:

```
colind[0] = 0;
bndtype[0] = 'U';
bndval[0] = 0.5;
XPRSchgbounds(prob,1,colind,bndtype,bndval);
```

See also examples `addmipsol.c`, `fixbv.c`, `repair.c`, `roundint.c` .

Further information

1. A column index may appear twice in the `colind` array so it is possible to change both the upper and lower bounds on a variable in one go.
2. `XPRSchgbounds` can be used from within callbacks to modify the local node bounds of a branch-and-bound search. In this case, the solver expects column indices to match the presolved problem. This function can be called from within the following callbacks: `XPRSaddcbprenode`, `XPRSaddcbnodelpsolved`, `XPRSaddcboptnode`, `XPRSaddcbcutmgr`.
3. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` defined in the library header file can be used to represent plus and minus infinity respectively in the bound (`bndval`) array.
4. If the upper bound on a binary variable is changed to be greater than 1 or the lower bound is changed to be less than 0 then the variable will become an integer variable.

Related topics

`XPRSgetlb`, `XPRSgetub`,

XPRSchgcoef

Purpose

Used to change a single coefficient in the matrix. If the coefficient does not already exist, a new coefficient will be added to the matrix. If many coefficients are being added to a row of the matrix, it may be more efficient to delete the old row of the matrix and add a new row.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgcoef(XPRSprob prob, int row, int col, double coef);
```

Arguments

`prob`	The current problem.
`row`	Row index for the coefficient.
`col`	Column index for the coefficient.
`coef`	New value for the coefficient. If `coef` is zero, any existing coefficient will be deleted.

Related controls

Double

`MATRIXTOL`	Tolerance on matrix elements.

Example

In the following, the element in row `2`, column `1` of the matrix is changed to `0.33`:

```
XPRSchgcoef(prob,2,1,0.33);
```

Further information

`XPRSchgmcoef` is more efficient than multiple calls to `XPRSchgcoef` and should be used in its place in such circumstances.

Related topics

`XPRSaddcols`, `XPRSaddrows`, `XPRSchgmcoef`, `XPRSchgmqobj`, `XPRSchgobj`, `XPRSchgqobj`, `XPRSchgrhs`, `XPRSgetcols`, `XPRSgetrows`.

XPRSchgcoltype

Purpose

Used to change the type of a specified set of columns in the matrix.

Topic areas

Problem Modification, MIP Entities

Synopsis

```
int XPRS_CC XPRSchgcoltype(XPRSprob prob, int ncols, const int colind[],
        const char coltype[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of columns to change.
`colind`	Integer array of length `ncols` containing the indices of the columns.
`coltype`	Character array of length `ncols` giving the new column types: `C` indicates a continuous column; `B` indicates a binary column; `I` indicates an integer column. `S` indicates a semi-continuous column. The semi-continuous lower bound will be set to `1.0`. `R` indicates a semi-integer column. The semi-integer lower bound will be set to `1.0`. `P` indicates a partial integer column. The partial integer limit will be set to `1.0`.

Example

The following changes columns 3 and 5 of the matrix to be integer and binary respectively:

```
colind[0] = 3; colind[1] = 5;
coltype[0] = 'I'; coltype[1] = 'B';
XPRSchgcoltype(prob,2,colind,coltype);
```

See also examples `els_managedcuts.c`, `els_usercuts.c`, `goalprog.c`, `trimloss.c`, `tsp.c` .

Further information

1. The column types can only be changed before the tree search is started.
2. Calling `XPRSchgcoltype` to change any variable into a binary variable causes the bounds previously defined for the variable to be deleted and replaced by bounds of 0 and 1.
3. Calling `XPRSchgcoltype` to change a continuous variable into an integer variable cause its lower bound to be rounded up to the nearest integer value and its upper bound to be rounded down to the nearest integer value.
4. After changing the type of a column to semi-continuous, semi-integer or partial integer, call `XPRSchgglblimit` to set the limit value.
5. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.
6. For performance reasons, avoid turning MIP entities into continuous columns by calling `XPRSchgcoltype` many times in the presence of SOS constraints. If an application has this requirement, it is recommended to first collect all such columns beforehand and then change their type in a single call to this function.

Related topics

`XPRSaddcols`, `XPRSchgrowtype`, `XPRSdelcols`, `XPRSgetcoltype`, `XPRSchgglblimit`.

XPRSchgglblimit

Purpose

Used to change semi-continuous or semi-integer lower bounds, or upper limits on partial integers.

Topic areas

Problem Modification, MIP Entities

Synopsis

```
int XPRS_CC XPRSchgglblimit(XPRSprob prob, int ncols, const int colind[],
      const double limit[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of column limits to change.
`colind`	Integer array of size `ncols` containing the indices of the semi-continuous, semi-integer or partial integer columns that should have their limits changed.
`limit`	Double array of length `ncols` giving the new limit values.

Further information

1. The new limits are not allowed to be negative.
2. Partial integer limits can be at most `2^28`.
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSchgcoltype`, `XPRSgetmipentities`.

XPRSchgmcoef, XPRSchgmcoef64

Purpose

Used to change multiple coefficients in the matrix. If any coefficient does not already exist, it will be added to the matrix. If many coefficients are being added to a row of the matrix, it may be more efficient to delete the old row of the matrix and add a new one.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgmcoef(XPRSprob prob, int ncoefs, const int rowind[],
       const int colind[], const double rowcoef[]);

int XPRS_CC XPRSchgmcoef64(XPRSprob prob, XPRSint64 ncoefs, const int
       rowind[], const int colind[], const double rowcoef[]);
```

Arguments

`prob`	The current problem.
`ncoefs`	Number of new coefficients.
`rowind`	Integer array of length `ncoefs` containing the row indices of the coefficients to be changed.
`colind`	Integer array of length `ncoefs` containing the column indices of the coefficients to be changed.
`rowcoef`	Double array of length `ncoefs` containing the new coefficient values. If an element of `rowcoef` is zero, the coefficient will be deleted.

Related controls

Double

`MATRIXTOL` Tolerance on matrix elements.

Example

```
rowind[0] = 0; rowind[1] = 3;
colind[0] = 1; colind[1] = 5;
rowcoef[0] = 2.0; rowcoef[1] = 0.0;
XPRSchgmcoef(prob,2,rowind,colind,rowcoef);
```

This changes two elements to values `2.0` and `0.0`.

Further information

`XPRSchgmcoef` is more efficient than repeated calls to `XPRSchgcoef` and should be used in its place if many coefficients are to be changed.

Related topics

`XPRSchgcoef`, `XPRSchgmqobj`, `XPRSchgobj`, `XPRSchgqobj`, `XPRSchgrhs`, `XPRSgetcols`, `XPRSgetrhs`.

XPRSchgmqobj, XPRSchgmqobj64

Purpose

Used to change multiple quadratic coefficients in the objective function. If any of the coefficients does not exist already, new coefficients will be added to the objective function.

Topic areas

Problem Modification, Quadratic

Synopsis

```
int XPRS_CC XPRSchgmqobj(XPRSprob prob, int ncoefs, const int objqcol1[],
       const int objqcol2[], const double objqcoef[]);

int XPRS_CC XPRSchgmqobj64(XPRSprob prob, XPRSint64 ncoefs, const int
       objqcol1[], const int objqcol2[], const double objqcoef[]);
```

Arguments

`prob`	The current problem.
`ncoefs`	The number of coefficients to change.
`objqcol1`	Integer array of size `ncol` containing the column index of the first variable in each quadratic term.
`objqcol2`	Integer array of size `ncol` containing the column index of the second variable in each quadratic term.
`objqcoef`	New values for the coefficients. If an entry in `objqcoef` is `0`, the corresponding entry will be deleted. These are the coefficients of the quadratic Hessian matrix.

Example

The following code results in an objective function with terms: $[15x_1^2 + 7x_1x_2]/2$

```
objqcol1[0] = 0; objqcol2[0] = 0; objqcoef[0] = 15;
objqcol1[1] = 1; objqcol2[1] = 0; objqcoef[1] = 3.5;
XPRSchgmqobj(prob,2,objqcol1,objqcol2,objqcoef);
```

Further information

1. If `objqcol1[t]` is not equal to `objqcol2[t]`, then both the matrix elements `(objqcol1[t], objqcol2[t])` and `(objqcol2[t], objqcol1[t])` are changed to leave the Hessian symmetric.
2. The quadratic matrix coefficients are implicitly divided by two. But since the coefficients for $x_i x_j$ (`i` not equal to `j`) appear twice, only the coefficients for x_i^2 should be multiplied by two in the `objqcoef` array to account for the implicit division. See the example above.
3. `XPRSchgmqobj` is more efficient than repeated calls to `XPRSchgqobj` and should be used in its place when several coefficients are to be changed.

Related topics

`XPRSchgcoef`, `XPRSchgmcoef`, `XPRSchgobj`, `XPRSchgqobj`, `XPRSgetqobj`.

XPRSchgobjn

Purpose

Modifies one or more coefficients of an objective function in a multi-objective problem. If the objective already exists, any coefficients not present in the `colind` and `objcoef` arrays will unchanged. If the objective does not exist, it will be added to the problem.

Topic areas

Problem Modification, Multiobjective

Synopsis

```
int XPRS_CC XPRSchgobjn(XPRSprob prob, int objidx, int ncols, const int
       colind[], const double objcoef[]);
```

Arguments

`prob`	The current problem.
`objidx`	Index of the objective function to add or modify.
`ncols`	Number of objective function coefficient elements to change.
`colind`	Integer array of length `ncols` containing the indices of the columns whose objective coefficients will change. An index of `-1` indicates that the fixed part of the objective function on the right hand side should change.
`objcoef`	Double array of length `ncols` giving the new objective function coefficients.

Example

Changing three coefficients of the first objective function:

```
colind[0] = 0; colind[1] = 2; colind[2] = 5;
objcoef[0] = 25.0; objcoef[1] = 5.3; objcoef[2] = 0.0;
XPRSchgobjn(prob,0,3,colind,objcoef);
```

See also example `goalprog.c`.

Further information

1. When `objidx=0`, this function is equivalent to `XPRSchgobj`.
2. Any objectives with `idx < objidx` that do not already exist will be added to the problem with all zero coefficients.
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddobj`, `XPRSgetobjn`, `XPRSdelobj`, `XPRSchgobj`.

XPRSchgobj

Purpose

Used to change the objective function coefficients.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgobj(XPRSprob prob, int ncols, const int colind[], const
       double objcoef[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of objective function coefficient elements to change.
`colind`	Integer array of length `ncols` containing the indices of the columns whose objective coefficients will change. An index of `-1` indicates that the fixed part of the objective function on the right hand side should change.
`objcoef`	Double array of length `ncols` giving the new objective function coefficients.

Example

Changing three coefficients of the objective function with `XPRSchgobj`:

```
colind[0] = 0; colind[1] = 2; colind[2] = 5;
objcoef[0] = 25.0; objcoef[1] = 5.3; objcoef[2] = 0.0;
XPRSchgobj(prob,3,colind,objcoef);
```

See also examples `els_managedcuts.c`, `els_usercuts.c`, `globjpar.c`, `goalprog.c` .

Further information

The value of the fixed part of the objective function can be obtained using the `OBJRHS` problem attribute.

Related topics

`XPRSchgcoef`, `XPRSchgmcoef`, `XPRSchgmqobj`, `XPRSchgqobj`, `XPRSgetobj`.

XPRSchgobjsense CHGOBJSENSE

Purpose

Changes the problem's objective function sense to minimize or maximize.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgobjsense(XPRSprob prob, int objsense);
CHGOBJSENSE [ min | max ]
```

Arguments

`prob`	The current problem.
`objsense`	`XPRS_OBJ_MINIMIZE` to change into a minimization, or `XPRS_OBJ_MAXIMIZE` to change into maximization problem.

Related topics

`XPRSlpoptimize`, `XPRSmipoptimize`.

Example

See also examples `fixbv.c`, `globjpar.c`, `glrhspar.c`, `knapsack.c`, `loadlp.c`, `Polygon_initialvalue.c`, `Polygon_textformula.c`, `Polygon_tokens.c`, `Polygon_userfunc.c`, `Polygon_userfunc_map.c`, `Polygon_userfunc_mapdelta.c`, `Polygon_userfunc_multimap.c`, `Polygon_userfunc_multimapdelta.c`, `Polygon_userfunc_vecmap.c`, `Polygon_userfunc_vecmapdelta.c`, `savesol.c`, `tableau.c` .

XPRSchgqobj

Purpose

Used to change a single quadratic coefficient in the objective function corresponding to the variable pair `(objqcol1,objqcol2)` of the Hessian matrix.

Topic areas

Problem Modification, Quadratic

Synopsis

```
int XPRS_CC XPRSchgqobj(XPRSprob prob, int objqcol1, int objqcol2, double
      objqcoef);
```

Arguments

`prob`	The current problem.
`objqcol1`	Column index for the first variable in the quadratic term.
`objqcol2`	Column index for the second variable in the quadratic term.
`objqcoef`	New value for the coefficient in the quadratic Hessian matrix. If an entry in `objqcoef` is `0`, the corresponding entry will be deleted.

Example

The following code adds the terms $[15x_1^2 + 7x_1x_2]/2$ to the objective function:

```
XPRSchgqobj(prob, 0, 0, 15);
XPRSchgqobj(prob, 0, 1, 3.5);
```

Further information

1. If `objqcol1` is not equal to `objqcol2`, then both the matrix elements `(objqcol1, objqcol2)` and `(objqcol2, objqcol1)` are changed to leave the Hessian symmetric.
2. The quadratic matrix coefficients are implicitly divided by two. But since the coefficients for $x_i x_j$ (`i` not equal to `j`) appear twice, only the coefficients for x_i^2 should be multiplied by two in the `objqcoef` argument to account for the implicit division. See the example above.

Related topics

`XPRSchgcoef`, `XPRSchgmcoef`, `XPRSchgmqobj`, `XPRSchgobj`, `XPRSgetqobj`.

XPRSchgqrowcoeff

Purpose

Changes a single quadratic coefficient in a row.

Topic areas

Problem Modification, Quadratic

Synopsis

```
int XPRS_CC XPRSchgqrowcoeff(XPRSprob prob, int row, int rowqcol1, int
       rowqcol2, double rowqcoef);
```

Arguments

`prob`	The current problem.
`row`	Index of the row where the quadratic matrix is to be changed.
`rowqcol1`	First index of the coefficient to be changed.
`rowqcol2`	Second index of the coefficient to be changed.
`rowqcoef`	The new coefficient.

Further information

This function may be used to add new nonzero coefficients, or even to define the whole quadratic expression with it. Doing that, however, is significantly less efficient than adding the whole expression with `XPRSaddqmatrix`.

Related topics

`XPRSloadqcqp`, `XPRSgetqrowcoeff`, `XPRSaddqmatrix`, `XPRSchgqrowcoeff`, `XPRSgetqrowqmatrix`, `XPRSgetqrowqmatrixtriplets`, `XPRSgetqrows`, `XPRSchgqobj`, `XPRSchgmqobj`, `XPRSgetqobj`.

XPRSchgrhs

Purpose

Used to change right–hand side values of the problem.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgrhs(XPRSprob prob, int nrows, const int rowind[], const
       double rhs[]);
```

Arguments

`prob`	The current problem.
`nrows`	Number of right hand side values to change.
`rowind`	Integer array of length `nrows` containing the indices of the rows on which the right hand side values will change.
`rhs`	Double array of length `nrows` giving the right hand side values.

Example

Here we change the three right hand sides in rows `2`, `6`, and `8` to new values:

```
rowind[0] = 2; rowind[1] = 8; rowind[2] = 6;
rhs[0] = 5.0; rhs[1] = 3.8; rhs[2] = 5.7;
XPRSchgrhs(prob,3,rowind,rhs);
```

See also examples `glrhspar.c`, `repair.c` .

Related topics

`XPRSchgcoef`, `XPRSchgmcoef`, `XPRSchgrhsrange`, `XPRSgetrhs`, `XPRSgetrhsrange`.

XPRSchgrhsrange

Purpose

Used to change the range for a row of the problem matrix.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgrhsrange(XPRSprob prob, int nrows, const int rowind[],
       const double rng[]);
```

Arguments

`prob`	The current problem.
`nrows`	Number of range elements to change.
`rowind`	Integer array of length `nrows` containing the indices of the rows on which the range elements will change.
`rng`	Double array of length `nrows` giving the range values.

Example

Here, the constraint $x + y \leqslant 10$ (with row index 5) in the problem is changed to $8 \leqslant x + y \leqslant 10$:

```
rowind[0] = 5; rng[0] = 2.0;
XPRSchgrhsrange(prob,1,rowind,rng);
```

See also example `repair.c`.

Further information

If the range specified on the row is r, what happens depends on the row type and value of r. It is possible to convert non-range rows using this routine.

Value of r	Row type	Effect
$r \geqslant 0$	$= b, \leqslant b$	$b - r \leqslant \sum a_j x_j \leqslant b$
$r \geqslant 0$	$\geqslant b$	$b \leqslant \sum a_j x_j \leqslant b + r$
$r < 0$	$= b, \leqslant b$	$b \leqslant \sum a_j x_j \leqslant b - r$
$r < 0$	$\geqslant b$	$b + r \leqslant \sum a_j x_j \leqslant b$

Related topics

`XPRSchgcoef`, `XPRSchgmcoef`, `XPRSchgrhs`, `XPRSgetrhsrange`.

XPRSchgrowtype

Purpose

Used to change the type of a row in the matrix.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSchgrowtype(XPRSprob prob, int nrows, const int rowind[],
        const char rowtype[]);
```

Arguments

`prob`	The current problem.
`nrows`	Number of rows to change.
`rowind`	Integer array of length `nrows` containing the indices of the rows.
`rowtype`	Character array of length `nrows` giving the new row types: `L` indicates a $\leqslant$ row; `E` indicates an = row; `G` indicates a $\geqslant$ row; `R` indicates a range row; `N` indicates a free row.

Example

Here row `4` is changed to an equality row:

```
rowind[0] = 4; rowtype[0] = 'E';
XPRSchgrowtype(prob,1,rowind,rowtype);
```

See also example `repair.c`.

Further information

A row can be changed to a range type row by first changing the row to an `R` or `L` type row and then changing the range on the row using `XPRSchgrhsrange`.

Related topics

`XPRSaddrows`, `XPRSchgcoltype`, `XPRSchgrhs`, `XPRSchgrhsrange`, `XPRSdelrows`, `XPRSgetrowtype`.

XPRScopycallbacks

Purpose

Copies callback functions defined for one problem to another.

Topic area

Callback

Synopsis

```
int XPRS_CC XPRScopycallbacks(XPRSprob dest, XPRSprob src);
```

Arguments

`dest`	The problem to which the callbacks are copied.
`src`	The problem from which the callbacks are copied.

Example

The following sets up a message callback function `callback` for problem `prob1` and then copies this to the problem `prob2`.

```
XPRScreateprob(&prob1);
XPRSaddcbmessage(prob1,callback,NULL,0);
XPRScreateprob(&prob2);
XPRScopycallbacks(prob2,prob1);
```

Related topics

`XPRScopycontrols`, `XPRScopyprob`.

XPRSclearrowflags

Purpose

Clears extra information attached to a range of rows.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSclearrowflags(XPRSprob prob, const int flags[], int first,
       int last);
```

Arguments

`prob`	The current problem
`flags`	Int array of length `last-first+1` including type of extra information to remove (see below)
`first`	First row index to be checked
`last`	Last row index to be checked

Further information

The `flags` array consists of a bitvector for each row defining types of information to remove:

`XPRS_ROWFLAG_QUADRATIC`	Remove all quadratic coefficients.
`XPRS_ROWFLAG_DELAYED`	The row will not be a delayed row.
`XPRS_ROWFLAG_MODELCUT`	The row will not be a model cut.
`XPRS_ROWFLAG_INDICATOR`	Remove indicators associated to the row.
`XPRS_ROWFLAG_NONLINEAR`	Remove any nonlinear coefficients.

Example

The following makes sure that the second and third rows are not indicators and removes any quadratic matrix from the second row.

```
int flags[2];
int flags[2];
flags[0] = XPRS_ROWFLAG_QUADRATIC+XPRS_ROWFLAG_INDICATOR;
flags[1] = XPRS_ROWFLAG_INDICATOR;
XPRSclearrowflags(prob,1,2,flags);
```

Related topics

`XPRSgetrowflags`

XPRScopycontrols

Purpose

Copies controls defined for one problem to another.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRScopycontrols(XPRSprob dest, XPRSprob src);
```

Arguments

dest	The problem to which the controls are copied.
src	The problem from which the controls are copied.

Example

The following turns off Presolve for problem prob1 and then copies this and other control values to the problem prob2 :

```
XPRScreateprob(&prob1);
XPRSsetintcontrol(prob1,XPRS_PRESOLVE,0);
XPRScreateprob(&prob2);
XPRScopycontrols(prob2,prob1);
```

See also example addmipsol.c.

Related topics

XPRScopycallbacks, XPRScopyprob.

XPRScopyprob

Purpose

Copies information defined for one problem to another.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRScopyprob(XPRSprob dest, XPRSprob src, const char *name);
```

Arguments

`dest`	The new problem pointer to which information is copied.
`src`	The old problem pointer from which information is copied.
`name`	A string of up to 1024 characters (including `NULL` terminator) containing the name for the problem copy. This must be unique when file writing is to be expected, and particularly for MIP problems.

Example

The following copies the problem, its controls and it callbacks from `prob1` to `prob2`:

```
XPRSprob prob1, prob2;
...
XPRScreateprob(&prob2);
XPRScopyprob(prob2,prob1,"MyProb");
XPRScopycontrols(prob2,prob1);
XPRScopycallbacks(prob2,prob1);
```

See also examples `addmipsol.c`, `globjpar.c`, `glrhspar.c` .

Further information

`XPRScopyprob` copies only the problem and controls and does not copy the callbacks associated to a problem. These must be copied separately using `XPRScopycallbacks`.

Related topics

`XPRScopycallbacks`, `XPRScopycontrols`, `XPRScreateprob`.

XPRScreateprob

Purpose

Sets up a new problem within the Optimizer.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRScreateprob(XPRSprob *p_prob);
```

Argument

p_prob	Pointer to a variable holding the new problem.

Return value

0	Success
255	Optimizer is not intialized (call XPRSinit)
other	An error occurred

Example

The following creates a problem which will contain myprob:

```
XPRSprob prob;
XPRSinit(NULL);
XPRScreateprob(&prob);
XPRSreadprob(prob,"myprob","");
```

See also examples addmipsol.c, ComplexUserFunctions.c, els_managedcuts.c, els_usercuts.c, fixbv.c, globjpar.c, glrhspar.c, goalprog.c, knapsack.c, loadlp.c, mipsolenum.c, mipsolpool.c, mostviolated.c, Polygon_initialvalue.c, Polygon_textformula.c, Polygon_tokens.c, Polygon_userfunc.c, Polygon_userfunc_map.c, Polygon_userfunc_mapdelta.c, Polygon_userfunc_multimap.c, Polygon_userfunc_multimapdelta.c, Polygon_userfunc_vecmap.c, Polygon_userfunc_vecmapdelta.c, repair.c, roundint.c, savesol.c, tableau.c, trimloss.c, tsp.c .

Further information

1. XPRScreateprob must be called after XPRSinit and before using the other Optimizer routines.
2. Any number of problems may be created in this way, depending on your license details. All problems should be removed using XPRSdestroyprob once you have finished working with them.
3. If XPRScreateprob cannot complete successfully, a nonzero value is returned and *p_prob is set to NULL (as a consequence, it is not possible to retrieve further error information using e.g. XPRSgetlasterror).

Related topics

XPRSdestroyprob, XPRScopyprob, XPRSinit.

XPRScrossoverlpsol

Purpose

Provides a basic optimal solution for a given solution of an LP problem. This function behaves like the crossover after the barrier algorithm.

Topic areas

Barrier, LP

Synopsis

```
int XPRS_CC XPRScrossoverlpsol(XPRSprob prob, int *p_status);
```

Arguments

`prob`	The current problem.
`p_status`	Pointer to an `int` where the status will be returned. The status is one of: `0` The crossover is successful. `1` The crossover is not performed because the problem has no solution.

Related controls

Integer

`ALGAFTERCROSSOVER`	Specifies which algorithm to use for cleaning up the solution.
`PREPROTECTDUAL`	Whether or not to protect the given dual solution during presolve.

Example

This example loads a problem, loads a solution for the problem and then uses `XPRScrossoverlpsol` to find a basic optimal solution.

```
XPRSreadprob(prob, "problem", "");
XPRSloadlpsol(prob, x, NULL, dual, NULL, &status);
XPRScrossoverlpsol(prob, &status);
```

A solution can also be loaded from an ASCII solution file using `XPRSreadslxsol`.

Further information

1. The crossover contains two phases: a crossover phase for finding a basic solution and a clean-up phase for finding a basic optimal solution. Setting `ALGAFTERCROSSOVER` to `0` will allow the crossover to skip the clean-up phase.
2. The given solution is expected to be feasible or nearly feasible, otherwise the crossover may take a long time to find a basic feasible solution. More importantly, the given solution is expected to have a small duality gap. A small duality gap indicates that the given solution is close to the optimal solution. If the given solution is far away from the optimal solution, the clean-up phase may need many simplex iterations to move to a basic optimal solution.

Related topics

`XPRSloadlpsol`, `XPRSreadslxsol`, Section 4.2.1.

XPRSdelcols

Purpose

Delete columns from a matrix.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSdelcols(XPRSprob prob, int ncols, const int colind[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of columns to delete.
`colind`	Integer array of length `ncols` containing the columns to delete.

Example

In this example, column `3` is deleted from the matrix:

```
colind[0] = 3;
XPRSdelcols(prob,1,colind);
```

Further information

1. After columns have been deleted from a problem, the numbers of the remaining columns are moved down so that the columns are always numbered from `0` to `COLS-1` where `COLS` is the problem attribute containing the number of non-deleted columns in the matrix.
2. If the problem has already been optimized, or an advanced basis has been loaded, and you delete a basis column the current basis will no longer be valid - the basis is "lost".
   If you go on to re-optimize the problem, a warning message is displayed (140) and the Optimizer automatically generates a corrected basis.
   You can avoid losing the basis by only deleting non-basic columns (see `XPRSgetbasis`), taking a basic column out of the basis first if necessary (see `XPRSgetpivots` and `XPRSpivot`).
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.
4. If a deleted column is referenced by a formula then it will be replaced by the constant 0 in that formula. If the formula is a PWL formula then instead an error is raised because the input column for a PWL must be a variable and cannot be a constant zero.
5. If a deleted column is referenced by a set then it will be removed from the set. Empty sets will not be deleted by this function.
6. If a column is referenced by a piecewise linear constraint then an attempt to delete it will raise an error.
7. If a column is referenced by a general constraint then an attempt to delete it will raise an error.

Related topics

`XPRSaddcols`, `XPRSdelrows`.

XPRSdelcpcuts

Purpose

During the branch and bound search, cuts are stored in the cut pool to be applied at descendant nodes. These cuts may be removed from a given node using `XPRSdelcuts`, but if this is to be applied in a large number of cases, it may be preferable to remove the cut completely from the cut pool. This is achieved using `XPRSdelcpcuts`.

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSdelcpcuts(XPRSprob prob, int cuttype, int interp, int
       ncuts, const XPRScut cutind[]);
```

Arguments

`prob`	The current problem.
`cuttype`	User defined cut type to match against.
`interp`	Way in which the cut `cuttype` is interpreted: `-1` match all cut types; `1` treat cut types as numbers; `2` treat cut types as bit-vectors (compare Section 9.2) - delete if any bit matches any bit set in `cuttype`; `3` treat cut types as bit-vectors (compare Section 9.2) - delete if all bits match those set in `cuttype`.
`ncuts`	The number of cuts to delete. A value of `-1` indicates delete all cuts.
`cutind`	Array containing pointers to the cuts which are to be deleted. This array may be `NULL` if `ncuts` is `-1`, otherwise it has length `ncuts`.

Related topics

`XPRSaddcuts`, `XPRSdelcuts`, `XPRSloadcuts`, Section 5.9.

XPRSdelcuts

Purpose

Deletes cuts from the matrix at the current node. Cuts from the parent node which have been automatically restored may be deleted as well as cuts added to the current node using `XPRSaddcuts` or `XPRSloadcuts`. The cuts to be deleted can be specified in a number of ways. If a cut is ruled out by any one of the criteria it will not be deleted.

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSdelcuts(XPRSprob prob, int basis, int cuttype, int interp,
       double delta, int ncuts, const XPRScut cutind[]);
```

Arguments

`prob`	The current problem.
`basis`	Ensures the basis will be valid if set to `1`. If set to `0`, cuts with non-basic slacks may be deleted.
`cuttype`	User defined type of the cut to be deleted.
`interp`	Way in which the cut `cuttype` is interpreted: `-1` match all cut types; `1` treat cut types as numbers; `2` treat cut types as bit-vectors (compare Section 9.2) - delete if any bit matches any bit set in `cuttype`; `3` treat cut types as bit-vectors (compare Section 9.2) - delete if all bits match those set in `cuttype`.
`delta`	Only delete cuts with an absolute slack value greater than `delta`. To delete all the cuts, this argument should be set to `XPRS_MINUSINFINITY`.
`ncuts`	Number of cuts to drop if a list of cuts is provided. A value of `-1` indicates all cuts.
`cutind`	Array containing pointers to the cuts which are to be deleted. This array may be `NULL` if `ncuts` is set to `-1` otherwise it has length `ncuts`.

Further information

1. It is usually best to drop only those cuts with basic slacks, otherwise the basis will no longer be valid and it may take many iterations to recover an optimal basis. If the `basis` parameter is set to `1`, this will ensure that cuts with non-basic slacks will not be deleted even if the other parameters specify that these cuts should be deleted. It is highly recommended that the `basis` parameter is always set to `1`.

2. The cuts to be deleted can also be specified by the size of the slack variable for the cut. Only those cuts with a slack value greater than the `delta` parameter will be deleted.

3. A list of indices of the cuts to be deleted can also be provided. The list of active cuts at a node can be obtained with the `XPRSgetcutlist` command.

4. This function should be called only from within callback functions set by either `XPRSaddcboptnode`, `XPRSaddcbprenode`, `XPRSaddcbpreintsol` or `XPRSaddcbnodelpsolved`.

Related topics

`XPRSaddcuts`, `XPRSdelcpcuts`, `XPRSgetcutlist`, `XPRSloadcuts`, Section 5.9.

XPRSdelgencons

Purpose

Delete general constraints from a problem.

Topic areas

Problem Creation, Piecewise Linear and General Constraints

Synopsis

```
int XPRS_CC XPRSdelgencons(XPRSprob prob, int ncons, const int conind[]);
```

Arguments

`prob`	The current problem.
`ncons`	Number of general constraints to delete.
`conind`	An integer array of length `ncons` containing the general constraints to delete.

Example

In this example, general constraints 0 and 2 are deleted from the problem:

```
conind[0] = 0; conind[1] = 2;
XPRSdelgencons(prob,2,conind);
```

Further information

1. After general constraints have been deleted from a problem, the indices of the remaining constraints are reduced down so that the general constraints are always numbered from 0 to `GENCONS-1` where `GENCONS` is the problem attribute containing the number of non-deleted general constraints in the problem.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddgencons`, `XPRSgetgencons`.

XPRSdelindicators

Purpose

Delete indicator constraints. This turns the specified rows into normal rows (not controlled by indicator variables).

Topic areas

Problem Creation, MIP Entities

Synopsis

```
int XPRS_CC XPRSdelindicators(XPRSprob prob, int first, int last);
```

Arguments

`prob`	The current problem.
`first`	First row in the range.
`last`	Last row in the range (inclusive).

Example

In this example, if any of the first two rows of the matrix is an indicator constraint, they are turned into normal rows:

```
XPRSdelindicators(prob,0,1);
```

Further information

1. This function has no effect on rows that are not indicator constraints.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSgetindicators`, `XPRSsetindicators`.

XPRSdelobj

Purpose

Removes an objective function from a multi-objective problem. Any objectives with `index > objidx` will be shifted down. Deleting the last objective function in the problem causes all the objective coefficients to be zeroed, but `OBJECTIVES` remains set to `1`.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSdelobj(XPRSprob prob, int objidx);
```

Arguments

`prob`	The current problem.
`objidx`	Index of the objective to remove.

Example

Removing the second objective function from a problem:

```
XPRSdelobj(prob, 1);
```

Related topics

`XPRSchgobjn`, `XPRSaddobj`, `XPRSgetobjn`.

XPRSdelpwlcons

Purpose

Delete piecewise linear constraints from a problem.

Topic areas

Problem Creation, Piecewise Linear and General Constraints

Synopsis

```
int XPRS_CC XPRSdelpwlcons(XPRSprob prob, int npwls, const int pwlind[]);
```

Arguments

`prob`	The current problem.
`npwls`	Number of piecewise linear constraints to delete.
`pwlind`	An integer array of length `npwls` containing the piecewise linear constraints to delete.

Example

In this example, piecewise linear constraints `0` and `2` are deleted from the problem:

```
pwlind[0] = 0; pwlind[1] = 2;
XPRSdelpwlcons(prob,2,pwlind);
```

Further information

1. After piecewise linear constraints have been deleted from a problem, the indices of the remaining constraints are reduced so that the piecewise linear constraints are always numbered from `0` to `PWLCONS-1` where `PWLCONS` is the problem attribute containing the number of non-deleted piecewise linear constraints in the problem.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddpwlcons`, `XPRSgetpwlcons`.

XPRSdelqmatrix

Purpose

Deletes the quadratic part of a row or of the objective function.

Topic areas

Problem Creation, Quadratic

Synopsis

```
int XPRS_CC XPRSdelqmatrix(XPRSprob prob, int row);
```

Arguments

`prob`	The current problem.
`row`	Index of row from which the quadratic part is to be deleted.

Further information

1. If a row index of `-1` is used, the function deletes the quadratic coefficients from the objective function.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddrows`, `XPRSdelcols`, `XPRSdelrows`.

XPRSdelrows

Purpose

Delete rows from a matrix.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSdelrows(XPRSprob prob, int nrows, const int rowind[]);
```

Arguments

`prob`	The current problem.
`nrows`	Number of rows to delete.
`rowind`	An integer array of length `nrows` containing the rows to delete.

Example

In this example, rows `0` and `10` are deleted from the matrix:

```
rowind[0] = 0; rowind[1] = 10;
XPRSdelrows(prob,2,rowind);
```

Further information

1. After rows have been deleted from a problem, the numbers of the remaining rows are moved down so that the rows are always numbered from `0` to `ROWS-1` where `ROWS` is the problem attribute containing the number of non-deleted rows in the matrix.
2. If the problem has already been optimized, or an advanced basis has been loaded, and you delete a row for which the slack column is non-basic, the current basis will no longer be valid - the basis is "lost".

   If you go on to re-optimize the problem, a warning message is displayed (140) and the Optimizer automatically generates a corrected basis.
   You can avoid losing the basis by only deleting basic rows (see `XPRSgetbasis`), bringing a non-basic row into the basis first if necessary (see `XPRSgetpivots` and `XPRSpivot`).
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddrows`, `XPRSdelcols`, `XPRSgetbasis`, `XPRSgetpivots`, `XPRSpivot`.

XPRSdelsets

Purpose

Delete sets from a problem.

Topic areas

Problem Creation, MIP Entities

Synopsis

```
int XPRS_CC XPRSdelsets(XPRSprob prob, int nsets, const int setind[]);
```

Arguments

`prob`	The current problem.
`nsets`	Number of sets to delete.
`setind`	An integer array of length `nsets` containing the sets to delete.

Example

In this example, sets `0` and `2` are deleted from the problem:

```
setind[0] = 0; setind[1] = 2;
XPRSdelsets(prob,2,setind);
```

Further information

1. After sets have been deleted from a problem, the numbers of the remaining sets are moved down so that the sets are always numbered from `0` to `SETS-1` where `SETS` is the problem attribute containing the number of non-deleted sets in the problem.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddsets`.

XPRSdestroyprob

Purpose

Removes a given problem and frees any memory associated with it following manipulation and optimization.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSdestroyprob(XPRSprob prob);
```

Argument

prob	The problem to be destroyed.

Example

The following creates, loads and solves a problem called myprob, before subsequently freeing any resources allocated to it:

```
XPRScreateprob(&prob);
XPRSreadprob(prob,"myprob","");
XPRSlpoptimize(prob,"");
XPRSdestroyprob(prob);
```

See also examples addmipsol.c, ComplexUserFunctions.c, els_managedcuts.c, els_usercuts.c, fixbv.c, globjpar.c, glrhspar.c, goalprog.c, knapsack.c, loadlp.c, mipsolenum.c, mipsolpool.c, Polygon_initialvalue.c, Polygon_textformula.c, Polygon_tokens.c, Polygon_userfunc.c, Polygon_userfunc_map.c, Polygon_userfunc_mapdelta.c, Polygon_userfunc_multimap.c, Polygon_userfunc_multimapdelta.c, Polygon_userfunc_vecmap.c, Polygon_userfunc_vecmapdelta.c, repair.c, roundint.c, savesol.c, tableau.c, trimloss.c, tsp.c .

Further information

After work is finished, all problems must be destroyed. If a NULL problem pointer is passed to XPRSdestroyprob, no error will result.

Related topics

XPRScreateprob, XPRSfree, XPRSinit.

XPRSdumpcontrols DUMPCONTROLS

Purpose

Displays the list of controls and their current value for those controls that have been set to a non default value.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSdumpcontrols(XPRSprob prob);
DUMPCONTROLS
```

Argument

`prob`	The problem for which controls are dumped.

Related topics

`SETDEFAULTS`, `SETDEFAULTCONTROL`

EXIT

Purpose

Terminates the Console Optimizer, returning a zero exit code to the operating system. Alias of QUIT.

Topic area

Misc

Synopsis

`EXIT`

Example

The command is called simply as:

```
EXIT
```

Further information

1. Fatal error conditions return nonzero exit values which may be of use to the host operating system. These are described in Chapter 11.
2. If you wish to return an exit code reflecting the final solution status, then use the `STOP` command instead.

Related topics

`STOP`, `QUIT`, `XPRSsave` (`SAVE`).

XPRSestimaterowdualranges

Purpose

Performs a dual side range sensitivity analysis, i.e. calculates estimates for the possible ranges for dual values.

Topic areas

Sensitivity Analysis, LP

Synopsis

```
int XPRS_CC XPRSestimaterowdualranges(XPRSprob prob, const int nrows, const
      int rowind[], const int iterlim, double mindual[], double maxdual[]);
```

Arguments

`prob`	The current problem.
`nrows`	The number of rows to analyze.
`rowind`	Row indices to analyze.
`iterlim`	Effort limit expressed as simplex iterations per row.
`mindual`	Estimated lower bounds on the possible dual ranges.
`maxdual`	Estimated upper bounds on the possible dual ranges.

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. This function may provide better results for individual row dual ranges when called for a larger number of rows, as it collect infomation for all rows while working on individual rows.
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSlpoptimize`, `XPRSstrongbranch`

XPRSfeaturequery

Purpose

Checks if the provided feature is available in the current license used by the optimizer.

Topic area

Licensing

Synopsis

```
int XPRS_CC XPRSfeaturequery(const char *feature, int *p_status);
```

Arguments

`feature`	The feature string to be checked in the license.
`p_status`	Return status of the check, a value of 1 indicates the feature is available.

XPRSfixglobals FIXGLOBALS

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSfixmipentities` and `FIXMIPENTITIES` instead.

Topic areas

Problem Modification, MIP Entities

Synopsis

```
int XPRS_CC XPRSfixglobals(XPRSprob prob, int options);
FIXGLOBALS [-flags]
```

XPRSfixmipentities FIXMIPENTITIES

Purpose

Fixes all the MIP entities to the values of the last found MIP solution. This is useful for finding the reduced costs for the continuous variables after the integer variables have been fixed to their optimal values.

Topic areas

Problem Modification, MIP Entities

Synopsis

```
int XPRS_CC XPRSfixmipentities(XPRSprob prob, int options);
FIXMIPENTITIES [-flags]
```

Arguments

`prob` The current problem.

`options` Options how to fix the MIP entities.

Bit	Meaning
0	If all MIP entities should be rounded to the nearest discrete value in the solution before being fixed.
1	If piecewise linear and general constraints should be kept in the problem with only the non-convex decisions (i.e. which part of a non-convex piecewise linear function or which variable attains a maximum) fixed. Otherwise all variables appearing in piecewise linear or general constraints will be fixed.

`flags` Flags to pass to FIXMIPENTITIES:

- `r` round all MIP entities to the nearest feasible value in the solution before being fixed;
- `t` keep piecewise linear and general constraints and only fix their non-convex decisions.

Example 1 (Library)

This example performs a tree search on problem `myprob` and then uses `XPRSfixmipentities` before solving the remaining linear problem:

```
XPRSreadprob(prob,"myprob","");
XPRSmipoptimize(prob," ");
XPRSfixmipentities(prob, 1);
XPRSlpoptimize(prob," ");
XPRSwriteprtsol(prob);
```

Example 2 (Console)

A similar set of commands at the console would be as follows:

```
READPROB
MIPOPTIMIZE
FIXMIPENTITIES -r
LPOPTIMIZE
PRINTSOL
```

Further information

1. Because of tolerances, it is possible for e.g. a binary variable to be slightly fractional in the MIP solution, where it might have the value `0.999999` instead of being at exactly `1.0`. With `ifround = 0`, such a binary will be fixed at `0.999999`, but with `ifround = 1`, it will be fixed at `1.0`.
2. This command is useful for inspecting the reduced costs of the continuous variables in a matrix after the MIP entities have been fixed. Sensitivity analysis can also be performed on the continuous variables in a MIP problem using `XPRSrhssa`, `XPRSobjsa` or `XPRSbndsa` after calling `XPRSfixmipentities` (`FIXMIPENTITIES`).
3. For nonlinear problems, one can set the initial point to the solution returned by the MIP search (via `XPRSnlpsetcurrentiv`) and then call `XPRSfixmipentities`and reoptimize the problem using a local solver (`XPRSnlpoptimize`). to obtain an approximation of the dual values

Related topics

`XPRSmipoptimize` (`MIPOPTIMIZE`).

XPRSfree

Purpose

Frees any allocated memory and closes all open files.

Topic area

Licensing

Synopsis

```
int XPRS_CC XPRSfree(void);
```

Example

The following frees resources allocated to the problem `prob` and then tidies up before exiting:

```
XPRSdestroyprob(prob);
XPRSfree();
return 0;
```

See also examples `addmipsol.c`, `ComplexUserFunctions.c`, `els_managedcuts.c`, `els_usercuts.c`, `fixbv.c`, `globjpar.c`, `glrhspar.c`, `goalprog.c`, `knapsack.c`, `loadlp.c`, `mipsolenum.c`, `mipsolpool.c`, `mostviolated.c`, `Polygon_initialvalue.c`, `Polygon_textformula.c`, `Polygon_tokens.c`, `Polygon_userfunc.c`, `Polygon_userfunc_map.c`, `Polygon_userfunc_mapdelta.c`, `Polygon_userfunc_multimap.c`, `Polygon_userfunc_multimapdelta.c`, `Polygon_userfunc_vecmap.c`, `Polygon_userfunc_vecmapdelta.c`, `repair.c`, `roundint.c`, `savesol.c`, `tableau.c`, `trimloss.c`, `tsp.c` .

Further information

After a call to `XPRSfree` no library functions may be used without first calling `XPRSinit` again.

Related topics

`XPRSdestroyprob`, `XPRSinit`.

XPRSftran

Purpose

Pre-multiplies a (column) vector provided by the user by the inverse of the current matrix.

Topic areas

Linear Algebra, LP

Synopsis

```
int XPRS_CC XPRSftran(XPRSprob prob, double vec[]);
```

Arguments

`prob`	The current problem.
`vec`	Double array of length `ROWS` containing the values which are to be multiplied by the basis inverse. The transformed values appear in the array.

Related controls

Double

`ETATOL`	Tolerance on eta elements.

Example

To get the (unscaled) tableau column of structural variable number `jcol`, assuming that all arrays have been dimensioned, do the following:

```
/* Min size of arrays: start: 2; rowind, rowcoef & y: nrow. */
/* Get column as loaded originally, in sparse format */
rc = XPRSgetcols(prob, start, rowind, rowcoef, nrow, &ncoefs,
                 jcol, jcol);

/* Unpack into the zeroed array */
for(i = 0; i < nrow; i++)
y[i] = 0.0;
for(ielt = 0; ielt < ncoefs; ielt++)
y[rowind[ielt]] = rowcoef[ielt];

rc = XPRSftran(prob,y);
```

Get the (unscaled) tableau column of the slack variable for row number `irow`, assuming that all arrays have been dimensioned.

```
/* Min size of arrays: y: nrow */
/* Set up the original slack column in full format */
for(i = 0; i < nrow; i++)
y[i] = 0.0;
y[irow] = 1.0;

rc = XPRSftran(prob,y);
```

Further information

If the matrix is in a presolved state, the function will work with the basis for the presolved problem.

Related topics

`XPRSbtran`.

XPRSgetattribinfo

Purpose

Accesses the id number and the type information of an attribute given its name. An attribute name may be for example XPRS_ROWS. Names are case-insensitive and may or may not have the XPRS_ prefix. The id number is the constant used to identify the attribute for calls to functions such as XPRSgetintattrib. The type information returned will be one of the below integer constants defined in the xprs.h header file.

The function will return an id number of 0 and a type value of XPRS_TYPE_NOTDEFINED if the name is not recognized as an attribute name. Note that this will occur if the name is a control name and not an attribute name.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetattribinfo(XPRSprob prob, const char* name, int* p_id,
       int* p_type);
```

Arguments

prob	The current problem.
name	The name of the attribute to be queried. Names are case-insensitive and may or may not have the XPRS_ prefix. A full list of all attributes may be found in Chapter 9, or from the list in the xprs.h header file.
p_id	Pointer to an integer where the id number will be returned.
p_type	Pointer to an integer where the type id will be returned. The value will be one of the following constants from xprs.h: XPRS_TYPE_NOTDEFINED — The name was not recognized. XPRS_TYPE_INT — 32-bit integer. XPRS_TYPE_INT64 — 64-bit integer. XPRS_TYPE_DOUBLE — Double precision floating point. XPRS_TYPE_STRING — String.

Example

The following code example obtains the id number and the type id of the control or attribute with name given by name. Note that the name happens to be a control name in this example:

```
const char *name = "presolve";
int id, type;
...
if(XPRSgetattribinfo(prob, name, &id,
                     &type) || id==0) {
  if(XPRSgetcontrolinfo(prob, name, &id,
                        &type) || id==0) {
    printf("Unrecognized name: %s\n", name);
  }
}
```

Related topics

XPRSgetcontrolinfo.

XPRSgetbanner

Purpose

Returns the banner and copyright message.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSgetbanner(char *banner);
```

Argument

banner	A buffer of at least XPRS_MAXBANNERLENGTH characters in which the null terminated banner string will be returned.

Example

The following calls XPRSgetbanner to return banner information at the start of the program:

```
char banner[XPRS_MAXBANNERLENGTH];
...
if(XPRSinit(NULL))
{
  /* The error message when XPRSinit fails is written to the banner. */
  XPRSgetbanner(banner);
  printf("%s\n", banner);
  return 1;
}
XPRSgetbanner(banner);
printf("%s\n", banner);
```

See also example ComplexUserFunctions.c.

Further information

This function can most usefully be employed to return extra information if a problem occurs with XPRSinit.

Related topics

XPRSinit.

XPRSgetbasis

Purpose

Returns the current basis into the user's data arrays.

Topic areas

Linear Algebra, LP

Synopsis

```
int XPRS_CC XPRSgetbasis(XPRSprob prob, int rowstat[], int colstat[]);
```

Arguments

`prob` The current problem.

`rowstat` Integer array of length `ORIGINALROWS` to the basis status of the slack, surplus or artificial variable associated with each row. The status will be one of:
- `XPRS_NONBASIC_LOWER (0)` slack, surplus or artificial is non-basic at lower bound;
- `XPRS_BASIC (1)` slack, surplus or artificial is basic;
- `XPRS_NONBASIC_UPPER (2)` slack or surplus is non-basic at upper bound.
- `XPRS_SUPERBASIC (3)` slack or surplus is super-basic.

May be `NULL` if not required.

`colstat` Integer array of length `ORIGINALCOLS` to hold the basis status of the columns in the constraint matrix. The status will be one of:
- `XPRS_NONBASIC_LOWER (0)` variable is non-basic at lower bound, or superbasic at zero if the variable has no lower bound;
- `XPRS_BASIC (1)` variable is basic;
- `XPRS_NONBASIC_UPPER (2)` variable is non-basic at upper bound;
- `XPRS_SUPERBASIC (3)` variable is super-basic.

May be `NULL` if not required.

Example

The following example minimizes a problem before saving the basis for later:

```
int rows, cols, *rowstat, *colstat;
...
XPRSgetintattrib(prob,XPRS_ORIGINALROWS,&rows);
XPRSgetintattrib(prob,XPRS_ORIGINALCOLS,&cols);
rowstat = (int *) malloc(sizeof(int)*rows);
colstat = (int *) malloc(sizeof(int)*cols);
XPRSlpoptimize(prob,"");
XPRSgetbasis(prob,rowstat,colstat);
```

See also examples `globjpar.c`, `glrhspar.c` .

Further information

The basis information is always for the original (unpresolved) model.

Related topics

`XPRSgetpresolvebasis`, `XPRSloadbasis`, `XPRSloadpresolvebasis`.

XPRSgetbasisval

Purpose

Returns the current basis status for a specific column or row.

Topic areas

Linear Algebra, LP

Synopsis

```
int XPRS_CC XPRSgetbasisval(XPRSprob prob, int row, int col, int
       *p_rowstat, int *p_colstat);
```

Arguments

`prob`	The current problem.
`row`	Row index to get the row basis status for.
`col`	Column index to get the column basis status for.
`p_rowstat`	Integer pointer where the value of the row basis status will be returned. May be `NULL` if not required.
`p_colstat`	Integer pointer where the value of the column basis status will be returned. May be `NULL` if not required.

Related topics

`XPRSgetbasis`, `XPRSgetpresolvebasis`, `XPRSloadbasis`, `XPRSloadpresolvebasis`.

XPRSgetcallbackduals

Purpose

Returns the dual values

from the solution associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbackduals(XPRSprob prob, int *p_available, double
        duals[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a dual solution is available. May be `NULL` if not required.
`duals`	Double array of length `last-first+1` where the values of the dual variables will be returned. May be `NULL` if not required.
`first`	First row whose dual value to return.
`last`	Last row whose dual value to return.

Related topics

`XPRSgetcallbacksolution`, `XPRSgetcallbackslacks`, `XPRSgetcallbackredcosts`, `XPRSgetcallbackpresolveduals`, `XPRSgetduals`.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Row indices should be in the range `0` to `INPUTROWS-1`.
2. See `XPRSgetcallbacksolution` for a list of callbacks where this function can be called. Note that a dual solution is not available in all contexts where a primal solution is available.

XPRSgetcallbackredcosts

Purpose

Returns the reduced costs

from the solution associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbackredcosts(XPRSprob prob, int *p_available, double
        djs[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a dual solution is available. May be `NULL` if not required.
`djs`	Double array of length `last-first+1` where the reduced costs of the variables will be returned. May be `NULL` if not required.
`first`	First column whose reduced cost to return.
`last`	Last column whose reduced cost to return.

Related topics

`XPRSgetcallbacksolution`, `XPRSgetcallbackslacks`, `XPRSgetcallbackduals`, `XPRSgetcallbackpresolveredcosts`, `XPRSgetredcosts`.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Column indices should be in the range `0` to `INPUTCOLS-1`.
2. See `XPRSgetcallbacksolution` for a list of callbacks where this function can be called. Note that reduced costs are not available in all contexts where a primal solution is available.

XPRSgetcallbackslacks

Purpose

Returns the slack values

from the solution associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbackslacks(XPRSprob prob, int *p_available, double
        slacks[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a solution is available. May be `NULL` if not required.
`slacks`	Double array of length `last-first+1` where the values of the slack variables will be returned. May be `NULL` if not required.
`first`	First row whose slack value to return.
`last`	Last row whose slack value to return.

Related topics

`XPRSgetcallbacksolution`, `XPRSgetcallbackduals`, `XPRSgetcallbackredcosts`, `XPRSgetcallbackpresolveslacks`, `XPRSgetslacks`.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Row indices should be in the range `0` to `INPUTROWS-1`.
2. See `XPRSgetcallbacksolution` for a list of callbacks where this function can be called.

XPRSgetcallbacksolution

Purpose

Returns the primal values

from the solution associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbacksolution(XPRSprob prob, int *p_available, double
       x[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a solution is available. May be `NULL` if not required.
`x`	Double array of length `last-first+1` where the values of the primal variables will be returned. May be `NULL` if not required.
`first`	First column in the solution to return.
`last`	Last column in the solution to return.

Related topics

`XPRSgetcallbackslacks`, `XPRSgetcallbackduals`, `XPRSgetcallbackredcosts`, `XPRSgetcallbackpresolvesolution`, `XPRSgetsolution`.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Column indices should be in the range `0` to `INPUTCOLS-1`.
2. This function may only be called within the following callbacks:

Callback	**Solution type returned**
bariteration	Barrier iteration solution
chgbranchobject	Node relaxation solution
intsol	New incumbent solution
nodelpsolved	Node relaxation solution
optnode	Node relaxation solution
preintsol	Candidate incumbent solution
slpcascadeend	SLP iteration solution
slpcascadestart	SLP iteration solution
slpintsol	New incumbent solution
slpiterend	SLP iteration solution
slpiterstart	SLP iteration solution
slpitervar	SLP iteration solution
slppreupdatelinearization	SLP iteration solution

Example

See also examples `addmipsol.c`, `els_managedcuts.c`, `els_usercuts.c`, `knapsack.c`, `savesol.c`, `tsp.c` .

XPRSgetcallbackpresolveduals

Purpose

Returns the dual values from the solution to the presolved problem associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbackpresolveduals(XPRSprob prob, int *p_available,
        double duals[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a dual solution is available. May be `NULL` if not required.
`duals`	Double array of length `last-first+1` where the values of the dual variables will be returned. May be `NULL` if not required.
`first`	First row whose dual value to return.
`last`	Last row whose dual value to return.

Related topics

`XPRSgetcallbackpresolvesolution`, `XPRSgetcallbackpresolveslacks`, `XPRSgetcallbackpresolveredcosts`, `XPRSgetcallbackduals`, `XPRSgetduals`.

Further information

1. Row indices should be in the range `0` to `ROWS-1`.
2. See `XPRSgetcallbacksolution` for a list of callbacks where this function can be called. Note that a dual solution is not available in all contexts where a primal solution is available.

XPRSgetcallbackpresolveredcosts

Purpose

Returns the reduced costs from the solution to the presolved problem associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbackpresolveredcosts(XPRSprob prob, int
      *p_available, double djs[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a dual solution is available. May be `NULL` if not required.
`djs`	Double array of length `last-first+1` where the reduced costs of the variables will be returned. May be `NULL` if not required.
`first`	First column whose reduced cost to return.
`last`	Last column whose reduced cost to return.

Related topics

`XPRSgetcallbackpresolvesolution`, `XPRSgetcallbackpresolveslacks`, `XPRSgetcallbackpresolveduals`, `XPRSgetcallbackredcosts`, `XPRSgetredcosts`.

Further information

1. Column indices should be in the range `0` to `COLS-1`.
2. See `XPRSgetcallbacksolution` for a list of callbacks where this function can be called. Note that reduced costs are not available in all contexts where a primal solution is available.

XPRSgetcallbackpresolveslacks

Purpose

Returns the slack values from the solution to the presolved problem associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbackpresolveslacks(XPRSprob prob, int *p_available,
      double slacks[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a solution is available. May be `NULL` if not required.
`slacks`	Double array of length `last-first+1` where the values of the slack variables will be returned. May be `NULL` if not required.
`first`	First row whose slack value to return.
`last`	Last row whose slack value to return.

Related topics

`XPRSgetcallbackpresolvesolution`, `XPRSgetcallbackpresolveduals`, `XPRSgetcallbackpresolveredcosts`, `XPRSgetcallbackslacks`, `XPRSgetslacks`.

Further information

1. Row indices should be in the range `0` to `ROWS-1`.
2. See `XPRSgetcallbacksolution` for a list of callbacks where this function can be called.

XPRSgetcallbackpresolvesolution

Purpose

Returns the solution to the presolved problem associated with the current callback.

Topic areas

Solution, Callback

Synopsis

```
int XPRS_CC XPRSgetcallbackpresolvesolution(XPRSprob prob, int
        *p_available, double x[], int first, int last);
```

Arguments

`prob`	The current problem.
`p_available`	This variable will be set to 1 if a solution is available. May be `NULL` if not required.
`x`	Double array of length `last-first+1` where the values of the primal variables will be returned. May be `NULL` if not required.
`first`	First column in the solution to return.
`last`	Last column in the solution to return.

Related topics

`XPRSgetcallbackpresolveslacks`, `XPRSgetcallbackpresolveduals`, `XPRSgetcallbackpresolveredcosts`, `XPRSgetcallbacksolution`, `XPRSgetsolution`.

Further information

1. Column indices should be in the range `0` to `COLS-1`.
2. See `XPRSgetcallbacksolution` for a list of callbacks where this function can be called.

Example

See also example `mostviolated.c`.

XPRSgetcheckedmode

Purpose

You can use this function to interrogate whether checking and validation of all Optimizer function calls is enabled for the current process. Checking and validation is enabled by default but can be disabled by `XPRSsetcheckedmode`.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSgetcheckedmode(int* p_checkedmode);
```

Argument

`p_checkedmode` Variable that is set to 0 if checking and validation of Optimizer function calls is disabled for the current process, non-zero otherwise.

Related topics

`XPRSsetcheckedmode`.

XPRSgetcoef

Purpose

Returns a single coefficient in the constraint matrix.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetcoef(XPRSprob prob, int row, int col, double *p_coef);
```

Arguments

`prob`	The current problem.
`row`	Row of the constraint matrix.
`col`	Column of the constraint matrix.
`p_coef`	Pointer to a double where the coefficient will be returned.

Further information

It is quite inefficient to get several coefficients with the `XPRSgetcoef` function. It is better to use `XPRSgetcols` or `XPRSgetrows`.

Related topics

`XPRSgetcols`, `XPRSgetrows`.

XPRSgetcols, XPRSgetcols64

Purpose

Returns the nonzeros in the constraint matrix for the columns in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetcols(XPRSprob prob, int start[], int rowind[], double
      rowcoef[], int maxcoefs, int *p_ncoefs, int first, int last);

int XPRS_CC XPRSgetcols64(XPRSprob prob, XPRSint64 start[], int rowind[],
      double rowcoef[], XPRSint64 maxcoefs, XPRSint64 *p_ncoefs, int first,
      int last);
```

Arguments

`prob`	The current problem.
`start`	Integer array which will be filled with the indices indicating the starting offsets in the `rowind` and `rowcoef` arrays for each requested column. It must be of length at least `last-first+2`. Column `i` starts at position `start[i]` in the `rowind` and `rowcoef` arrays, and has `start[i+1]-start[i]` elements in it. May be `NULL` if not required.
`rowind`	Integer array of length `maxcoefs` which will be filled with the row indices of the nonzero coefficents for each column. May be `NULL` if not required.
`rowcoef`	Double array of length `maxcoefs` which will be filled with the nonzero coefficient values. May be `NULL` if not required.
`maxcoefs`	The size of the `rowind` and `rowcoef` arrays. This is the maximum number of nonzero coefficients that the Optimizer is allowed to return.
`p_ncoefs`	Pointer to an integer where the number of nonzero coefficients in the selected columns will be returned. If `p_ncoefs` exceeds `maxcoefs`, only the `maxcoefs` first nonzero coefficients will be returned. May be `NULL` if not required.
`first`	First column in the range.
`last`	Last column in the range.

Example

The following examples retrieves the number of nonzero coefficients in all columns of the problem:

```
int ncoefs, cols, first = 0, last;
...
XPRSgetintattrib(prob,XPRS_COLS,&cols);
last = cols-1;
XPRSgetcols(prob,NULL,NULL,NULL,0,&ncoefs,first,last);
```

See also example `tableau.c`.

Further information

It is possible to obtain just the number of elements in the range of columns by replacing `start`, `rowind` and `rowcoef` by `NULL`, as in the example. In this case, `maxcoefs` must be set to `0` to indicate that the length of arrays passed is zero. This is demonstrated in the example above.

Related topics

`XPRSgetrows`.

XPRSgetcoltype

Purpose

Returns the column types for the columns in a given range.

Topic areas

Problem Information, MIP Entities

Synopsis

```
int XPRS_CC XPRSgetcoltype(XPRSprob prob, char coltype[], int first, int
      last);
```

Arguments

`prob`	The current problem.
`coltype`	Character array of length `last-first+1` where the column types will be returned: `C` indicates a continuous variable; `I` indicates an integer variable; `B` indicates a binary variable; `S` indicates a semi-continuous variable; `R` indicates a semi-continuous integer variable; `P` indicates a partial integer variable.
`first`	First column in the range.
`last`	Last column in the range.

Example

This example finds the types for all columns in the matrix and prints them to the console:

```
int cols, i;
char *types;
...
XPRSgetintattrib(prob,XPRS_COLS,&cols);
types = (char *)malloc(sizeof(char)*cols);
XPRSgetcoltype(prob,types,0,cols-1);

for(i=0;i<cols;i++) printf("%c\n",types[i]);
```

Related topics

`XPRSchgcoltype`, `XPRSgetrowtype`.

XPRSgetcontrolinfo

Purpose

Accesses the id number and the type information of a control given its name. A control name may be for example `XPRS_PRESOLVE`. Names are case-insensitive and may or may not have the `XPRS_` prefix. The id number is the constant used to identify the control for calls to functions such as `XPRSgetintcontrol`.

The function will return an id number of `0` and a type value of `XPRS_TYPE_NOTDEFINED` if the name is not recognized as a control name. Note that this will occur if the name is an attribute name and not a control name.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetcontrolinfo(XPRSprob prob, const char* name, int* p_id,
       int* p_type);
```

Arguments

`prob`	The current problem.
`name`	The name of the control to be queried. Names are case-insensitive and may or may not have the `XPRS_` prefix. A full list of all controls may be found in 9, or from the list in the `xprs.h` header file.
`p_id`	Pointer to an integer where the id number will be returned.
`p_type`	Pointer to an integer where the type information will be returned. The returned value will be one of the following constants from `xprs.h`:

XPRS_TYPE_NOTDEFINED	The name was not recognized.
XPRS_TYPE_INT	32-bit integer.
XPRS_TYPE_INT64	64-bit integer.
XPRS_TYPE_DOUBLE	Double precision floating point.
XPRS_TYPE_STRING	String.

Example

The following code example obtains the id number and the type information of the control or attribute with name given by `name`. Note that the name happens to be a control name in this example:

```
const char *name = "presolve";
int id, type;
...
if(XPRSgetattribinfo(prob, name, &id,
                     &type) || id==0) {
  if(XPRSgetcontrolinfo(prob, name, &id,
                        &type) || id==0) {
    printf("Unrecognized name: %s\n", name);
  }
}
```

Related topics

`XPRSgetattribinfo`.

XPRSgetcpcutlist

Purpose

Returns a list of cut indices from the cut pool.

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSgetcpcutlist(XPRSprob prob, int cuttype, int interp, double
      delta, int *p_ncuts, int maxcuts, XPRScut cutind[], double viol[]);
```

Arguments

`prob`	The current problem.
`cuttype`	The user defined type of the cuts to be returned.
`interp`	Way in which the cut type is interpreted: `-1` get all cuts; `1` treat cut types as numbers; `2` treat cut types as bit-vectors (compare Section 9.2) - get cut if any bit matches any bit set in `cuttype`; `3` treat cut types as bit-vectors (compare Section 9.2) - get cut if all bits match those set in `cuttype`.
`delta`	Only those cuts with a signed violation greater than delta will be returned.
`p_ncuts`	Pointer to the integer where the number of cuts of type `cuttype` in the cut pool will be returned.
`maxcuts`	Maximum number of cuts to be returned.
`cutind`	Array of length `maxcuts` where the pointers to the cuts will be returned.
`viol`	Double array of length `maxcuts` where the values of the signed violations of the cuts will be returned.

Further information

1. The violated cuts can be obtained by setting the `delta` parameter to the size of the (signed) violation required. If unviolated cuts are required as well, `delta` may be set to `XPRS_MINUSINFINITY` which is defined in the library header file.
2. If the number of active cuts is greater than `maxcuts`, only `maxcuts` cuts will be returned and `p_ncuts` will be set to the number of active cuts. If `p_ncuts` is less than `maxcuts`, then only `p_ncuts` positions will be filled in `cutind`.
3. In case of a cut of type `'L'`, the violation equals the negative of the slack associated with the row of the cut. In case of a cut of type `'G'`, the violation equals the slack associated with the row of the cut. For cuts of type `'E'`, the violation equals the absolute value of the slack.
4. Please note that the violations returned are absolute violations, while feasibility is checked by the Optimizer in the scaled problem.

Related topics

`XPRSdelcpcuts`, `XPRSgetcpcuts`, `XPRSgetcutlist`, `XPRSloadcuts`, `XPRSgetcutmap`, `XPRSgetcutslack`, Section 5.9.

XPRSgetcpcuts, XPRSgetcpcuts64

Purpose

Returns cuts from the cut pool. A list of cut pointers in the array `rowind` must be passed to the routine. The columns and elements of the cut will be returned in the regions pointed to by the `colind` and `cutcoef` parameters. The columns and elements will be stored contiguously and the starting point of each cut will be returned in the region pointed to by the `start` parameter.

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSgetcpcuts(XPRSprob prob, const XPRScut rowind[], int ncuts,
      int maxcoefs, int cuttype[], char rowtype[], int start[], int
      colind[], double cutcoef[], double rhs[]);

int XPRS_CC XPRSgetcpcuts64(XPRSprob prob, const XPRScut rowind[], int
      ncuts, XPRSint64 maxcoefs, int cuttype[], char rowtype[], XPRSint64
      start[], int colind[], double cutcoef[], double rhs[]);
```

Arguments

`prob`	The current problem.
`rowind`	Array of length `ncuts` containing the pointers to the cuts.
`ncuts`	Number of cuts to be returned.
`maxcoefs`	Maximum number of column indices of the cuts to be returned.
`cuttype`	Integer array of length at least `ncuts` where the cut types will be returned. May be NULL if not required.
`rowtype`	Character array of length at least `ncuts` where the sense of the cuts (`L`, `G`, or `E`) will be returned. May be NULL if not required.
`start`	Integer array of length at least `ncuts+1` containing the offsets into the `colind` and `cutcoef` arrays. The last element indicates where cut `ncuts+1` would start. May be NULL if not required.
`colind`	Integer array of length `maxcoefs` where the column indices of the cuts will be returned. May be NULL if not required.
`cutcoef`	Double array of length `maxcoefs` where the matrix values will be returned. May be NULL if not required.
`rhs`	Double array of length at least `ncuts` where the right hand side elements for the cuts will be returned. May be NULL if not required.

Example

The following example gets the first two cuts:

```
int cuttype[2], start[3];
int *colind;
int rowind[] = { 0, 1 };
double rhs[2];
double *cutcoef;
char * rowtype;
...
XPRSgetcpcuts(prob,rowind,2,0,NULL,NULL,start,NULL,NULL,NULL);
colind = (int*) malloc(start[2]*sizeof(int));
cutcoef = (double*) malloc(start[2]*sizeof(double));
XPRSgetcpcuts(prob,rowind,2,0,cuttype,rowtype,start,colind,cutcoef,rhs);
```

Further information

It is possible to obtain just the number of nonzeros in the range of queried cuts by calling the functions with all output arays except for start equaling `NULL` and checking the value of `start[ncuts]`. In this case, `maxcoefs` must be set to `0` to indicate that the length of arrays passed is `0`.

Related topics

`XPRSgetcpcutlist`, `XPRSgetcutlist`, 5.9.

XPRSgetcutlist

Purpose
Retrieves a list of cut pointers for the cuts active at the current node.

Topic area
Cuts

Synopsis
```
int XPRS_CC XPRSgetcutlist(XPRSprob prob, int cuttype, int interp, int
      *p_ncuts, int maxcuts, XPRScut cutind[]);
```

Arguments

`prob`	The current problem.
`cuttype`	User defined type of the cuts to be returned. A value of `-1` indicates return all active cuts.
`interp`	Way in which the cut type is interpreted: `-1` get all cuts; `1` treat cut types as numbers; `2` treat cut types as bit-vectors (compare Section 9.2) - get cut if any bit matches any bit set in `cuttype`; `3` treat cut types as bit-vectors (compare Section 9.2) - get cut if all bits match those set in `cuttype`.
`p_ncuts`	Pointer to the integer where the number of active cuts of type `cuttype` will be returned.
`maxcuts`	Maximum number of cuts to be retrieved.
`cutind`	Array of length `maxcuts` where the pointers to the cuts will be returned.

Further information
If the number of active cuts is greater than `maxcuts`, then `maxcuts` cuts will be returned and `p_ncuts` will be set to the number of active cuts. If `p_ncuts` is less than `maxcuts`, then only `p_ncuts` positions will be filled in `cutind`.

Related topics
`XPRSgetcpcutlist`, `XPRSgetcpcuts`, Section 5.9.

XPRSgetcutmap

Purpose

Used to return in which rows a list of cuts are currently loaded into the Optimizer. This is useful for example to retrieve the duals associated with active cuts.

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSgetcutmap(XPRSprob prob, int ncuts, const XPRScut cutind[],
      int cutmap[]);
```

Arguments

`prob`	The current problem.
`ncuts`	Number of cuts in the cutind array.
`cutind`	Pointer array to the cuts for which the row index is requested.
`cutmap`	Integer array of length `ncuts`, where the row indices are returned.

Further information

For cuts currently not loaded into the problem, a row index of `−1` is returned.

Related topics

`XPRSgetcpcutlist`, `XPRSdelcpcuts`, `XPRSgetcutlist`, `XPRSloadcuts`, `XPRSgetcutslack`, `XPRSgetcpcuts`, Section 5.9.

XPRSgetcutslack

Purpose

Used to calculate the slack value of a cut with respect to the current LP relaxation solution. The slack is calculated from the cut itself, and might be requested for any cut (even if it is not currently loaded into the problem).

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSgetcutslack(XPRSprob prob, XPRScut cutind, double*
        p_slack);
```

Arguments

`prob`	The current problem.
`cutind`	Pointer of the cut for which the slack is to be calculated.
`p_slack`	Double pointer where the value of the slack is returned.

Related topics

`XPRSgetcpcutlist`, `XPRSdelcpcuts`, `XPRSgetcutlist`, `XPRSloadcuts`, `XPRSgetcutmap`, `XPRSgetcpcuts`, Section 5.9.

XPRSgetdaysleft

Purpose

Returns the number of days left until the license expires.

Topic area

Licensing

Synopsis

```
int XPRS_CC XPRSgetdaysleft(int *p_daysleft);
```

Argument

`p_daysleft` Pointer to an integer where the number of days is to be returned. For a permanent license, the return value will be `XPRS_MAXINT`

Example

The following calls `XPRSgetdaysleft` to print information about the license:

```
int daysleft;
...
XPRSinit(NULL);
if(XPRSgetdaysleft(&daysleft) != 0) {
  printf("An error occurred\n");
} else if (daysleft==XPRS_MAXINT) {
  printf("License will never expire\n");
} else {
  printf("License expires in %d days\n", daysleft);
}
```

Related topics

`XPRSgetbanner`.

XPRSgetdblattrib

Purpose

Enables users to retrieve the values of various double problem attributes. Problem attributes are set during loading and optimization of a problem.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetdblattrib(XPRSprob prob, int attrib, double *p_value);
```

Arguments

`prob`	The current problem.
`attrib`	Problem attribute whose value is to be returned. A full list of all available problem attributes may be found in Chapter 10, or from the list in the `xprs.h` header file.
`p_value`	Pointer to a double where the value of the problem attribute will be returned.

Example

The following obtains the optimal value of the objective function and displays it to the console:

```
double lpobjval;
...
XPRSlpoptimize(prob,"");
XPRSgetdblattrib(prob,XPRS_LPOBJVAL,&lpobjval);
printf("The maximum profit is %f\n",lpobjval);
```

See also examples `addmipsol.c`, `fixbv.c`, `globjpar.c`, `glrhspar.c`, `knapsack.c`, `loadlp.c`, `roundint.c`, `savesol.c`, `tableau.c`, `trimloss.c` .

Related topics

`XPRSgetintattrib`, `XPRSgetstrattrib`.

XPRSgetdblcontrol

Purpose

Retrieves the value of a given double control parameter.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetdblcontrol(XPRSprob prob, int control, double *p_value);
```

Arguments

`prob`	The current problem.
`control`	Control parameter whose value is to be returned. A full list of all controls may be found in Chapter 9, or from the list in the `xprs.h` header file.
`p_value`	Pointer to the location where the control value will be returned.

Example

The following returns the integer feasibility tolerance:

```
XPRSgetdblcontrol(prob,XPRS_MIPTOL,&miptol);
```

See also examples `addmipsol.c`, `knapsack.c`, `mostviolated.c` .

Related topics

`XPRSsetdblcontrol`, `XPRSgetintcontrol`, `XPRSgetstrcontrol`.

XPRSgetdirs

Purpose

Used to return the directives that have been loaded into a matrix. Priorities, forced branching directions and pseudo costs can be returned. If called after presolve, `XPRSgetdirs` will get the directives for the presolved problem.

Topic areas

Problem Information, Branching

Synopsis

```
int XPRS_CC XPRSgetdirs(XPRSprob prob, int *p_ndir, int indices[], int
       prios[], char branchdirs[], double uppseudo[], double downpseudo[]);
```

Arguments

`prob`	The current problem.
`p_ndir`	Pointer to an integer where the number of directives will be returned.
`indices`	Integer array of length `p_ndir` containing the column numbers (`0`, `1`, `2`,...) or negative values corresponding to special ordered sets (the first set numbered `-1`, the second numbered `-2`,...). May be NULL if not required.
`prios`	Integer array of length `p_ndir` containing the priorities for the columns and sets, where columns/sets with smallest priority will be branched on first. May be NULL if not required.
`branchdirs`	Character array of length `p_ndir` specifying the branching direction for each column or set: `U` the entity is to be forced up; `D` the entity is to be forced down; `N` not specified. May be NULL if not required.
`uppseudo`	Double array of length `p_ndir` containing the up pseudo costs for the columns and sets. May be NULL if not required.
`downpseudo`	Double array of length `p_ndir` containing the down pseudo costs for the columns and sets. May be NULL if not required.

Further information

The value `p_ndir` denotes the number of directives, at most `MIPENTS`, obtainable with `XPRSgetintattrib(prob,XPRS_MIPENTS,& mipents);`.

Related topics

`XPRSloaddirs`, `XPRSloadpresolvedirs`.

XPRSgetdualray

Purpose

Retrieves a dual ray (dual unbounded direction) for the current problem, if the problem is found to be infeasible.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetdualray(XPRSprob prob, double ray[], int *p_hasray);
```

Arguments

`prob`	The current problem.
`ray`	Double array of length `ROWS` to hold the ray. May be NULL if not required.
`p_hasray`	This variable will be set to 1 if the Optimizer is able to return a dual ray, 0 otherwise.

Example

The following code tries to retrieve a dual ray:

```
int rows;
double *dualRay;
int hasray;
...
XPRSgetintattrib(prob, XPRS_ROWS, &rows);
dualRay = malloc(rows*sizeof(double));
XPRSgetdualray(prob, dualRay, &hasray);
if(!hasray) printf("Could not retrieve a dual ray\n");
```

Further information

1. It is possible to retrieve a dual ray only when, after solving an LP problem, the final status (`LPSTATUS`) is `XPRS_LP_INFEAS`.
2. Dual rays are not post-solved. If the problem is in a presolved state, the dual ray that is returned will be for the presolved problem. If the problem was solved with presolve on and has been restored to the original state (the default behavior), this function will not be able to return a ray. To ensure that a dual ray can be obtained, it is recommended to solve a problem with presolve turned off (`PRESOLVE` = 0).

Related topics

`XPRSgetprimalray`.

XPRSgetgencons, XPRSgetgencons64

Purpose

Returns the general constraints `y = f(x1, ..., xn, c1, ..., cm)` in a given range.

Topic areas

Problem Information, Piecewise Linear and General Constraints

Synopsis

```
int XPRS_CC XPRSgetgencons(XPRSprob prob, int contype[], int resultant[],
      int colstart[], int colind[], int maxcols, int *p_ncols, int
      valstart[], double val[], int maxvals, int *p_nvals, int first, int
      last);

int XPRS_CC XPRSgetgencons64(XPRSprob prob, int contype[], int resultant[],
      XPRSint64 colstart[], int colind[], XPRSint64 maxcols, XPRSint64
      *p_ncols, XPRSint64 valstart[], double val[], XPRSint64 maxvals,
      XPRSint64 *p_nvals, int first, int last);
```

Arguments

`prob`	The current problem.
`contype`	`NULL` if not required or an integer array of length at least `last-first+1` which will be filled with the types of the general constraints: `XPRS_GENCONS_MAX (0)` indicates a `maximum` constraint; `XPRS_GENCONS_MIN (1)` indicates a `minimum` constraint; `XPRS_GENCONS_AND (2)` indicates an `and` constraint. `XPRS_GENCONS_OR (3)` indicates an `or` constraint; `XPRS_GENCONS_ABS (4)` indicates an `absolute value` constraint.
`resultant`	Integer array which will be filled with the indices of the output variables `y`. It must be of length at least `last-first+1`. May be `NULL` if not required.
`colstart`	Integer array of length at least `last-first+2` which will be filled with the start index of each general constraint in the `colind` array. May be `NULL` if not required.
`colind`	Integer array which will be filled with the indices of the input variables `xi`. May be `NULL` if not required.
`maxcols`	Maximum number of input columns to be retrieved.
`p_ncols`	Pointer to return the number of input columns in the `colind` array. If the number of input columns is greater than `maxcols`, then only `maxcols` elements will be returned. May be `NULL` if not required.
`valstart`	Integer array of length at least `last-first+2` which will be filled with the start index of each general constraint in the `val` array. May be `NULL` if not required.
`val`	Integer array which will be filled with the constant values `ci`. May be `NULL` if not required.
`maxvals`	Maximum number of constant values to be retrieved.
`p_nvals`	Pointer to return the number of constant values in the `val` array. If the number of constant values is greater than `maxvals`, then only `maxvals` elements will be returned.May be `NULL` if not required.
`first`	First general constraint in the range.
`last`	Last general constraint in the range.

Example

The following example retrieves all general constraints:

```
int ngencons;
int *contype;
int *resultant;
int *colstart;
int *colind;
int maxcols;
int ncols;
int *valstart;
int *val;
int maxvals;
int nvals;
...
XPRSgetdblattrib(prob, XPRS_GENCONS, &ngencons);
XPRSgetgencons(prob, NULL, NULL, NULL, NULL, 0, &maxcols, NULL, NULL, 0, &maxvals, 0, ngencons
contype = (int*) malloc(ngencons*sizeof(int));
resultant = (int*) malloc(ngencons*sizeof(int));
colstart = (int*) malloc((ngencons+1)*sizeof(int));
colind = (int*) malloc(maxcols*sizeof(int));
valstart = (int*) malloc((ngencons+1)*sizeof(int));
val = (double*) malloc(maxvals*sizeof(double));
XPRSgetgencons(prob, contype, resultant, colstart, colind, maxcols, &ncols, valstart, val, max
...
```

Further information

1. It is possible to obtain just the number of input columns and/or constant values in the range of general constraints by calling this function with `maxcols` and `maxvals` set to `0`, in which case the required size for the arrays will be returned in `p_ncols` and `p_nvals` (one of them may be `NULL` if only the other is required). In this case it is also possible to query `contype` and `resultant` or they can be left `NULL`.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddgencons`, `XPRSdelgencons`.

XPRSgetglobal, XPRSgetglobal64

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSgetmipentities` instead.

Topic areas

Problem Information, MIP Entities

Synopsis

```
int XPRS_CC XPRSgetglobal(XPRSprob prob, int *p_nentities, int *p_nsets,
      char coltype[], int colind[], double limit[], char settype[], int
      start[], int setcols[], double refval[]);

int XPRS_CC XPRSgetglobal64(XPRSprob prob, int *p_nentities, int *p_nsets,
      char coltype[], int colind[], double limit[], char settype[],
      XPRSint64 start[], int setcols[], double refval[]);
```

XPRSgetmipentities, XPRSgetmipentities64

Purpose

Retrieves integr and entity information about a problem. It must be called before `XPRSmipoptimize` if the presolve option is used.

Topic areas

Problem Information, MIP Entities

Synopsis

```
int XPRS_CC XPRSgetmipentities(XPRSprob prob, int *p_nentities, int
       *p_nsets, char coltype[], int colind[], double limit[], char
       settype[], int start[], int setcols[], double refval[]);

int XPRS_CC XPRSgetmipentities64(XPRSprob prob, int *p_nentities, int
       *p_nsets, char coltype[], int colind[], double limit[], char
       settype[], XPRSint64 start[], int setcols[], double refval[]);
```

Arguments

`prob` The current problem.

`p_nentities` Pointer to the integer where the number of binary, integer, semi-continuous, semi-continuous integer and partial integer entities will be returned. This is equal to the problem attribute `MIPENTS`. May be `NULL`.

`p_nsets` Pointer to the integer where the number of SOS1 and SOS2 sets will be returned. It can be retrieved from the problem attribute `SETS`. May be `NULL`.

`coltype` Character array of length `p_nentities` where the entity types will be returned. The types will be one of:

- `B` binary variables;
- `I` integer variables;
- `P` partial integer variables;
- `S` semi-continuous variables;
- `R` semi-continuous integer variables.

`colind` Integer array of length `p_nentities` where the column indices of the MIP entities will be returned.

`limit` Double array of length `p_nentities` where the limits for the partial integer variables and lower bounds for the semi-continuous and semi-continuous integer variables will be returned (any entries in the positions corresponding to binary and integer variables will be meaningless).

`settype` Character array of length `p_nsets` where the set types will be returned. The set types will be one of:

- `1` SOS1 type sets;
- `2` SOS2 type sets.

`start` Integer array where the offsets into the `setcols` and `refval` arrays indicating the start of the sets will be returned. This array must be of length `p_nsets+1`, the final element will contain the offset where set `p_nsets+1` would start and equals the length of the `setcols` and `refval` arrays, `SETMEMBERS`.

`setcols` Integer array of length `SETMEMBERS` where the columns in each set will be returned.

`refval` Double array of length `SETMEMBERS` where the reference row entries for each member of the sets will be returned. These define the order for SOS2 constraints and may be used in branching for both types.

Example

The following obtains the MIP entities and their types in the arrays `colind` and `qrtype`:

```
int nentities, nsets, *colind;
char *coltype;
...
XPRSgetmipentities(prob,&nentities,&nsets,NULL,NULL,NULL,NULL,
              NULL,NULL,NULL);
colind = malloc(nentities*sizeof(int));
coltype = malloc(nentities*sizeof(char));
XPRSgetmipentities(prob,&nentities,&nsets,coltype,ngcols,NULL,
              NULL,NULL,NULL,NULL);
```

See also examples addmipsol.c, fixbv.c, mostviolated.c, repair.c, roundint.c .

Further information

Any of the arguments except prob, p_nentities and p_nsets may be NULL if not required.

Related topics

XPRSloadmip, XPRSloadmiqp.

XPRSgetiisdata

Purpose

Returns information for an Irreducible Infeasible Set: size, variables and constraints (row and column vectors), and conflicting sides of the variables. For pure linear problems there is also information on duals, reduced costs and isolations.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSgetiisdata(XPRSprob prob, int iis, int *p_nrows, int
       *p_ncols, int rowind[], int colind[], char contype[], char bndtype[],
       double duals[], double djs[], char isolationrows[], char
       isolationcols[]);
```

Arguments

`prob`	The current problem.
`iis`	The ordinal number of the IIS to get data for.
`p_nrows`	Pointer to an integer where the number of rows in the IIS will be returned.
`p_ncols`	Pointer to an integer where the number of bounds in the IIS will be returned.
`rowind`	Indices of rows in the IIS. Can be NULL if not required.
`colind`	Indices of bounds (columns) in the IIS. Can be NULL if not required.
`contype`	Sense of rows in the IIS: `L` for less or equal row; `G` for greater or equal row. `E` for an equality row (for a non LP IIS); `1` for a SOS1 row; `2` for a SOS2 row; `W` for a piecewise linear constraint; `X` for a general constraint; `I` for an indicator row. Can be NULL if not required.
`bndtype`	Sense of bound in the IIS: `U` for upper bound; `L` for lower bound. `F` for fixed columns (for a non LP IIS); `B` for a binary column; `I` for an integer column; `P` for a partial integer columns; `S` for a semi-continuous column; `R` for a semi-continuous integer column. Can be NULL if not required.
`duals`	The dual multipliers associated with the rows. Can be NULL if not required.
`djs`	The dual multipliers (reduced costs) associated with the bounds. Can be NULL if not required.
`isolationrows`	The isolation status of the rows: `-1` if isolation information is not available for row (run iis isolations); `0` if row is not in isolation; `1` if row is in isolation. Can be NULL if not required.

`isolationcols` The isolation status of the bounds:

- `-1` if isolation information is not available for column (run iis isolations);
- `0` if column is not in isolation;
- `1` if column is in isolation.

Can be NULL if not required.

Example

This example first retrieves the size of IIS 1, then gets the detailed information for the IIS.

```
XPRSgetiisdata(myprob, 1, &nrow, &ncol, NULL, NULL, NULL, NULL,
NULL,NULL,NULL,NULL);

rows = malloc(nrow*sizeof(int));
cols = malloc(ncol*sizeof(int));
contype = malloc(nrow);
bndtype = malloc(ncol);
duals = malloc(nrow*sizeof(double));
djs = malloc(ncol*sizeof(double));
isolationrows = malloc(nrow);
isolationcols = malloc(ncol);
XPRSgetiisdata(myprob, 1, &nrow, &ncol, rows, cols, contype,
               bndtype, duals, djs, isolationrows, isolationcols);
```

Further information

1. Calling `IIS` from the console automatically prints some of the above IIS information to the screen, along with a detailed progress log. Extra information can be printed with the `IIS -p` command.
2. IISs are numbered from 1 to `NUMIIS`. Index number 0 refers to the IIS approximation.
3. If `rowind` and `colind` are both NULL, then only the `p_nrows` and `p_ncols` are returned.
4. The arrays may be NULL if not required. However, arrays `contype`, `duals` and `isolationrows` are only returned if `rowind` is not NULL. Similarly, arrays `bndtype`, `djs` and `isolationcols` are only returned if `colind` is not NULL.
5. All the non NULL arrays should be of length `p_nrows` or `p_ncols`, respectively.
6. For the initial IIS approximation of linear problems (`iis` = 0) the number of rows and columns with a nonzero Lagrange multiplier (dual/reduced cost respectively) are returned. Please note that, in such cases, it might be necessary to call `XPRSiisstatus` to retrieve the necessary size of the return arrays.
7. If there are Special Ordered Sets in the IIS, their number is included in the `rowind` array.
8. For non LP IISs, some column indices may appear more than once in the `colind` array, for example an integrality and a bound restriction for the same column.
9. Duals, reduced cost and isolation information are not available for nonlinear IIS problems, and for those the arrays are filled with zero values if they are provided.

Related topics

`XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisstatus`, `XPRSiiswrite`, `IIS`, Section A.6.

XPRSgetindex

Purpose

Returns the index for a specified row or column name.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetindex(XPRSprob prob, int type, const char *name, int
       *p_index);
```

Arguments

prob	The current problem.
type	XPRS_NAMES_ROW (=1) if row index is required; XPRS_NAMES_COLUMN (=2) if column index is required; XPRS_NAMES_SET (=3) if set index is required; XPRS_NAMES_PWLCONS (=4) if piecewise linear constraint index is required; XPRS_NAMES_GENCONS (=5) if general constraint index is required; XPRS_NAMES_OBJECTIVE (=6) if objective index is required; XPRS_NAMES_USERFUNC (=7) if user function index is required; XPRS_NAMES_INTERNALFUNC (=8) if an internal function index is required.
name	Null terminated string.
p_index	Pointer of the integer where the row or column index number will be returned. A value of -1 will be returned if the row or column does not exist.

Further information

Names are case sensitive, and the names of internal functions are upper case, e.g., COS.

Example

The following example loads problem and checks to see if "n 0203" is the name of a row or column:

```
int seqr, seqc;
...
XPRSreadprob(prob,"problem","");

XPRSgetindex(prob,1,"n   0203", &seqr);
XPRSgetindex(prob,2,"n   0203", &seqc);
if(seqr==-1 && seqc ==-1) printf("n   0203 not there\n");
if(seqr!= -1) printf("n   0203 is row %d\n",seqr);
if(seqc!= -1) printf"n   0203 is column %d\n",seqc);
```

See also examples ComplexUserFunctions.c, globjpar.c, glrhspar.c .

Related topics

XPRSaddnames.

XPRSgetindicators

Purpose

Returns the indicator constraint condition (indicator variable and complement flag) associated to the rows in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetindicators(XPRSprob prob, int colind[], int
       complement[], int first, int last);
```

Arguments

prob	The current problem.
colind	Integer array of length last-first+1 where the column indices of the indicator variables are to be placed. May be NULL.
complement	Integer array of length last-first+1 where the indicator complement flags will be returned: 0 not an indicator constraint (in this case the corresponding entry in the colind array is ignored); 1 for indicator constraints with condition "bin = 1"; -1 for indicator constraints with condition "bin = 0". May be NULL.
first	First row in the range.
last	Last row in the range (inclusive).

Example

The following example retrieves information about all indicator constraints in the matrix and prints a list of their indices.

```
int i, rows;
double *colind, *complement;
...
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
colind = malloc(rows*(sizeof(int)));
complement = malloc(rows*(sizeof(int)));
XPRSgetindicators(prob,colind,complement,0,rows-1);

printf("Indicator rows:");
for(i=0; i<rows; i++) if(complement[i]!=0) printf(" %d", i);
printf("\n");
```

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use XPRSpostsolve to restore the problem to a non-solving state.

Related topics

XPRSsetindicators, XPRSdelindicators.

XPRSgetinfeas

Purpose

Returns a list of infeasible primal and dual variables.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetinfeas(XPRSprob prob, int *p_nprimalcols, int
       *p_nprimalrows, int *p_ndualrows, int *p_ndualcols, int x[], int
       slack[], int duals[], int djs[]);
```

Arguments

`prob`	The current problem.
`p_nprimalcols`	Pointer to an integer where the number of primal infeasible variables is returned.
`p_nprimalrows`	Pointer to an integer where the number of primal infeasible rows is returned.
`p_ndualrows`	Pointer to an integer where the number of dual infeasible rows is returned.
`p_ndualcols`	Pointer to an integer where the number of dual infeasible variables is returned.
`x`	Integer array of length `p_nprimalcols` where the primal infeasible variables will be returned. May be `NULL` if not required.
`slack`	Integer array of length `p_nprimalrows` where the primal infeasible rows will be returned. May be `NULL` if not required.
`duals`	Integer array of length `p_ndualrows` where the dual infeasible rows will be returned. May be `NULL` if not required.
`djs`	Integer array of length `p_ndualcols` where the dual infeasible variables will be returned. May be `NULL` if not required.

Error values

91	A current problem is not available.
422	A solution is not available.

Related controls

Double

`FEASTOL`	Tolerance on RHS.
`OPTIMALITYTOL`	Reduced cost tolerance.

Example

In this example, `XPRSgetinfeas` is first called with nulled integer arrays to get the number of infeasible entries. Then space is allocated for the arrays and the function is again called to fill them in:

```
int nprimalcols, nprimalrows, ndualrows, ndualcols, *x, *slack, *duals, *djs;
...
XPRSgetinfeas(prob, &nprimalcols, &nprimalrows, &ndualrows, &ndualcols,
              NULL, NULL, NULL, NULL);
x = malloc(nprimalcols * sizeof(*x));
slack = malloc(nprimalrows * sizeof(*slack));
duals = malloc(ndualrows * sizeof(*duals));
djs = malloc(ndualcols * sizeof(*djs));
XPRSgetinfeas(prob, &nprimalcols, &nprimalrows, &ndualrows, &ndualcols,
              x, slack, duals, djs);
```

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. To find the infeasibilities in a previously saved solution, the solution must first be loaded into memory with the `XPRSreadbinsol` (`READBINSOL`) function.
3. If any of the last four arguments are set to `NULL`, the corresponding number of infeasibilities is still returned.

Related topics

`XPRSgetscaledinfeas`, `XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisstatus`, `XPRSiiswrite`, `IIS`.

XPRSgetintattrib, XPRSgetintattrib64

Purpose

Enables users to recover the values of various integer problem attributes. Problem attributes are set during loading and optimization of a problem.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetintattrib(XPRSprob prob, int attrib, int *p_value);

int XPRS_CC XPRSgetintattrib64(XPRSprob prob, int attrib, XPRSint64
       *p_value);
```

Arguments

`prob`	The current problem.
`attrib`	Problem attribute whose value is to be returned. A full list of all problem attributes may be found in Chapter 10, or from the list in the `xprs.h` header file.
`p_value`	Pointer to an integer where the value of the problem attribute will be returned.

Example

The following obtains the number of columns in the matrix and allocates space to obtain lower bounds for each column:

```
int cols;
double *lb;
...
XPRSgetintattrib(prob,XPRS_COLS,&cols);
lb = (double *) malloc(sizeof(double)*cols);
XPRSgetlb(prob,lb,0,cols-1);
```

See also examples `addmipsol.c`, `ComplexUserFunctions.c`, `els_managedcuts.c`, `els_usercuts.c`, `fixbv.c`, `globjpar.c`, `glrhspar.c`, `goalprog.c`, `knapsack.c`, `loadlp.c`, `mipsolenum.c`, `mipsolpool.c`, `mostviolated.c`, `Polygon_initialvalue.c`, `Polygon_textformula.c`, `Polygon_tokens.c`, `Polygon_userfunc.c`, `Polygon_userfunc_map.c`, `Polygon_userfunc_mapdelta.c`, `Polygon_userfunc_multimap.c`, `Polygon_userfunc_multimapdelta.c`, `Polygon_userfunc_vecmap.c`, `Polygon_userfunc_vecmapdelta.c`, `repair.c`, `roundint.c`, `savesol.c`, `tableau.c`, `trimloss.c`, `tsp.c` .

Related topics

`XPRSgetdblattrib`, `XPRSgetstrattrib`.

XPRSgetintcontrol, XPRSgetintcontrol64

Purpose

Enables users to recover the values of various integer control parameters

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetintcontrol(XPRSprob prob, int control, int *p_value);

int XPRS_CC XPRSgetintcontrol64(XPRSprob prob, int control, XPRSint64
       *p_value);
```

Arguments

prob	The current problem.
control	Control parameter whose value is to be returned. A full list of all controls may be found in Chapter 9, or from the list in the xprs.h header file.
p_value	Pointer to an integer where the value of the control will be returned.

Example

The following obtains the value of DEFAULTALG and outputs it to screen:

```
int defaultalg;
...
XPRSlpoptimize(prob,"");
XPRSgetintcontrol(prob,XPRS_DEFAULTALG,&defaultalg);
printf("DEFAULTALG is %d\n",defaultalg);
```

Further information

Some control parameters, such as SCALING, are bitmaps. Each bit controls a different behavior. If set, bit 0 has value 1, bit 1 has value 2, bit 2 has value 4, and so on.

Related topics

XPRSsetintcontrol, XPRSgetdblcontrol, XPRSgetstrcontrol.

XPRSgetlastbarsol

Purpose

Used to obtain the last barrier solution values following optimization that used the barrier solver.

Topic areas

Solution, Barrier

Synopsis

```
int XPRS_CC XPRSgetlastbarsol(XPRSprob prob, double x[], double slack[],
        double duals[], double djs[], int *p_status);
```

Arguments

`prob`	The current problem.
`x`	Double array of length `ORIGINALCOLS` where the values of the primal variables will be returned. May be `NULL` if not required.
`slack`	Double array of length `ORIGINALROWS` where the values of the slack variables will be returned. May be `NULL` if not required.
`duals`	Double array of length `ORIGINALROWS` where the values of the dual variables ($c_B^T B^{-1}$) will be returned. May be `NULL` if not required.
`djs`	Double array of length `ORIGINALCOLS` where the reduced cost for each variable ($c^T - c_B^T B^{-1} A$) will be returned. May be `NULL` if not required.
`p_status`	Status of the last barrier solve. Value matches that of XPRS_LPSTATUS should the solve have been stopped immediately after the barrier.

Further information

1. If the barrier solver has not been used, `p_status` will return `XPRS_LP_UNSOLVED`.
2. The barrier solution or the solution candidate is always available if the status is not `XPRS_LP_UNSOLVED`.
3. The last barrier solution is available until the next solve, and is not invalidated by otherwise working with the problem.

Related topics

`XPRSgetsolution`, `XPRSgetcallbacksolution`.

XPRSgetlasterror

Purpose

Returns the error message corresponding to the last error encountered by a library function.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSgetlasterror(XPRSprob prob, char *errmsg);
```

Arguments

`prob`	The current problem.
`errmsg`	A 512 character buffer where the last error message will be returned. If the message is longer than 512 characters then it will be truncated. The message will always be terminated by a `NUL` character.

Example

The following shows how this function might be used in error-checking:

```
void error(XPRSprob myprob, char *function)
{
  char errmsg[512];
  XPRSgetlasterror(myprob,errmsg);
  printf("Function %s did not execute correctly: %s\n",
         function, errmsg);
  XPRSdestroyprob(myprob);
  XPRSfree();
  exit(1);
}
```

where the main function might contain lines such as:

```
XPRSprob prob;
...
if(XPRSreadprob(prob,"myprob",""))
  error(prob,"XPRSreadprob");
```

See also examples `addmipsol.c`, `ComplexUserFunctions.c`, `fixbv.c`, `globjpar.c`, `glrhspar.c`, `goalprog.c`, `knapsack.c`, `loadlp.c`, `mipsolenum.c`, `mipsolpool.c`, `mostviolated.c`, `Polygon_initialvalue.c`, `Polygon_textformula.c`, `Polygon_tokens.c`, `Polygon_userfunc.c`, `Polygon_userfunc_map.c`, `Polygon_userfunc_mapdelta.c`, `Polygon_userfunc_multimap.c`, `Polygon_userfunc_multimapdelta.c`, `Polygon_userfunc_vecmap.c`, `Polygon_userfunc_vecmapdelta.c`, `repair.c`, `roundint.c`, `savesol.c`, `tableau.c`, `trimloss.c`, `tsp.c` .

Related topics

`ERRORCODE`, `XPRSaddcbmessage`, `XPRSsetlogfile`, Chapter 11.

XPRSgetlb

Purpose

Returns the lower bounds for the columns in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetlb(XPRSprob prob, double lb[], int first, int last);
```

Arguments

`prob`	The current problem.
`lb`	Double array of length `last-first+1` where the lower bounds are to be placed.
`first`	First column in the range.
`last`	Last column in the range.

Example

The following example retrieves the lower bounds for the columns of the current problem:

```
int cols;
double *lb;
...
XPRSgetintattrib(prob,XPRS_COLS,&cols);
lb = (double *) malloc(sizeof(double)*cols);
XPRSgetlb(prob,lb,0,cols-1);
```

See also example `repair.c`.

Further information

Values greater than or equal to `XPRS_PLUSINFINITY` should be interpreted as infinite; values less than or equal to `XPRS_MINUSINFINITY` should be interpreted as infinite and negative.

Related topics

`XPRSchgbounds`, `XPRSgetub`.

XPRSgetlicerrmsg

Purpose

Retrieves an error message describing the last licensing error, if any occurred.

Topic area

Licensing

Synopsis

```
int XPRS_CC XPRSgetlicerrmsg(char *buffer, int maxbytes);
```

Arguments

`buffer`	Buffer long enough to hold the error message (including a null terminator).
`maxbytes`	Length of the buffer.

Example

The following calls `XPRSgetlicerrmsg` to find out why `XPRSinit` failed:

```
char message[512];
...
if(XPRSinit(NULL))
{
  XPRSgetlicerrmsg(message,512);
  printf("%s\n", message);
}
```

See also examples `addmipsol.c`, `els_managedcuts.c`, `els_usercuts.c`, `fixbv.c`, `globjpar.c`, `glrhspar.c`, `goalprog.c`, `knapsack.c`, `loadlp.c`, `mipsolenum.c`, `mipsolpool.c`, `mostviolated.c`, `repair.c`, `roundint.c`, `savesol.c`, `tableau.c`, `trimloss.c`, `tsp.c` .

Further information

1. The error message includes an error code, which in case the user wishes to use it is also returned by the function. If there was no licensing error the function returns 0.
2. It's recommended that you pass a buffer of at least 2048 bytes as licensing errors can be qute long. If the error message is too large to fit in the buffer, the first `maxbytes-1` characters will be returned.

Related topics

`XPRSinit`.

XPRSgetlpsol

Purpose

Used to obtain the LP solution values following optimization.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetlpsol(XPRSprob prob, double x[], double slack[], double
      duals[], double djs[]);
```

Arguments

`prob`	The current problem.
`x`	Double array of length `ORIGINALCOLS` where the values of the primal variables will be returned. May be `NULL` if not required.
`slack`	Double array of length `ORIGINALROWS` where the values of the slack variables will be returned. May be `NULL` if not required.
`duals`	Double array of length `ORIGINALROWS` where the values of the dual variables ($c_B^T B^{-1}$) will be returned. May be `NULL` if not required.
`djs`	Double array of length `ORIGINALCOLS` where the reduced cost for each variable ($c^T - c_B^T B^{-1} A$) will be returned. May be `NULL` if not required.

Example

The following sequence of commands will get the LP solution (`x`) at the root node of a MIP and the optimal MIP solution (`y`):

```
int cols;
double *x, *y;
...
XPRSmipoptimize(prob,"l");
XPRSgetintattrib(prob,XPRS_ORIGINALCOLS,&cols);
x = malloc(cols*sizeof(double));
XPRSgetlpsol(prob,x,NULL,NULL,NULL);
XPRSmipoptimize(prob,"");
y = malloc(cols*sizeof(double));
XPRSgetsolution(prob,NULL,y,0,cols-1);
```

See also example `fixbv.c`.

Further information

1. After solving an LP, it is recommended to use `XPRSgetsolution` to retrieve the solution.
2. When an integer solution is found during a tree search, it is always set up as a solution to the current node; therefore the integer solution is available as the current node. It is recommended to retrieved the solution with `XPRSgetcallbacksolution` and `XPRSgetcallbackpresolvesolution`.
3. The function always returns data in the original space, even if the problem is currently in presolved state. Use `XPRSgetpresolvesol` if you need the solution in terms of the presolved model.
4. If the matrix is modified after calling `XPRSlpoptimize`, then the solution will no longer be available.
5. If the problem has been presolved, then `XPRSgetlpsol` returns the solution to the original problem. The only way to obtain the presolved solution is to call the related function, `XPRSgetpresolvesol`.

Related topics

`XPRSgetsolution`, `XPRSgetcallbacksolution`, `XPRSgetpresolvesol`, `XPRSwriteprtsol`, `XPRSwritesol`, `XPRSpostsolvesol`.

XPRSgetlpsolval

Purpose

This subroutine is deprecated and will be removed in a future release. Use `XPRSgetsolution` or `XPRSgetcallbacksolution` and related functions instead.
Used to obtain a single LP solution value following optimization.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetlpsolval(XPRSprob prob, int col, int row, double *p_x,
       double *p_slack, double *p_dual, double *p_dj);
```

Arguments

prob	The current problem.
col	Column index of the variable for which to return the solution value.
row	Row index of the constraint for which to return the solution value.
p_x	Double pointer where the value of the primal variable will be returned. May be NULL if not required.
p_slack	Double pointer where the value of the slack variable will be returned. May be NULL if not required.
p_dual	Double pointer where the value of the dual variable ($c_B^T B^{-1}$) will be returned. May be NULL if not required.
p_dj	Double pointer where the reduced costs for the variable ($c^T - c_B^T B^{-1} A$) will be returned. May be NULL if not required.

Further information

This function is currently not supported if the problem is in a presolved state. Please make sure the problem is postsolved (XPRSpostsolve and XPRSnlppostsolve).

Related topics

XPRSgetsolution, XPRSgetcallbacksolution, XPRSgetcallbackpresolvesolution, XPRSwriteprtsol, XPRSwritesol.

XPRSgetmessagestatus

Purpose

Retrieves the current suppression status of a message.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSgetmessagestatus(XPRSprob prob, int msgcode, int
       *p_status);
```

Arguments

`prob`	The problem to check for the suppression status of the message error code. Use `NULL` to check for the global suppression status of the message msgcode.
`msgcode`	The id number of the message. Refer to Chapter 11 for a list of possible message numbers.
`p_status`	Non-zero if the message is not suppressed; `0` otherwise.

Further information

If a message is suppressed globally then the message will always have `p_status` return zero from `XPRSgetmessagestatus` when prob is non-NULL.

Related topics

`XPRSsetmessagestatus`.

XPRSgetmipsol

Purpose

This subroutine is deprecated and will be removed in a future release. Use `XPRSgetsolution` and `XPRSgetslacks` instead.
Used to obtain the solution values of the last MIP solution that was found.

Topic area

Solution

Synopsis

```
int XPRS_CC XPRSgetmipsol(XPRSprob prob, double x[], double slack[]);
```

Arguments

`prob`	The current problem.
`x`	Double array of length `ORIGINALCOLS` where the values of the primal variables will be returned. May be `NULL` if not required.
`slack`	Double array of length `ORIGINALROWS` where the values of the slack variables will be returned. May be `NULL` if not required.

Example

The following sequence of commands will get the solution (`x`) of the last MIP solution for a problem:

```
int cols;
double *x;
...
XPRSmipoptimize(prob,"");
XPRSgetintattrib(prob,XPRS_ORIGINALCOLS,&cols);
x = malloc(cols*sizeof(double));
XPRSgetmipsol(prob,x,NULL);
```

Further information

1. **Warning:** If allocating space for the MIP solution the row and column sizes must be obtained for the original problem and not for the presolve problem. They can be obtained before optimizing or after calling `XPRSpostsolve` for the case where the tree search has not completed.
2. During a `intsol` or `preintsol` callback, in order to retrieve the corresponding integer solution, use either `XPRSgetcallbacksolution` or `XPRSgetcallbackpresolvesolution`, not `XPRSgetmipsol` (see the documentation of these callbacks for an explanation).
3. After solving the root LP of a MIP with `XPRSmipoptimize(prob, "l")`, you can access the solution by calling `XPRSgetlpsol`.

Related topics

`XPRSgetsolution`, `XPRSgetslacks`, `XPRSgetlpsol`, `XPRSwriteprtsol`, `XPRSwritesol`.

XPRSgetmipsolval

Purpose

This subroutine is deprecated and will be removed in a future release. Use `XPRSgetsolution` and `XPRSgetslacks` instead.
Used to obtain a single solution value of the last MIP solution that was found.

Topic area

Solution

Synopsis

```
int XPRS_CC XPRSgetmipsolval(XPRSprob prob, int col, int row, double *p_x,
        double *p_slack);
```

Arguments

`prob`	The current problem.
`col`	Column index of the variable for which to return the solution value.
`row`	Row index of the constraint for which to return the solution value.
`p_x`	Double pointer where the value of the primal variable will be returned. May be `NULL` if not required.
`p_slack`	Double pointer where the value of the slack variable will be returned. May be `NULL` if not required.

Related topics

`XPRSgetsolution`, `XPRSgetslacks`, `XPRSwriteprtsol`, `XPRSwritesol`.

XPRSgetmqobj, XPRSgetmqobj64

Purpose

Returns the nonzeros in the quadratic objective coefficients matrix for the columns in a given range. To achieve maximum efficiency, `XPRSgetmqobj` returns the lower triangular part of this matrix only.

Topic areas

Problem Information, Quadratic

Synopsis

```
int XPRS_CC XPRSgetmqobj (XPRSprob prob, int start[], int colind[], double
       objqcoef[], int maxcoefs, int *p_ncoefs, int first, int last);

int XPRS_CC XPRSgetmqobj64 (XPRSprob prob, XPRSint64 start[], int colind[],
       double objqcoef[], XPRSint64 maxcoefs, XPRSint64 *p_ncoefs, int
       first, int last);
```

Arguments

`prob`	The current problem.
`start`	Integer array which will be filled with indices indicating the starting offsets in the `colind` and `objqcoef` arrays for each requested column. It must be length of at least `last-first+2`. Column `i` starts at position `start[i]` in the `colind` and `objqcoef` arrays, and has `start[i+1]-start[i]` elements in it. May be NULL if `maxcoefs` is `0`.
`colind`	Integer array of length `maxcoefs` which will be filled with the column indices of the nonzero elements in the lower triangular part of `Q`. May be `NULL` if `maxcoefs` is `0`.
`objqcoef`	Double array of length `maxcoefs` which will be filled with the nonzero element values. May be `NULL` if `maxcoefs` is `0`.
`maxcoefs`	The maximum number of elements to be returned (size of the arrays).
`p_ncoefs`	Pointer to an integer where the number of nonzero quadratic objective coefficients will be returned. If the number of nonzero coefficients is greater than `maxcoefs`, then only `maxcoefs` elements will be returned. If `p_ncoefs` is smaller than `maxcoefs`, then only `p_ncoefs` will be returned. May be `NULL`.
`first`	First column in the range.
`last`	Last column in the range.

Further information

The objective function is of the form $c^T x + 0.5 x^T Q x$. Note that only the upper or lower triangular part of the `Q` matrix is returned.

Related topics

`XPRSchgmqobj`, `XPRSchgqobj`, `XPRSgetqobj`.

XPRSgetobjn

Purpose

For a given objective function, returns the objective coefficients for the columns in a given range.

Topic areas

Problem Information, Multiobjective

Synopsis

```
int XPRS_CC XPRSgetobjn(XPRSprob prob, int objidx, double objcoef[], int
       first, int last);
```

Arguments

`prob`	The current problem.
`objidx`	Index of the objective function whose coefficients to return.
`objcoef`	Double array of length `last-first+1` where the objective function coefficients are to be placed.
`first`	First column in the range.
`last`	Last column in the range.

Example

The following example retrieves the coefficients of the first objective function in the current problem:

```
int cols;
double *objcoef;
...
XPRSgetintattrib(prob,XPRS_COLS,&cols);
objcoef = (double *) malloc(sizeof(double)*cols);
XPRSgetobjn(prob, 0, objcoef, 0, cols-1);
```

Further information

1. This function always returns the objective coefficients from the original problem: it will raise an error if the problem is presolved.
2. The coefficients of the main objective (`objidx=0`) can also be retrieved using `XPRSgetobj`.

Related topics

`XPRSchgobjn`, `XPRSaddobj`, `XPRSdelobj`, `XPRSgetobj`.

XPRSgetobjdblattrib

Purpose

Retrieves the value of a given double attribute associated with a multi-objective solve. When solving a multi-objective problem, several objectives might be optimized in sequence. After each solve, the problem attributes are captured so that they can be queried afterwards.

Topic areas

Controls and Attributes, Multiobjective

Synopsis

```
int XPRS_CC XPRSgetobjdblattrib(XPRSprob prob, int solveidx, int attrib,
       double* p_value);
```

Arguments

`prob`	The current problem.
`solveidx`	Index of the solve to query. Must be between `0` and `SOLVEDOBJS-1`.
`attrib`	Problem attribute whose value is to be returned. A full list of all problem attributes may be found in 10, or from the list in the `xprs.h` header file.
`p_value`	Pointer to a double where attribute value will be returned.

Example

The following obtains the primal dual integral calculated for each solve and displays it to the console:

```
double pdi;
int nsolves, i;
XPRSgetintattrib(prob, XPRS_SOLVEDOBJS, &nsolves);
XPRSlpoptimize(prob,"");
for (i = 0; i < nsolves; i++) {
  XPRSgetobjdblattrib(prob,i,XPRS_PRIMALDUALINTEGRAL,&pdi);
  printf("PDI for solve %d: %f\n",i,pdi);
}
```

Related topics

`XPRSgetobjintattrib`, `XPRScalcobjn`.

Further information

To query the optimal objective values for a multi-objective solve, it is better to use `XPRScalcobjn`. Passing `XPRS_LPOBJVAL` or `XPRS_MIPOBJVAL` to `XPRSgetobjdblattrib` will give the objective value immediately after this objective was solved. The final objective value may be worse due to re-optimizing a subsequent objective. Also, if you are using weighted objective functions with the same priority, `XPRSgetobjdblattrib` will return the value of the weighted objective function for a given solve, whereas `XPRScalcobjn` always returns the unweighted value of a single objective function.

XPRSgetobjdblcontrol

Purpose

Retrieves the value of a given double control parameter associated with an objective function. These parameters control how the objective is treated during multi-objective optimization.

Topic areas

Controls and Attributes, Multiobjective

Synopsis

```
int XPRS_CC XPRSgetobjdblcontrol(XPRSprob prob, int objidx, int control,
        double* p_value);
```

Arguments

prob	The current problem.
objidx	Index of the objective to query.
control	Control parameter whose value is to be returned. Must be one of: XPRS_OBJECTIVE_WEIGHT get the weight of the given objective; XPRS_OBJECTIVE_ABSTOL get the absolute tolerance of the given objective; XPRS_OBJECTIVE_RELTOL get the relative tolerance of the given objective; XPRS_OBJECTIVE_RHS get the constant term of the given objective.
p_value	Pointer to a double where the control value will be returned.

Example

The following obtains the weight of the first objective and outputs it to screen:

```
double weight;
XPRSgetobjdblcontrol(prob,1,XPRS_OBJECTIVE_WEIGHT,&weight);
printf("Weight is %g\n",weight);
```

Related topics

XPRSsetobjdblcontrol, XPRSgetobjintcontrol, XPRSsetobjintcontrol.

XPRSgetobjintattrib, XPRSgetobjintattrib64

Purpose

Retrieves the value of a given integer attribute associated with a multi-objective solve. When solving a multi-objective problem, several objectives might be optimized in sequence. After each solve, the problem attributes are captured so that they can be queried afterwards.

Topic areas

Controls and Attributes, Multiobjective

Synopsis

```
int XPRS_CC XPRSgetobjintattrib(XPRSprob prob, int solveidx, int attrib,
       int* p_value);
int XPRS_CC XPRSgetobjintattrib64(XPRSprob prob, int solveidx, int attrib,
       XPRSint64* p_value);
```

Arguments

prob	The current problem.
solveidx	Index of the solve to query. Must be between 0 and SOLVEDOBJS-1.
attrib	Problem attribute whose value is to be returned. A full list of all problem attributes may be found in 10, or from the list in the xprs.h header file.
p_value	Pointer to an integer where attribute value will be returned.

Example

The following obtains the status of each LP solve, and displays it to the console:

```
int nsolves, status, i;
XPRSlpoptimize(prob,"");
XPRSgetintattrib(prob, XPRS_SOLVEDOBJS, &nsolves);
for (i = 0; i < nsolves; i++) {
  XPRSgetobjintattrib(prob,i,XPRS_LPSTATUS,&status);
  printf("Status of solve %d: %d\n",i,status);
}
```

Related topics

XPRSgetobjdblattrib.

XPRSgetobjintcontrol

Purpose

Retrieves the value of a given integer control parameter associated with an objective. These parameters control how the objective is treated during multi-objective optimization.

Topic areas

Controls and Attributes, Multiobjective

Synopsis

```
int XPRS_CC XPRSgetobjintcontrol(XPRSprob prob, int objidx, int control,
      int* p_value);
```

Arguments

`prob`	The current problem.
`objidx`	Index of the objective to query.
`control`	Control parameter whose value is to be returned. Must be one of: `XPRS_OBJECTIVE_PRIORITY` get the priority of the given objective.
`p_value`	Pointer to an integer where the control value will be returned.

Example

The following obtains the priority of the first objective and outputs it to screen:

```
int priority;
XPRSgetobjintcontrol(prob,1,XPRS_OBJECTIVE_PRIORITY,&priority);
printf("Priority is %d\n",priority);
```

Related topics

`XPRSsetobjintcontrol`, `XPRSgetobjdblcontrol`, `XPRSsetobjdblcontrol`.

XPRSgetnamelist

Purpose

Returns the names for the rows, columns, sets, piecewise linear constraints, general constraints or objectives in a given range. The names will be returned in a character buffer, with no trailing whitespace and with each name being separated by a NULL character.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRSgetnamelist(XPRSprob prob, int type, char names[], int
       maxbytes, int * p_nbytes, int first, int last);
```

Arguments

`prob`	The current problem.
`type`	`XPRS_NAMES_ROW (=1)` if row names are required; `XPRS_NAMES_COLUMN (=2)` if column names are required; `XPRS_NAMES_SET (=3)` if set names are required; `XPRS_NAMES_PWLCONS (=4)` if piecewise linear constraint names are required; `XPRS_NAMES_GENCONS (=5)` if general constraint names are required; `XPRS_NAMES_OBJECTIVE (=6)` if objective function names are required.
`names`	A buffer into which the names will be returned as a sequence of null-terminated strings. The buffer should be of length `maxbytes` bytes. May be NULL if `maxbytes` is `0`.
`maxbytes`	The maximum number of bytes that may be written to the buffer `names`.
`p_nbytes`	A pointer to a variable into which will be written the number of bytes required to contain the names in the specified range. May be NULL if not required.
`first`	First row, column, set, piecewise linear or general constraint in the range.
`last`	Last row, column, set, piecewise linear or general constraint in the range.

Example

The following example retrieves and outputs the row and column names for the current problem.

```
int i, o, cols, rows, cnames_len, rnames_len;
char *cnames, *rnames;
...
/* Get problem size */
XPRSgetintattrib(prob,XPRS_COLS,&cols);
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
/* Request number of bytes required to retrieve the names */
XPRSgetnamelist(prob,XPRS_NAMES_ROW,NULL,0,&rnames_len,0,rows-1);
XPRSgetnamelist(prob,XPRS_NAMES_COLUMN,NULL,0,&cnames_len,0,cols-1);

/* Now allocate buffers big enough then fetch the names */
cnames = (char *) malloc(sizeof(char)*cnames_len);
rnames = (char *) malloc(sizeof(char)*rnames_len);
XPRSgetnamelist(prob,XPRS_NAMES_ROW,rnames,rnames_len,NULL,0,rows-1);
XPRSgetnamelist(prob,XPRS_NAMES_COLUMN,cnames,cnames_len,NULL,0,cols-1);

/* Output row names */
o=0;
for (i=0;i<rows;i++) {
  printf("Row #%d: %s\n", i, rnames+o);
  o += strlen(rnames+o)+1;
}
/* Output column names */
o=0;
```

```
for (i=0;i<cols;i++) {
  printf("Column #%d: %s\n", i, cnames+o);
  o += strlen(cnames+o)+1;
}
```

See also examples `repair.c`, `tableau.c` .

Related topics

`XPRSaddnames`.

XPRSgetnamelistobject

Purpose

This subroutine is deprecated and will be removed in a future release. The names list API is scheduled for removal.

Returns the `XPRSnamelist` object for the rows, columns or sets of a problem. The names stored in this object can be queried using the `XPRS_nml_` functions.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRSgetnamelistobject(XPRSprob prob, int type, XPRSnamelist
      *p_nml);
```

Arguments

`prob`	The current problem.
`type`	`XPRS_NAMES_ROW (=1)` if the row name list is required; `XPRS_NAMES_COLUMN (=2)` if the column name list is required; `XPRS_NAMES_SET (=3)` if the set name list is required; `XPRS_NAMES_PWLCONS (=4)` if piecewise linear constraint name list is required; `XPRS_NAMES_GENCONS (=5)` if general constraint name list is required; `XPRS_NAMES_OBJECTIVE (=6)` if objective name list is required.
`p_nml`	Pointer to a variable holding the name list contained by the problem.

Further information

The `XPRSnamelist` object is a map of names to and from indices.

Related topics

None.

XPRSgetnames

Purpose

This subroutine is deprecated and will be removed in a future release. Use `XPRSgetnamelist` *instead.* Returns the names for the rows, columns, sets, piecewise linear constraints, general constraints or objectives in a given range. The names will be returned in a character buffer, each name being separated by a null character.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XPRSgetnames(XPRSprob prob, int type, char names[], int first,
       int last);
```

Arguments

`prob`	The current problem.
`type`	`XPRS_NAMES_ROW (=1)` if row names are required; `XPRS_NAMES_COLUMN (=2)` if column names are required; `XPRS_NAMES_SET (=3)` if set names are required; `XPRS_NAMES_PWLCONS (=4)` if piecewise linear constraint names are required; `XPRS_NAMES_GENCONS (=5)` if general constraint names are required; `XPRS_NAMES_OBJECTIVE (=6)` if objective function names are required;
`names`	Buffer long enough to hold the names. Since each name is `8*NAMELENGTH` characters long (plus a null terminator), the array, `names`, would be required to be at least as long as `(first-last+1)`*(`8*NAMELENGTH+1`) characters. The names of the row/column/set `first+i` will be written into the `names` buffer starting at position `i*8*NAMELENGTH+i`.
`first`	First row, column, set, piecewise linear or general constraint in the range.
`last`	Last row, column, set, piecewise linear or general constraint in the range.

Further information

The names returned in `names` will be padded with spaces to `8*NAMELENGTH` characters.

Example

The following example retrieves the row and column names of the current problem:

```
int cols, rows, nl;
...
XPRSgetintattrib(prob,XPRS_COLS,&cols);
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
XPRSgetintattrib(prob,XPRS_NAMELENGTH,&nl);

cnames = (char *) malloc(sizeof(char)*(8*nl+1)*cols);
rnames = (char *) malloc(sizeof(char)*(8*nl+1)*rows);
XPRSgetnames(prob,XPRS_NAMES_ROW,rnames,0,rows-1);
XPRSgetnames(prob,XPRS_NAMES_COLUMN,cnames,0,cols-1);
```

To display `names[i]`, use

```
int namelength;
...

XPRSgetintattrib(prob,XPRS_NAMELENGTH,&namelength);
printf("%s",names + i*(8*namelength+1));
```

Related topics

`XPRSaddnames`, `XPRSgetnamelist`.

XPRSgetobj

Purpose

Returns the objective function coefficients for the columns in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetobj(XPRSprob prob, double objcoef[], int first, int
      last);
```

Arguments

`prob`	The current problem.
`objcoef`	Double array of length `last-first+1` where the objective function coefficients are to be placed.
`first`	First column in the range.
`last`	Last column in the range.

Example

The following example retrieves the objective function coefficients of the current problem:

```
int cols;
double *objcoef;
...
XPRSgetintattrib(prob,XPRS_COLS,&cols);
objcoef = (double *) malloc(sizeof(double)*cols);
XPRSgetobj(prob, objcoef, 0, cols-1);
```

See also example `knapsack.c`.

Related topics

`XPRSchgobj`.

XPRSgetobjecttypename

Purpose

Function to access the type name of an object referenced using the generic Optimizer object pointer `XPRSobject`.

Topic areas

Controls and Attributes, Multiobjective

Synopsis

```
int XPRS_CC XPRSgetobjecttypename(XPRSobject xprsobj, const char **p_name);
```

Arguments

`xprsobj`	The object for which the type name will be retrieved.
`p_name`	Pointer to a char pointer returning a reference to the null terminated string containing the object's type name. For example, if the object is of type `XPRSprob` then the returned pointer points to the string "`XPRSprob`".

Further information

This function is intended to be used typically from within the message callback function registered with the `XPRS_ge_addcbmsghandler` function. In such cases the user will need to identify the type of object sending the message since the message callback is passed only a generic pointer to the Optimizer object (`XPRSobject`) sending the message.

Related topics

`XPRS_ge_addcbmsghandler`.

XPRSgetpivotorder

Purpose
Returns the pivot order of the basic variables.

Topic areas
Simplex, LP

Synopsis
```
int XPRS_CC XPRSgetpivotorder(XPRSprob prob, int pivotorder[]);
```

Arguments
`prob`	The current problem.
`pivotorder`	Integer array of length `ROWS` where the pivot order will be returned.

Example
The following returns the pivot order of the variables into an array `pPivot`:
```
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
pPivot = malloc(rows*(sizeof(int)));
XPRSgetpivotorder(prob,pPivot);
```

See also example `tableau.c`.

Further information
Row indices are in the range `0` to `ROWS-1`, whilst columns are in the range `ROWS+SPAREROWS` to `ROWS+SPAREROWS+COLS-1`.

Related topics
`XPRSgetpivots`, `XPRSpivot`.

XPRSgetpivots

Purpose

Returns a list of potential leaving variables if a specified variable enters the basis.

Topic areas

Simplex, LP

Synopsis

```
int XPRS_CC XPRSgetpivots(XPRSprob prob, int enter, int outlist[], double
       x[], double *p_objval, int *p_npivots, int maxpivots);
```

Arguments

`prob`	The current problem.
`enter`	Index of the specified row or column to enter basis.
`outlist`	Integer array of length at least `maxpivots` to hold list of potential leaving variables. May be `NULL` if not required.
`x`	Double array of length `ROWS+SPAREROWS+COLS` to hold the values of all the variables that would result if `enter` entered the basis. May be `NULL` if not required.
`p_objval`	Pointer to a double where the objective function value that would result if `enter` entered the basis will be returned.
`p_npivots`	Pointer to an integer where the actual number of potential leaving variables will be returned.
`maxpivots`	Maximum number of potential leaving variables to return.

Error value

`425`	Indicates `enter` is invalid (out of range or already basic).

Example

The following retrieves a list of up to 5 potential leaving variables if variable 6 enters the basis:

```
int npivots, outlist[5];
double objective;
...
XPRSgetpivots(prob,6,outlist,NULL,&objective,&npivots,5);
```

Further information

1. If the variable `enter` enters the basis and the problem is degenerate then several basic variables are candidates for leaving the basis, and the number of potential candidates is returned enter `p_npivots`. A list of at most `maxpivots` of these candidates is returned enter `outlist` which must be at least `maxpivots` long. If variable `enter` were to be pivoted enter, then because the problem is degenerate, the resulting values of the objective function and all the variables do not depend on which of the candidates from `outlist` is chosen to leave the basis. The value of the objective is returned in `p_objval` and the values of the variables into `x`.
2. Row indices are in the range `0` to `ROWS-1`, whilst columns are in the range `ROWS+SPAREROWS` to `ROWS+SPAREROWS+COLS-1`.

Related topics

`XPRSgetpivotorder`, `XPRSpivot`.

XPRSgetpresolvebasis

Purpose

Returns the current basis from memory into the user's data areas. If the problem is presolved, the presolved basis will be returned. Otherwise the original basis will be returned.

Topic areas

Simplex, LP

Synopsis

```
int XPRS_CC XPRSgetpresolvebasis(XPRSprob prob, int rowstat[], int
       colstat[]);
```

Arguments

`prob`	The current problem.
`rowstat`	Integer array of length `ROWS` to the basis status of the stack, surplus or artificial variable associated with each row. The status will be one of: `XPRS_NONBASIC_LOWER (0)` slack, surplus or artificial is non-basic at lower bound; `XPRS_BASIC (1)` slack, surplus or artificial is basic; `XPRS_NONBASIC_UPPER (2)` slack or surplus is non-basic at upper bound. May be `NULL` if not required.
`colstat`	Integer array of length `COLS` to hold the basis status of the columns in the constraint matrix. The status will be one of: `XPRS_NONBASIC_LOWER (0)` variable is non-basic at lower bound, or superbasic at zero if the variable has no lower bound; `XPRS_BASIC (1)` variable is basic; `XPRS_NONBASIC_UPPER (2)` variable is at upper bound; `XPRS_SUPERBASIC (3)` variable is super-basic. May be `NULL` if not required.

Example

The following obtains and outputs basis information on a presolved problem prior to the tree search:

```
XPRSprob prob;
int i, cols, *colstat;
...
XPRSreadprob(prob,"mymip","");
XPRSmipoptimize(prob,"l");
XPRSgetintattrib(prob,XPRS_COLS,&cols);
colstat = malloc(cols*sizeof(int));
XPRSgetpresolvebasis(prob,NULL,colstat);
for(i=0;i<cols;i++)
 printf("Column %d: %d\n", i, colstat[i]);
XPRSmipoptimize(prob);
```

Related topics

`XPRSgetbasis`, `XPRSloadbasis`, `XPRSloadpresolvebasis`.

XPRSgetpresolvemap

Purpose

Returns the mapping of the row and column numbers from the presolve problem back to the original problem.

Topic areas

Problem Information, Presolve

Synopsis

```
int XPRS_CC XPRSgetpresolvemap(XPRSprob prob, int rowmap[], int colmap[]);
```

Arguments

`prob`	The current problem.
`rowmap`	Integer array of length `ROWS` where the row maps will be returned.
`colmap`	Integer array of length `COLS` where the column maps will be returned.

Example

The following reads in a (Mixed) Integer Programming problem and gets the mapping for the rows and columns back to the original problem following optimization of the linear relaxation. The elimination operations of the presolve are turned off so that a one-to-one mapping between the presolve problem and the original problem.

```
XPRSreadprob(prob,"MyProb","");
XPRSsetintcontrol(prob,XPRS_PRESOLVEOPS,255);
XPRSmipoptimize(prob,"l");
XPRSgetintattrib(prob,XPRS_COLS,&cols);
colmap = malloc(cols*sizeof(int));
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
rowmap = malloc(rows*sizeof(int));
XPRSgetpresolvemap(prob,rowmap,colmap);
```

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. The presolved problem can contain rows or columns that do not map to anything in the original problem. An example of this are cuts created during the MIP solve and temporarily added to the presolved problem. It is also possible that the presolver will introduce new rows or columns. For any row or column that does not have a mapping to a row or column in the original problem, the corresponding entry in the returned `rowmap` and `colmap` arrays will be `-1`.

Related topics

5.3.

XPRSgetpresolvesol

Purpose

Returns the solution for the presolved problem from memory.

Topic areas

Solution, Presolve

Synopsis

```
int XPRS_CC XPRSgetpresolvesol(XPRSprob prob, double x[], double slack[],
        double duals[], double djs[]);
```

Arguments

`prob`	The current problem.
`x`	Double array of length `COLS` where the values of the primal variables will be returned. May be `NULL` if not required.
`slack`	Double array of length `ROWS` where the values of the slack variables will be returned. May be `NULL` if not required.
`duals`	Double array of length `ROWS` where the values of the dual variables will be returned. May be `NULL` if not required.
`djs`	Double array of length `COLS` where the reduced cost for each variable will be returned. May be `NULL` if not required.

Example

The following reads in a (Mixed) Integer Programming problem and displays the solution to the presolved problem following optimization of the linear relaxation:

```
XPRSreadprob(prob,"MyProb","");
XPRSmipoptimize(prob,"l");
XPRSgetintattrib(prob,XPRS_COLS,&cols);
x = malloc(cols*sizeof(double));
XPRSgetpresolvesol(prob,x,NULL,NULL,NULL);
for(i=0;i<cols;i++)
printf("Presolved x(%d) = %g\n",i,x[i]);
XPRSmipoptimize(prob,"");
```

Further information

1. In callbacks, it is recommended to use `XPRSgetcallbackpresolvesolution` instead of `XPRSgetpresolvesol`.
2. This function applies to linear and convex quadratic (including second order conic) problems only.
3. If the problem has not been presolved, the solution in memory will be returned.
4. The solution to the original problem should be returned using the related function `XPRSgetlpsol`.
5. If called during a callback during MIP tree search the solution of the current node will be returned.
6. When an integer solution is found during a tree search, it is always set up as a solution to the current node; therefore the integer solution is available as the current node solution and can be retrieved with `XPRSgetlpsol` and `XPRSgetpresolvesol`.

Related topics

`XPRSgetsolution`, `XPRSgetcallbacksolution`, `XPRSgetlpsol`, 5.3, `XPRSpostsolvesol`.

XPRSgetprimalray

Purpose

Retrieves a primal ray (primal unbounded direction) for the current problem, if the problem is found to be unbounded.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetprimalray(XPRSprob prob, double ray[], int *p_hasray);
```

Arguments

prob	The current problem.
ray	Double array of length COLS to hold the ray. May be NULL if not required.
p_hasray	This variable will be set to 1 if the Optimizer is able to return a primal ray, 0 otherwise.

Example

The following code tries to retrieve a primal ray:

```
int cols;
double *primalRay;
int hasray;
...
XPRSgetintattrib(prob, XPRS_COLS, &cols);
primalRay = malloc(cols*sizeof(double));
XPRSgetprimalray(prob, primalRay, &hasray);
if(!hasray) printf("Could not retrieve a primal ray\n");
```

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. It is possible to retrieve a primal ray only when, after solving an LP problem, the final status (LPSTATUS) is XPRS_LP_UNBOUNDED.
3. Primal rays are not post-solved. If the problem is in a presolved state, the primal ray that is returned will be for the presolved problem. If the problem was solved with presolve on and has been restored to the original state (the default behavior), this function will not be able to return a ray. To ensure that a primal ray can be obtained, it is recommended to solve a problem with presolve turned off (PRESOLVE = 0).

Related topics

XPRSgetdualray.

XPRSgetprobname

Purpose

Returns the current problem name.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetprobname(XPRSprob prob, char *name);
```

Arguments

`prob`	The current problem.
`name`	A buffer of `MAXPROBNAMELENGTH`+1 bytes to contain the current problem name.

Related topics

`XPRSsetprobname`, `MAXPROBNAMELENGTH`.

Example

See also example `addmipsol.c`.

XPRSgetpwlcons, XPRSgetpwlcons64

Purpose

Returns the piecewise linear constraints `y = f(x)` in a given range.

Topic areas

Problem Information, Piecewise Linear and General Constraints

Synopsis

```
int XPRS_CC XPRSgetpwlcons(XPRSprob prob, int colind[], int resultant[],
      int start[], double xval[], double yval[], int maxpoints, int
      *p_npoints, int first, int last);

int XPRS_CC XPRSgetpwlcons64(XPRSprob prob, int colind[], int resultant[],
      XPRSint64 start[], double xval[], double yval[], XPRSint64 maxpoints,
      XPRSint64 *p_npoints, int first, int last);
```

Arguments

`prob`	The current problem.
`colind`	Integer array which will be filled with the indices of the input variables `x`. It must be of length at least `last-first+1`. May be `NULL` if not required.
`resultant`	Integer array which will be filled with the indices of the output variables `y`. It must be of length at least `last-first+1`. May be `NULL` if not required.
`start`	Integer array which will be filled with the start indices of the different constraints in the breakpoint arrays. It must be of length at least `last-first+2`. The `x`-values of the breakpoints of piecewise linear constraint `i < last` will be given in `xval[start[i]]` to `xval[start[i+1]]`. May be `NULL` if not required.
`xval`	Double array of length `maxpoints` which will be filled with the `x`-values of the breakpoints. May be `NULL` if not required.
`yval`	Double array of length `maxpoints` which will be filled with the `y`-values of the breakpoints. May be `NULL` if not required.
`maxpoints`	Maximum number of breakpoints to be retrieved.
`p_npoints`	Pointer to return the number of breakpoints in the selected constraints. If the number of breakpoints is greater than `maxpoints`, then only `maxpoints` elements will be returned in the `xval` and `yval` arrays. May be `NULL` if not required.
`first`	First piecewise linear constraint in the range.
`last`	Last piecewise linear constraint in the range.

Example

The following example retrieves all variables and breakpoints in the first two piecewise linear constraints:

```
int *colind;
int *resultant;
int *start;
double *xval;
double *yval;
int maxpoints;
int npoints;
...
XPRSgetpwlcons(prob, NULL, NULL, NULL, NULL, NULL, 0, &maxpoints, 0, 1);
colind = (int*) malloc(2*sizeof(int));
resultant = (int*) malloc(2*sizeof(int));
start = (int*) malloc(3*sizeof(int));
xval = (double*) malloc(maxpoints*sizeof(double);
yval = (double*) malloc(maxpoints*sizeof(double);
```

```
XPRSgetpwlcons(prob, colind, resultant, start, xval, yval, maxpoints, &npoints, 0, 1);
...
```

Further information

1. It is possible to obtain just the number of breakpoints in the range of piecewise linear constraints by calling this function with `maxpoints` set to `0`, in which case the required size for the breakpoint arrays will be returned in `p_npoints`. In this case `colind` and `resultant` can also be queried or set to `NULL`.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddpwlcons`, `XPRSdelpwlcons`.

XPRSgetqobj

Purpose

Returns a single quadratic objective function coefficient corresponding to the variable pair (objqcol1, objqcol2) of the Hessian matrix.

Topic areas

Problem Information, Quadratic

Synopsis

```
int XPRS_CC XPRSgetqobj(XPRSprob prob, int objqcol1, int objqcol2, double
        *p_objqcoef);
```

Arguments

prob	The current problem.
objqcol1	Column index for the first variable in the quadratic term.
objqcol2	Column index for the second variable in the quadratic term.
p_objqcoef	Pointer to a double value where the objective function coefficient is to be placed.

Example

The following returns the coefficient of the x_0^2 term in the objective function, placing it in the variable value:

```
double value;
...
XPRSgetqobj(prob,0,0,&value);
```

Further information

p_objqcoef is the coefficient in the quadratic Hessian matrix. For example, if the objective function has the term $[3x_1x_2 + 3x_2x_1]/2$ the value retrieved by XPRSgetqobj is 3.0 and if the objective function has the term $[6x_1^2]/2$ the value retrieved by XPRSgetqobj is 6.0.

Related topics

XPRSgetmqobj, XPRSchgqobj, XPRSchgmqobj.

XPRSgetqrowcoeff

Purpose

Returns a single quadratic constraint coefficient corresponding to the variable pair (rowqcol1, rowqcol2) of the Hessian of a given constraint.

Topic areas

Problem Information, Quadratic

Synopsis

```
int XPRS_CC XPRSgetqrowcoeff (XPRSprob prob, int row, int rowqcol1, int
       rowqcol2, double *p_rowqcoef);
```

Arguments

prob	The current problem.
row	The quadratic row where the coefficient is to be looked up.
rowqcol1	Column index for the first variable in the quadratic term.
rowqcol2	Column index for the second variable in the quadratic term.
p_rowqcoef	Pointer to a double value where the objective function coefficient is to be placed.

Example

The following returns the coefficient of the x_0^2 term in the second row, placing it in the variable value :

```
double value;
...
XPRSgetqrowcoeff(prob,1,0,0,&value);
```

Further information

The coefficient returned corresponds to the Hessian of the constraint. That means the for constraint x + [x^2 + 6 xy] <= 10 XPRSgetqrowcoeff would return 1 as the coefficient of x^2 and 3 as the coefficient of xy.

Related topics

XPRSloadqcqp, XPRSaddqmatrix, XPRSchgqrowcoeff, XPRSgetqrowqmatrix, XPRSgetqrowqmatrixtriplets, XPRSgetqrows, XPRSchgqobj, XPRSchgmqobj, XPRSgetqobj.

XPRSgetqrowqmatrix

Purpose

Returns the nonzeros in a quadratic constraint coefficients matrix for the columns in a given range. To achieve maximum efficiency, `XPRSgetqrowqmatrix` returns the lower triangular part of this matrix only.

Topic areas

Problem Information, Quadratic

Synopsis

```
int XPRS_CC XPRSgetqrowqmatrix(XPRSprob prob, int row, int start[], int
       colind[], double rowqcoef[], int maxcoefs, int * p_ncoefs, int first,
       int last);
```

Arguments

`prob`	The current problem.
`row`	Index of the row for which the quadratic coefficients are to be returned.
`start`	Integer array which will be filled with indices indicating the starting offsets in the `colind` and `rowqcoef` arrays for each requested column. It must be length of at least `last-first+2`. Column `i` starts at position `start[i]` in the `colind` and `rowqcoef` arrays, and has `start[i+1]-start[i]` elements in it. May be `NULL` if maxcoefs is 0.
`colind`	Integer array of length maxcoefs which will be filled with the column indices of the nonzero elements in the lower triangular part of Q. May be `NULL` if maxcoefs is 0.
`rowqcoef`	Double array of length maxcoefs which will be filled with the nonzero element values. May be `NULL` if maxcoefs is 0.
`maxcoefs`	Number of elements to be saved in colind and rowqcoef. If `maxcoefs < *p_ncoefs`, only `maxcoefs` elements are written.
`p_ncoefs`	Pointer to the integer where the number of nonzero elements in the queried columns will be returned. If the number of nonzero elements is greater than maxcoefs, then only maxcoefs elements will be returned. If `p_ncoefs` is smaller than maxcoefs, then only `p_ncoefs` will be returned. May be `NULL`.
`first`	First column in the range.
`last`	Last column in the range.

Related topics

`XPRSloadqcqp`, `XPRSgetqrowcoeff`, `XPRSaddqmatrix`, `XPRSchgqrowcoeff`, `XPRSgetqrowqmatrixtriplets`, `XPRSgetqrows`, `XPRSchgqobj`, `XPRSchgmqobj`, `XPRSgetqobj`.

XPRSgetqrowqmatrixtriplets

Purpose

Returns the nonzeros in a quadratic constraint coefficients matrix as triplets (index pairs with coefficients). To achieve maximum efficiency, `XPRSgetqrowqmatrixtriplets` returns the lower triangular part of this matrix only.

Topic areas

Problem Information, Quadratic

Synopsis

```
int XPRS_CC XPRSgetqrowqmatrixtriplets(XPRSprob prob, int row, int *
        p_ncoefs, int rowqcol1[], int rowqcol2[], double rowqcoef[]);
```

Arguments

`prob`	The current problem.
`row`	Index of the row for which the quadratic coefficients are to be returned.
`p_ncoefs`	Argument used to return the number of quadratic coefficients in the row. May be `NULL` if not required.
`rowqcol1`	First index in the triplets. May be `NULL` if not required.
`rowqcol2`	Second index in the triplets. May be `NULL` if not required.
`rowqcoef`	Coefficients in the triplets. May be `NULL` if not required.

Further information

If a row index of `-1` is used, the function returns the quadratic coefficients for the objective function.

Related topics

`XPRSloadqcqp`, `XPRSgetqrowcoeff`, `XPRSaddqmatrix`, `XPRSchgqrowcoeff`, `XPRSgetqrowqmatrix`, `XPRSgetqrows`, `XPRSchgqobj`, `XPRSchgmqobj`, `XPRSgetqobj`.

XPRSgetqrows

Purpose

Returns the list indices of the rows that have quadratic coefficients.

Topic areas

Problem Information, Quadratic

Synopsis

```
int XPRS_CC XPRSgetqrows(XPRSprob prob, int * p_nrows, int rowind[]);
```

Arguments

`prob`	The current problem.
`p_nrows`	Used to return the number of quadratic constraints in the matrix.
`rowind`	Array of length `*p_nrows` used to return the indices of rows with quadratic coefficients in them. May be NULL if not required.

Related topics

`XPRSloadqcqp`, `XPRSgetqrowcoeff`, `XPRSaddqmatrix`, `XPRSchgqrowcoeff`, `XPRSgetqrowqmatrix`, `XPRSgetqrowqmatrixtriplets`, `XPRSchgqobj`, `XPRSchgmqobj`, `XPRSgetqobj`.

XPRSgetrhs

Purpose

Returns the right hand side elements for the rows in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetrhs(XPRSprob prob, double rhs[], int first, int last);
```

Arguments

`prob`	The current problem.
`rhs`	Double array of length `last-first+1` where the right hand side elements are to be placed.
`first`	First row in the range.
`last`	Last row in the range.

Example

The following example retrieves the right hand side values of the problem:

```
int rows;
double *rhs;
...
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
rhs = (double *) malloc(sizeof(double)*rows);
XPRSgetrhs(prob,rhs,0,rows-1);
```

See also example `repair.c`.

Further information

For range rows, `rhs` is populated with the upper limit of the row range.

Related topics

`XPRSchgrhs`, `XPRSchgrhsrange`, `XPRSgetrhsrange`.

XPRSgetrhsrange

Purpose

Returns the right hand side range values for the rows in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetrhsrange(XPRSprob prob, double rng[], int first, int
      last);
```

Arguments

`prob`	The current problem.
`rng`	Double array of length `last-first+1` where the right hand side range values are to be placed.
`first`	First row in the range.
`last`	Last row in the range.

Example

The following returns right hand side range values for all rows in the matrix:

```
int rows;
double *rng;
...
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
rng = malloc(rows*sizeof(double));
XPRSgetrhsrange(prob,rng,0,rows);
```

See also example `repair.c`.

Further information

Range values are defined as the upper limit of the row range minus the lower limit, which is always non-negative. For equality rows, the range is `0`. For inequality rows and nonbinding rows, the range is `XPRS_INFINITY`.

Related topics

`XPRSchgrhs`, `XPRSchgrhsrange`, `XPRSgetrhs`.

XPRSgetrowflags

Purpose

Retrieve if a range of rows have been set up as special rows.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetrowflags(XPRSprob prob, int flags[], int first, int
        last);
```

Arguments

`prob`	The current problem
`flags`	Int array of length `last-first+1` where type of information (see below) will be returned
`first`	First row index to be checked
`last`	Last row index to be checked

Further information

The `flags` array returns a bitvector for each row defining all the information that is currently attached to that row:

`XPRS_ROWFLAG_QUADRATIC`	The row has quadratic coefficients.
`XPRS_ROWFLAG_DELAYED`	The row is marked as a delayed row.
`XPRS_ROWFLAG_MODELCUT`	The row is marked as a model cut.
`XPRS_ROWFLAG_INDICATOR`	The row is used as an indicator.
`XPRS_ROWFLAG_NONLINEAR`	The row has nonlinear coefficients.

Example

The following example will print all three messages if the row at index 1 of the problem is an indicator constraint involving a quadratic matrix:

```
int flags[2];
...
XPRSgetrowflags(prob,1,2,flags);
if (flags[0] & (XPRS_ROWFLAG_QUADRATIC)) { printf("the second row is quadratic" }
if (flags[0] & (XPRS_ROWFLAG_INDICATOR)) { printf("the second row is an indicator" }
if (flags[0] & (XPRS_ROWFLAG_QUADRATIC+XPRS_ROWFLAG_INDICATOR)) { printf("the second row is a
```

Related topics

`XPRSclearrowflags`

XPRSgetrows, XPRSgetrows64

Purpose

Returns the nonzeros in the constraint matrix for the rows in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetrows(XPRSprob prob, int start[], int colind[], double
      colcoef[], int maxcoefs, int *p_ncoefs, int first, int last);

int XPRS_CC XPRSgetrows64(XPRSprob prob, XPRSint64 start[], int colind[],
      double colcoef[], XPRSint64 maxcoefs, XPRSint64 *p_ncoefs, int first,
      int last);
```

Arguments

prob	The current problem.
start	Integer array which will be filled with the indices indicating the starting offsets in the colind and colcoef arrays for each requested row. It must be of length at least last-first+2. Row i starts at position start[i] in the colind and colcoef arrays, and has start[i+1]-start[i] elements in it. May be NULL if not required.
colind	Integer arrays of length maxcoefs which will be filled with the column indices of the nonzero elements for each row. May be NULL if not required.
colcoef	Double array of length maxcoefs which will be filled with the nonzero element values. May be NULL if not required.
maxcoefs	Maximum number of elements to be retrieved.
p_ncoefs	Pointer to the integer where the number of nonzero elements in the selected rows will be returned. If the number of nonzero elements is greater than maxcoefs, then only maxcoefs elements will be returned. If p_ncoefs is smaller than maxcoefs, then only p_ncoefs will be returned. May be NULL if not required.
first	First row in the range.
last	Last row in the range.

Example

The following example returns and displays at most six nonzero matrix entries in the first two rows:

```
int maxcoefs=6, ncoefs, nreturnedels, start[3], colind[6];
double colcoef[6];
...
XPRSgetrows(prob,start,colind,colcoef,maxcoefs,&ncoefs,0,1);

nreturnedels = ncoefs > maxcoefs ? maxcoefs : ncoefs;
for(i=0;i<nreturnedels;i++) printf("\t%2.1f\n",dmtval[i]);
```

Further information

It is possible to obtain just the number of elements in the range of columns by replacing start, colind and colcoef by NULL. In this case, maxcoefs must be set to 0 to indicate that the length of arrays passed is 0.

Related topics

XPRSgetcols, XPRSgetrowtype.

XPRSgetrowtype

Purpose

Returns the row types for the rows in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetrowtype(XPRSprob prob, char rowtype[], int first, int
      last);
```

Arguments

`prob`	The current problem.
`rowtype`	Character array of length `last-first+1` characters where the row types will be returned: `N` indicates a free constraint; `L` indicates a $\leqslant$ constraint; `E` indicates an = constraint; `G` indicates a $\geqslant$ constraint; `R` indicates a range constraint.
`first`	First row in the range.
`last`	Last row in the range.

Example

The following example retrieves row types into an array `rowtype`:

```
int rows;
char *rowtype;
...
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
rowtype = (char *) malloc(sizeof(char)*rows);
XPRSgetrowtype(prob,rowtype,0,rows-1);
```

See also example `repair.c`.

Related topics

`XPRSchgrowtype`, `XPRSgetrows`.

XPRSgetscale

Purpose

Returns the the current scaling of the matrix.

Topic area

Numerics

Synopsis

```
int XPRS_CC XPRSgetscale(XPRSprob prob, int rowscale[], int colscale[]);
```

Arguments

prob	The current problem.
rowscale	Integer array of size ROWS that will contain the powers of 2 with which the rows are currently scaled.
colscale	Integer array of size COLS that will contain the powers of 2 with which the columns are currently scaled.

Related topics

XPRSscale (SCALE).

XPRSgetscaledinfeas

Purpose

Returns a list of scaled infeasible primal and dual variables for the original problem. If the problem is currently presolved, it is postsolved before the function returns.

Topic areas

LP, Numerics

Synopsis

```
int XPRS_CC XPRSgetscaledinfeas(XPRSprob prob, int *p_nprimalcols, int
       *p_nprimalrows, int *p_ndualrows, int *p_ndualcols, int x[], int
       slack[], int duals[], int djs[]);
```

Arguments

prob	The current problem.
p_nprimalcols	Number of primal infeasible variables.
p_nprimalrows	Number of primal infeasible rows.
p_ndualrows	Number of dual infeasible rows.
p_ndualcols	Number of dual infeasible variables.
x	Integer array of length p_nprimalcols where the primal infeasible variables will be returned. May be NULL if not required.
slack	Integer array of length p_nprimalrows where the primal infeasible rows will be returned. May be NULL if not required.
duals	Integer array of length p_ndualrows where the dual infeasible rows will be returned. May be NULL if not required.
djs	Integer array of length p_ndualcols where the dual infeasible variables will be returned. May be NULL if not required.

Error value

422	A solution is not available.

Related controls

Double

FEASTOL	Tolerance on RHS.
OPTIMALITYTOL	Reduced cost tolerance.

Example

In this example, XPRSgetscaledinfeas is first called with nulled integer arrays to get the number of infeasible entries. Then space is allocated for the arrays and the function is again called to fill them in.

```
int *x, *slack, *duals, *djs, nprimalcols, nprimalrows, ndualrows, ndualcols;
...
XPRSgetscaledinfeas(prob, &nprimalcols, &nprimalrows, &ndualrows, &ndualcols,
                    NULL, NULL, NULL, NULL);

x = malloc(nprimalcols * sizeof(int));
slack = malloc(nprimalrows * sizeof(int));
duals = malloc(ndualrows * sizeof(int));
djs = malloc(ndualcols * sizeof(int));
XPRSgetscaledinfeas(prob, &nprimalcols, &nprimalrows, &ndualrows, &ndualcols,
                    x, slack, duals, djs);
```

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. If any of the last four arguments are set to NULL, the corresponding number of infeasibilities is still returned.

Related topics

XPRSgetinfeas, XPRSgetiisdata, XPRSiisall, XPRSiisclear, XPRSiisfirst, XPRSiisisolations, XPRSiisnext, XPRSiisstatus, XPRSiiswrite, IIS.

XPRSgetstrattrib, XPRSgetstringattrib

Purpose

Enables users to recover the values of various string problem attributes. Problem attributes are set during loading and optimization of a problem.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetstrattrib(XPRSprob prob, int attrib, char *value);

int XPRS_CC XPRSgetstringattrib(XPRSprob prob, int attrib, char *value, int
      maxbytes, int* p_nbytes);
```

Arguments

`prob`	The current problem.
`attrib`	Problem attribute whose value is to be returned. A full list of all problem attributes may be found in 10, or from the list in the `xprs.h` header file.
`value`	Pointer to a string where the value of the attribute (plus null terminator) will be returned.
`maxbytes`	Maximum number of bytes to be written into the cgval argument.
`p_nbytes`	Returns the length of the string control including the null terminator.

Related topics

`XPRSgetdblattrib`, `XPRSgetintattrib`.

XPRSgetstrcontrol, XPRSgetstringcontrol

Purpose

Returns the value of a given string control parameters.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSgetstrcontrol(XPRSprob prob, int control, char *value);

int XPRS_CC XPRSgetstringcontrol(XPRSprob prob, int control, char *value,
       int maxbytes, int* p_nbytes);
```

Arguments

prob	The current problem.
control	Control parameter whose value is to be returned. A full list of all controls may be found in 9, or from the list in the xprs.h header file.
value	Pointer to a string where the value of the control (plus null terminator) will be returned.
maxbytes	Maximum number of bytes to be written into the value argument.
p_nbytes	Returns the length of the string control including the null terminator.

Related topics

XPRSgetdblcontrol, XPRSgetintcontrol, XPRSsetstrcontrol.

XPRSgetub

Purpose

Returns the upper bounds for the columns in a given range.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSgetub(XPRSprob prob, double ub[], int first, int last);
```

Arguments

`prob`	The current problem.
`ub`	Double array of length `last-first+1` where the upper bounds are to be placed.
`first`	First column in the range.
`last`	Last column in the range.

Example

The following example retrieves the upper bounds for the columns of the current problem:

```
int cols;
double *ub;
...
XPRSgetintattrib(prob, XPRS_COLS, &cols);
ub = (double *) malloc(sizeof(double)*cols);
XPRSgetub(prob, ub, 0, cols-1);
```

See also example `repair.c`.

Further information

Values greater than or equal to `XPRS_PLUSINFINITY` should be interpreted as infinite; values less than or equal to `XPRS_MINUSINFINITY` should be interpreted as infinite and negative.

Related topics

`XPRSchgbounds`, `XPRSgetlb`.

XPRSgetunbvec

Purpose

Returns the index vector which causes the primal simplex or dual simplex algorithm to determine that a matrix is primal or dual unbounded respectively.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetunbvec(XPRSprob prob, int *p_seq);
```

Arguments

prob	The current problem.
p_seq	Pointer to an integer where the vector causing the problem to be detected as being primal or dual unbounded will be returned. In the dual simplex case, the vector is the leaving row for which the dual simplex detected dual unboundedness. In the primal simplex case, the vector is the entering row p_seq (if p_seq is in the range 0 to ROWS-1) or column (variable) p_seq-ROWS-SPAREROWS (if p_seq is between ROWS+SPAREROWS and ROWS+SPAREROWS+COLS-1) for which the primal simplex detected primal unboundedness.

Error value

91	A current problem is not available.

Further information

When solving using the dual simplex method, if the problem is primal infeasible then XPRSgetunbvec returns the pivot row where dual unboundedness was detected. Also note that when solving using the dual simplex method, if the problem is primal unbounded then XPRSgetunbvec returns -1 since the problem is dual infeasible and not dual unbounded.

Related topics

XPRSgetinfeas, XPRSlpoptimize.

XPRSgetversion

Purpose

Returns the full Optimizer version number in the form 15.10.03, where 15 is the major release, 10 is the minor release, and 03 is the build number.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSgetversion(char *version);
```

Argument

`version`	Buffer long enough to hold the version string (plus a null terminator). This should be at least 16 characters.

Related controls

Integer

`VERSION`	The Optimizer version number

Example

The following calls `XPRSgetversion` to return version information at the start of the program:

```
char version[16];
XPRSgetversion(version);
printf("Xpress Optimizer version %s\n",version);
```

Further information

1. This function supersedes the `VERSION` control, which only returns the first two parts of the version number. Release 2004 versions of the Optimizer have a three-part version number.
2. This function can be called without calling `XPRSinit()` first.

Related topics

`XPRSinit`, `XPRSgetversionnumbers`.

XPRSgetversionnumbers

Purpose

Returns the Optimizer version numbers split into major, minor, and build number.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSgetversionnumbers(int *p_major, int *p_minor, int
       *p_build);
```

Arguments

`p_major`	Pointer to integer to receive the major version number. Can be `NULL`.
`p_minor`	Pointer to integer to receive the minor version number. This is a number from 0 to 99. Can be `NULL`.
`p_build`	Pointer to integer to receive the build number. This is a number from 0 to 99. Can be `NULL`.

Related controls

Integer

`VERSION`	The Optimizer version number

Example

The following calls `XPRSgetversionnumbers` to return version information at the start of the program:

```
int major, minor, build;
XPRSgetversionnumbers(&amp:major, &amp;minor, &amp;build);
printf("Xpress Optimizer version %d.%02d.%02d\n", major, minor, build);
```

Further information

1. This function always returns 0.
2. This function supersedes the `VERSION` control, which only returns the first two parts of the version number. Release 2004 versions of the Optimizer have a three-part version number.
3. This function can be called without calling `XPRSinit()` first.

Related topics

`XPRSinit`.

HELP

Purpose

Provides quick reference help for console users of the Optimizer.

Topic area

Misc

Synopsis

```
HELP
HELP commands
HELP controls
HELP attributes
HELP [command-name]
HELP [control-name]
HELP [attribute-name]
```

Example

This command is used by calling it at the Console Optimizer command line:

```
HELP MAXTIME
```

Related topics

None.

IIS

Purpose

Provides the Irreducible Infeasible Set (IIS) functionality for the console.

Topic area

Infeasibility

Synopsis

```
IIS [-flags]
```

Arguments

`IIS`	Finds an IIS.
`IIS -a`	Performs an automated search for a set of independent IISs.
`IIS -c`	Resets the search for IISs (deletes already found ones).
`IIS -e [num] fn`	Writes a CSV file named `fn` containing the IIS data of IIS `num`.
`IIS -f`	Generates only an approximation of an IIS. This functionality is available only for linear problems.
`IIS -i num`	Performs the isolation identification for IIS with ordinal number `num` (only for linear problems).
`IIS -n`	Attempts to find another (independent) IIS (only for linear problems).
`IIS -p [num]`	Prints the IIS with ordinal number `num` to the screen.
`IIS -s`	Returns statistics on the IISs found.
`IIS -w [num] fn type`	Writes an LP or MPS file named `fn` containing the IIS subproblem of IIS `num` depending on the `type` flags.

Example 1 (Console)

This example reads in an infeasible problem, executes an automated search for the IISs, prints the IIS to the screen and then displays a summary on the results.

```
READPROB PROB.LP
IIS -a -s
```

Example 2 (Console)

This example reads in an infeasible problem, identifies an IIS and its isolations, then writes the IIS as an LP for easier viewing and as a CSV file to contain the supplementary information.

```
READPROB PROB.LP
IIS
IIS -i -p 1
IIS -w 1 "IIS.LP"  lp
IIS -e 1 "IIS.CSV"
```

Further information

1. This function applies to all the problem types that the Xpress Solver can solve, including linear, convex and nonconvex quadratic, and general nonlinear problems, with or without discrete entities (integer variables, piecewise linear functions, special ordered sets, semicontinuous variables). User functions are not supported.
2. The IISs are numbered from `1` to `NUMIIS`. If no IIS number is provided, the functions take the last IIS identified as default. When applicable, IIS `0` refers to the initial infeasible IIS (the IIS approximation).
3. A model may have several infeasibilities. Repairing a single IIS may not make the model feasible. For this reason the Optimizer attempts to find an IIS for each of the infeasibilities in a model. You may call the `IIS -n` function repeatedly, or use the `IIS -a` function to retrieve all IISs at once.
4. An IIS isolation (for linear problems) is a special constraint or bound in an IIS. Removing an IIS isolation constraint or bound will remove all infeasibilities in the IIS without increasing the infeasibilities in any row or column outside the IIS, thus in any other IISs. The IIS isolations thus indicate the likely cause of each independent infeasibility and give an indication of which constraint or bound to drop or modify. It is not always possible to find IIS isolations. IIS isolations are only available for linear problems.
5. Generally, one should first look for rows or columns in the IIS which are both in isolation, and have a high dual multiplier relative to the others.
6. Initial infeasible subproblem: The subproblem identified after the sensitivity filter is referred to as initial infeasible subproblem. Its size is crucial to the running time of the deletion filter and it contains all the infeasibilities of the first phase simplex algorithm, thus if the corresponding rows and bounds are removed the problem becomes feasible.
7. `IIS -f` performs the initial sensitivity analysis on rows and columns to reduce the problem size, and sets up the initial infeasible subproblem. This subproblem speeds up the generation of IISs significantly, however in itself it may serve as an approximation of an IIS, since its identification typically takes only a fraction of time compared to the identification of an IIS.
8. The `num` parameter cannot be zero for `IIS -i`: the concept of isolations is meaningless for the initial infeasible subproblem.
9. If `IIS -n [num]` is called, the return status is `1` if less than `num` IISs have been found and zero otherwise. The total number of IISs found is stored in `NUMIIS`.
10. The type flags passed to `IIS -w` are directly passed to the `WRITEPROB` command.
11. The LP or MPS files created by `IIS -w` corresponding to an IIS contain no objective function, since infeasibility is independent from the objective.
12. Please note that there are problems on the boundary of being infeasible or not. For such problems, feasibility or infeasibility often depends on tolerances or even on scaling. This phenomenon makes it possible that after writing an IIS out as an LP file and reading it back, it may report feasibility. As a first check it is advised to consider the following options:
    (a) Turn presolve off (e.g. in console `presolve = 0`) since the nature of an IIS makes it necessary that during their identification the presolve is turned off.
    (b) Use the primal simplex method to solve the problem (e.g. in console `lpoptimize -p`).
    (c) Tighten the feasibility tolerances `FEASTOL` or `MIPTOL`.
13. Note that the sense of the original objective function plays no role in an IIS.
14. The supplementary information provided in the CSV file created by `IIS -e` is identical to that returned by the `XPRSgetiisdata` function.
15. In the case of hitting a time or resource limit, the IIS approximation and the IISs generated so far are always available. In such cases the last IIS may not be minimal.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisprint`, `XPRSiisstatus`, `XPRSiiswrite`.

XPRSiisall

Purpose

Performs an automated search for independent Irreducible Infeasible Sets (IIS) in an infeasible problem.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSiisall(XPRSprob prob);
```

Argument

prob	The current problem.

Related controls

Integer

MAXIIS	Number of Irreducible Infeasible Sets to be found.

Example

This example searches for IISs and then queries the problem attribute NUMIIS to determine how many were found:

```
int iis;
...
XPRSiisall(prob);
XPRSgetintattrib(prob, XPRS_NUMIIS, &iis);
printf("number of IISs = %d\n", iis);
```

Further information

1. This function applies to all problem types, including general nonlinear problems. For anything but linear problems its functionality is identical to XPRSiisfirst.
2. Calling IIS -a from the console has the same effect as this function.
3. A model may have several infeasibilities. Repairing a single IIS may not make the model feasible. For this reason the Optimizer can find an IIS for each of the infeasibilities in a model. If the control MAXIIS is set to a positive integer value then the XPRSiisall command will stop if MAXIIS IISs have been found. By default the control MAXIIS is set to -1, in which case an IIS is found for each of the infeasibilities in the model.
4. The problem attribute NUMIIS allows the user to recover the number of IISs found in a particular search. Alternatively, the XPRSiisstatus function may be used to retrieve the number of IISs found by XPRSiisfirst (IIS), XPRSiisnext (IIS -n) or XPRSiisall (IIS -a) functions.

Related topics

XPRSgetiisdata, XPRSiisclear, XPRSiisfirst, XPRSiisisolations, XPRSiisnext, XPRSiisprint, XPRSiisstatus, XPRSiiswrite, IIS.

XPRSiisclear

Purpose

Resets the search for Irreducible Infeasible Sets (IIS).

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSiisclear(XPRSprob prob);
```

Argument

`prob`	The current problem.

Example

```
XPRSiisclear(prob);
```

Further information

1. Calling `IIS -c` from the console has the same effect as this function.
2. The information stored internally about the IISs identified by `XPRSiisfirst`, `XPRSiisnext` or `XPRSiisall` are cleared. Functions `XPRSgetiisdata`, `XPRSiisstatus`, `XPRSiiswrite` and `XPRSiisisolations` cannot be called until the IIS identification procedure is started again.
3. This function is called automatically by `XPRSiisfirst` and `XPRSiisall`.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisprint`, `XPRSiisstatus`, `XPRSiiswrite`, `IIS`.

XPRSiisfirst

Purpose

Initiates a search for an Irreducible Infeasible Set (IIS) in an infeasible problem.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSiisfirst(XPRSprob prob, int mode, int *p_status);
```

Arguments

`prob`	The current problem.
`mode`	The IIS search mode: 0 stops after finding the initial infeasible subproblem; 1 find an IIS, emphasizing simplicity of the IIS; 2 find an IIS, emphasizing a quick result.
`p_status`	The status after the search: 0 success; 1 feasible problem; 2 error; 3 timeout or interruption.

Example

This call looks for the first IIS.

```
XPRSiisfirst(myprob,1,&status);
```

Further information

1. This function applies to all problem types.
2. Calling `IIS` from the console has the same effect as this function.
3. A model may have several infeasibilities. Repairing a single IIS may not make the model feasible. For this reason the Optimizer can find an IIS for each of the infeasibilities in a model. For the generation of several independent IISs use functions `XPRSiisnext` (`IIS -n`) or `XPRSiisall` (`IIS -a`).
4. IIS sensitivity filter: after an optimal but infeasible first phase primal simplex, it is possible to identify a subproblem containing all the infeasibilities (corresponding to the given basis) to reduce the size of the IIS working problem dramatically, i.e., rows with zero duals (thus with artificials of zero reduced cost) and columns that have zero reduced costs may be deleted. Moreover, for rows and columns with nonzero costs, the sign of the cost is used to relax equality rows either to less than or greater than equal rows, and to drop either possible upper or lower bounds on columns.
5. Initial infeasible subproblem: The subproblem identified after the sensitivity filter is referred to as initial infeasible subproblem. Its size is crucial to the running time of the deletion filter and it contains all the infeasibilities of the first phase simplex, thus if the corresponding rows and bounds are removed the problem becomes feasible.
6. `XPRSiisfirst` performs the initial sensitivity analysis on rows and columns to reduce the problem size, and sets up the initial infeasible subproblem. This subproblem significantly speeds up the generation of IISs, however in itself it may serve as an approximation of an IIS, since its identification typically takes only a fraction of time compared to the identification of an IIS.
7. The IIS approximation and the IISs generated so far are always available.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisprint`, `XPRSiisstatus`, `XPRSiiswrite`, `IIS`.

XPRSiisisolations

Purpose

Performs the isolation identification procedure for an Irreducible Infeasible Set (IIS). This function applies only to linear problems.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSiisisolations(XPRSprob prob, int iis);
```

Arguments

`prob`	The current problem.
`iis`	The number of the IIS identified by either `XPRSiisfirst` (`IIS`), `XPRSiisnext` (`IIS -n`) or `XPRSiisall` (`IIS -a`) in which the isolations should be identified.

Example

This example finds the first IIS and searches for the isolations in that IIS.

```
XPRSiisfirst(prob,1,&status);
XPRSiisisolations (prob,1);
```

Further information

1. Calling `IIS -i [iis]` from the console has the same effect as this function.
2. An IIS isolation is a special constraint or bound in an IIS. Removing an IIS isolation constraint or bound will remove all infeasibilities in the IIS without increasing the infeasibilities in any row or column outside the IIS, thus in any other IISs. The IIS isolations thus indicate the likely cause of each independent infeasibility and give an indication of which constraint or bound to drop or modify. It is not always possible to find IIS isolations.
3. Generally, one should first look for rows or columns in the IIS which are both in isolation, and have a high dual multiplier relative to the others.
4. The `iis` parameter cannot be zero: the concept of isolations is meaningless for the initial infeasible subproblem.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisnext`, `XPRSiisprint`, `XPRSiisstatus`, `XPRSiiswrite`, `IIS`.

XPRSiisnext

Purpose

Continues the search for further Irreducible Infeasible Sets (IIS), or calls `XPRSiisfirst` (`IIS`) if no IIS has been identified yet.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSiisnext(XPRSprob prob, int *p_status);
```

Arguments

`prob`	The current problem.	
`p_status`	The status after the search:	
	`0`	success;
	`1`	no more IIS could be found, or problem is feasible if no `XPRSiisfirst` call preceded;
	`2`	on error (when the function returns nonzero).

Example

This looks for a further IIS.

```
XPRSiisnext(prob,&status);
```

Further information

1. Returning multiple IISs is supported only for linear problems.
2. Calling `IIS -n` from the console has the same effect as this function.
3. A model may have several infeasibilities. Repairing a single IIS may not make the model feasible. For this reason the Optimizer attempts to find an IIS for each of the infeasibilities in a model. You may call the `XPRSiisnext` function repeatedly, or use the `XPRSiisall` (`IIS -a`) function to retrieve all IIS at once.
4. This function is not affected by the control `MAXIIS`.
5. If the problem has been modified since the last call to `XPRSiisfirst` or `XPRSiisnext`, the generation process has to be started from scratch.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisprint`, `XPRSiisstatus`, `XPRSiiswrite`, `IIS`.

XPRSiisprint

Purpose

Prints a given Irreducible Infeasible Set (IIS) in the log. If 0 is passed as the IIS number parameter, the initial infeasible subproblem is printed.

Topic areas

Infeasibility, Logging

Synopsis

```
int XPRS_CC XPRSiisprint(XPRSprob prob, int iis);
```

Arguments

`prob`	The current problem.
`iis`	The ordinal number of the IIS to be printed.

Further information

Calling `IIS -p [iis]` from the console has the same effect as this function.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisstatus`, `IIS`.

XPRSiisstatus

Purpose

Returns statistics on the Irreducible Infeasible Sets (IIS) found so far by `XPRSiisfirst` (`IIS`), `XPRSiisnext` (`IIS -n`) or `XPRSiisall` (`IIS -a`).

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSiisstatus(XPRSprob prob, int *p_niis, int nrows[], int
       ncols[], double suminfeas[], int numinfeas[]);
```

Arguments

`prob`	The current problem.
`p_niis`	The number of IISs found so far.
`nrows`	Number of rows in the IISs.
`ncols`	Number of bounds in the IISs.
`suminfeas`	The sum of infeasibilities in the IISs after the first phase simplex.
`numinfeas`	The number of infeasible variables in the IISs after the first phase simplex.

Example

This example first retrieves the number of IISs found so far, and then retrieves their main properties. Note that the arrays have size `count+1`, since the first index is reserved for the initial infeasible subset.

```
XPRSiisstatus(myprob,&count,NULL,NULL,NULL,NULL);
nrows   = malloc((count+1)*sizeof(int));
ncols   = malloc((count+1)*sizeof(int));
suminfeas = malloc((count+1)*sizeof(double));
numinfeas = malloc((count+1)*sizeof(int));
XPRSiisstatus(myprob,&count,nrows,ncols,suminfeas,numinfeas);
```

Further information

1. Calling `IIS -s` from the console has the same effect as this function.
2. All arrays should be of dimension `p_niis+1`. The arrays are `0` based, index 0 corresponding to the initial infeasible subproblem.
3. The arrays may be NULL if not required.
4. For the initial infeasible problem (IIS 0) the subproblem size is returned (which may be different from the number of bounds), while for the IISs the number of bounds is returned (usually much smaller than the number of columns in the IIS).
5. Note that the values in `suminfeas` and `numinfeas` depend on the IIS, and in turn on the basis where the simplex algorithm has stopped. Different IISs can have different error measures.
6. `p_niis` is set to `-1` if the search for IISs has not yet started.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisprint`, `XPRSiiswrite`, `IIS`.

XPRSiiswrite

Purpose

Writes an LP/MPS/CSV file containing a given Irreducible Infeasible Set (IIS). If 0 is passed as the IIS number parameter, the initial infeasible subproblem is written.

Topic areas

Infeasibility, File IO

Synopsis

```
int XPRS_CC XPRSiiswrite(XPRSprob prob, int iis, const char *filename, int
      filetype, const char *flags);
```

Arguments

`prob`	The current problem.
`iis`	The ordinal number of the IIS to be written.
`filename`	The name of the file to be created.
`filetype`	Type of file to be created: `0` creates an lp/mps file containing the IIS as a linear programming problem; `1` creates a comma separated (csv) file containing the description and supplementary information on the given IIS.
`flags`	Flags passed to the `XPRSwriteprob` function.

Example

This writes the first IIS (if one exists and is already found) as an lp file.

```
XPRSiiswrite(prob,1,"iis.lp",0,"l")
```

Further information

1. Calling `IIS -w [iis] filename` and `IIS -e [iis] filename` from the console have the same effect as this function.
2. Please note that there are problems on the boundary of being infeasible or not. For such problems, feasibility or infeasibility often depends on tolerances or even on scaling. This phenomenon makes it possible that after writing an IIS out as an LP file and reading it back, it may report feasibility. As a first check it is advised to consider the following options:
   1. turn presolve off (e.g. in console `presolve = 0`) since the nature of an IIS makes it necessary that during their identification the presolve is turned off.
   2. use the primal simplex method to solve the problem (e.g. in console `LPOPTIMIZE -p`).
   3. Tighten the feasibility tolerances `FEASTOL` or `MIPTOL`.
3. Note that the original sense of the original objective function plays no role in an IIS.
4. Even though an attempt is made to identify the most infeasible IISs first by the `XPRSiisfirst` (`IIS`), `XPRSiisnext` (`IIS -n`) and `XPRSiisall` (`IIS -a`) functions, it is also possible that an IIS becomes just infeasible in problems that are otherwise highly infeasible. In such cases, you may try to deal with the more stable IISs first, and consider to use the infeasibility breaker tool if only slight infeasibilities remain.
5. The LP or MPS files created by `XPRSiiswrite` corresponding to an IIS contain no objective function, since infeasibility is independent of the objective.

Related topics

`XPRSgetiisdata`, `XPRSiisall`, `XPRSiisclear`, `XPRSiisfirst`, `XPRSiisisolations`, `XPRSiisnext`, `XPRSiisprint`, `XPRSiisstatus`, `IIS`.

XPRSinit

Purpose

Initializes the Optimizer library. This must be called before any other library routines.

Topic area

Licensing

Synopsis

```
int XPRS_CC XPRSinit(const char *path);
```

Argument

`path`	The directory where the FICO Xpress license file is located. Users should employ a value of `NULL` unless otherwise advised, allowing the standard initialization directories to be checked.

Return value

0	Success
1	An error occurred

Example

The following is the usual way of calling `XPRSinit`:

```
if(XPRSinit(NULL)) printf("Problem with XPRSinit\n");
```

See also examples `addmipsol.c`, `ComplexUserFunctions.c`, `els_managedcuts.c`, `els_usercuts.c`, `fixbv.c`, `globjpar.c`, `glrhspar.c`, `goalprog.c`, `knapsack.c`, `loadlp.c`, `mipsolenum.c`, `mipsolpool.c`, `mostviolated.c`, `Polygon_initialvalue.c`, `Polygon_textformula.c`, `Polygon_tokens.c`, `Polygon_userfunc.c`, `Polygon_userfunc_map.c`, `Polygon_userfunc_mapdelta.c`, `Polygon_userfunc_multimap.c`, `Polygon_userfunc_multimapdelta.c`, `Polygon_userfunc_vecmap.c`, `Polygon_userfunc_vecmapdelta.c`, `repair.c`, `roundint.c`, `savesol.c`, `tableau.c`, `trimloss.c`, `tsp.c` .

Further information

1. Whilst error checking should always be used on all library function calls, it is especially important to do so with the initialization functions, since a majority of errors encountered by users are caused at the initialization stage. Any nonzero return code indicates that no license could be found. In such circumstances the application should be made to exit. It is possible to retrieve a message describing the error by calling `XPRSgetlicerrmsg`.
2. In multi-threaded applications where all threads are equal, `XPRSinit` may be called by each thread prior to using the library. Whilst the process of initialization will be carried out only once, this guarantees that the library functions will be available to each thread as necessary. In applications with a clear master thread, spawning other Optimizer threads, initialization need only be called by the master thread.

Related topics

`XPRScreateprob`, `XPRSfree`, `XPRSgetlicerrmsg`.

XPRSinterrupt

Purpose

Interrupts the Optimizer algorithms.

Topic area

Solution Process

Synopsis

```
int XPRS_CC XPRSinterrupt(XPRSprob prob, int reason);
```

Arguments

`prob`	The current problem.
`reason`	The reason for stopping. Possible reasons are: `XPRS_STOP_NONE` do not stop; `XPRS_STOP_TIMELIMIT` time limit hit; `XPRS_STOP_WORKLIMIT` work limit hit; `XPRS_STOP_CTRLC` control C hit; `XPRS_STOP_NODELIMIT` node limit hit; `XPRS_STOP_ITERLIMIT` iteration limit hit; `XPRS_STOP_MIPGAP` MIP gap is sufficiently small; `XPRS_STOP_SOLLIMIT` solution limit hit; `XPRS_STOP_USER` user interrupt; `>= 1000` user defined value.

Further information

1. The `XPRSinterrupt` command can be called from any callback.
2. The range of values below 1000 is reserved for future extension. The behavior of the function is undefined if `reason` is smaller than 1000 and not one of the values listed above.

Related topics

None.

Example

See also examples `els_managedcuts.c`, `els_usercuts.c`, `knapsack.c`, `mostviolated.c`, `roundint.c`, `savesol.c`, `tsp.c` .

XPRSloadbasis

Purpose

Loads a basis from the user's areas.

Topic areas

Data Input, LP

Synopsis

```
int XPRS_CC XPRSloadbasis(XPRSprob prob, const int rowstat[], const int
        colstat[]);
```

Arguments

`prob`	The current problem.
`rowstat`	Integer array of length `ROWS` containing the basis status of the slack, surplus or artificial variable associated with each row. The status must be one of: `XPRS_NONBASIC_LOWER (0)` slack, surplus or artificial is non-basic at lower bound; `XPRS_BASIC (1)` slack, surplus or artificial is basic; `XPRS_NONBASIC_UPPER (2)` slack or surplus is non-basic at upper bound. `XPRS_SUPERBASIC (3)` slack or surplus is super-basic.
`colstat`	Integer array of length `COLS` containing the basis status of each of the columns in the constraint matrix. The status must be one of: `XPRS_NONBASIC_LOWER (0)` variable is non-basic at lower bound or superbasic at zero if the variable has no lower bound; `XPRS_BASIC (1)` variable is basic; `XPRS_NONBASIC_UPPER (2)` variable is at upper bound; `XPRS_SUPERBASIC (3)` variable is super-basic.

Example

This example loads a problem and then reloads a (previously optimized) basis from a similar problem to speed up the optimization:

```
XPRSreadprob(prob,"problem","");
XPRSloadbasis(prob,rowstat,colstat);
XPRSlpoptimize(prob,"");
```

See also examples `globjpar.c`, `glrhspar.c` .

Further information

If the problem has been altered since saving an advanced basis, you may want to alter the basis as follows before loading it:

- Make new variables non-basic at their lower bound (`colstat[icol]=0`), unless a variable has an infinite lower bound and a finite upper bound, in which case make the variable non-basic at its upper bound (`colstat[icol]=2`);
- Make new constraints basic (`rowstat[jrow]=1`);
- Try not to delete basic variables, or non-basic constraints.

Related topics

`XPRSgetbasis`, `XPRSgetpresolvebasis`, `XPRSloadpresolvebasis`.

XPRSloadbranchdirs

Purpose

Loads directives into the current problem to specify which MIP entities the Optimizer should continue to branch on when a node solution is integer feasible.

Topic areas

Data Input, Branching

Synopsis

```
int XPRS_CC XPRSloadbranchdirs(XPRSprob prob, int ncols, const int
       colind[], const int dir[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of directives.
`colind`	Integer array of length `ncols` containing the column numbers. A negative value indicates a set number (the first set being -1, the second -2, and so on).
`dir`	Integer array of length `ncols` containing either 0 or 1 for the entities given in `colind`. Entities for which dir is set to 1 will be branched on until fixed before an integer feasible solution is returned. If `dir` is NULL, the branching directive will be set for all entities in `colind`.

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSloaddirs`, `XPRSreaddirs`, A.5.

XPRSloadcuts

Purpose

Loads cuts from the cut pool into the matrix. Without calling `XPRSloadcuts` the cuts will remain in the cut pool but will not be active at the node. Cuts loaded at a node remain active at all descendant nodes unless they are deleted using `XPRSdelcuts`.

Topic areas

Data Input, Cuts

Synopsis

```
int XPRS_CC XPRSloadcuts(XPRSprob prob, int cuttype, int interp, int ncuts,
        const XPRScut cutind[]);
```

Arguments

`prob`	The current problem.
`cuttype`	Cut type.
`interp`	The way in which the cut type is interpreted: `-1` load all cuts; `1` treat cut types as numbers; `2` treat cut types as bit-vectors (compare Section 9.2) - load cut if any bit matches any bit set in `cuttype`; `3` treat cut types as bit-vectors (compare Section 9.2) - `0` load cut if all bits match those set in `cuttype`.
`ncuts`	Number of cuts to load.
`cutind`	Array of length `ncuts` containing pointers to the cuts to be loaded into the matrix. These are pointers returned by either `XPRSstorecuts` or `XPRSgetcpcutlist`.

Further information

This function should be called only from within callback functions set by either `XPRSaddcboptnode`, `XPRSaddcbprenode`, `XPRSaddcbnodelpsolved` or `XPRSaddcbpreintsol`.

Related topics

`XPRSaddcuts`, `XPRSdelcpcuts`, `XPRSdelcuts`, `XPRSgetcpcutlist`, 5.9.

XPRSloaddelayedrows

Purpose

Specifies that a set of rows in the matrix will be treated as delayed rows during a tree search. These are rows that must be satisfied for any integer solution, but will not be loaded into the active set of constraints until required.

Topic areas

Problem Input, Cuts

Synopsis

```
int XPRS_CC XPRSloaddelayedrows(XPRSprob prob, int nrows, const int
       rowind[]);
```

Arguments

`prob`	The current problem.
`nrows`	The number of delayed rows.
`rowind`	An array of row indices to treat as delayed rows.

Example

This sets the first six matrix rows as delayed rows in the MIP problem `prob`.

```
int rowind[] = {0,1,2,3,4,5}
...
XPRSloaddelayedrows(prob,6,rowind);
XPRSmipoptimize(prob,"");
```

See also example `tsp.c`.

Further information

1. Delayed rows must be set up before solving the problem. Any delayed rows will be removed from the matrix after presolve and added to a special pool. A delayed row will be added back into the active matrix only when such a row is violated by an integer solution found by the Optimizer.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSloadmodelcuts`.

XPRSloaddirs

Purpose

Loads directives into the matrix.

Topic area

Data Input

Synopsis

```
int XPRS_CC XPRSloaddirs(XPRSprob prob, int ndirs, const int colind[],
       const int priority[], const char dir[], const double uppseudo[],
       const double downpseudo[]);
```

Arguments

`prob`	The current problem.
`ndirs`	Number of directives.
`colind`	Integer array of length `ndirs` containing the column numbers. A negative value indicates a set number (the first set being `-1`, the second `-2`, and so on).
`priority`	Integer array of length `ndirs` containing the priorities for the columns or sets. Priorities must be between 0 and 1000, where columns/sets with smallest priority will be branched on first. May be `NULL` if not required.
`dir`	Character array of length `ndirs` specifying the branching direction for each column or set: `U` the entity is to be forced up; `D` the entity is to be forced down; `N` not specified. May be `NULL` if not required.
`uppseudo`	Double array of length `ndirs` containing the up pseudo costs for the columns or sets. May be `NULL` if not required.
`downpseudo`	Double array of length `ndirs` containing the down pseudo costs for the columns or sets. May be `NULL` if not required.

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSgetdirs`, `XPRSloadpresolvedirs`, `XPRSreaddirs`.

Example

See also example `trimloss.c`.

XPRSloadglobal, XPRSloadglobal64

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSloadmip` instead.

Topic areas

Problem Input, MIP Entities

Synopsis

```
int XPRS_CC XPRSloadglobal(XPRSprob prob, const char *probname, int ncols,
      int nrows, const char rowtype[], const double rhs[], const double
      rng[], const double objcoef[], const int start[], const int collen[],
      const int rowind[], const double rowcoef[], const double lb[], const
      double ub[], int nentities, int nsets, const char coltype[], const
      int entind[], const double limit[], const char settype[], const int
      setstart[], const int setind[], const double refval[]);

int XPRS_CC XPRSloadglobal64(XPRSprob prob, const char *probname, int
      ncols, int nrows, const char rowtype[], const double rhs[], const
      double rng[], const double objcoef[], const XPRSint64 start[], const
      int collen[], const int rowind[], const double rowcoef[], const
      double lb[], const double ub[], int nentities, int nsets, const char
      coltype[], const int entind[], const double limit[], const char
      settype[], const XPRSint64 setstart[], const int setind[], const
      double refval[]);
```

XPRSloadmip, XPRSloadmip64

Purpose

Used to load a MIP problem into the Optimizer data structures. Integer, binary, partial integer, semi-continuous and semi-continuous integer variables can be defined, together with sets of type 1 and 2. The reference row values for the set members are passed as an array rather than specifying a reference row.

Topic areas

Problem Input, MIP Entities

Synopsis

```
int XPRS_CC XPRSloadmip(XPRSprob prob, const char *probname, int ncols, int
      nrows, const char rowtype[], const double rhs[], const double rng[],
      const double objcoef[], const int start[], const int collen[], const
      int rowind[], const double rowcoef[], const double lb[], const double
      ub[], int nentities, int nsets, const char coltype[], const int
      entind[], const double limit[], const char settype[], const int
      setstart[], const int setind[], const double refval[]);

int XPRS_CC XPRSloadmip64(XPRSprob prob, const char *probname, int ncols,
      int nrows, const char rowtype[], const double rhs[], const double
      rng[], const double objcoef[], const XPRSint64 start[], const int
      collen[], const int rowind[], const double rowcoef[], const double
      lb[], const double ub[], int nentities, int nsets, const char
      coltype[], const int entind[], const double limit[], const char
      settype[], const XPRSint64 setstart[], const int setind[], const
      double refval[]);
```

Arguments

`prob`	The current problem.
`probname`	A string of up to `MAXPROBNAMELENGTH` characters containing a name for the problem.
`ncols`	Number of structural columns in the matrix.
`nrows`	Number of rows in the matrix not (including the objective row). Objective coefficients must be supplied in the `objcoef` array, and the objective function should not be included in any of the other arrays.
`rowtype`	Character array of length `nrows` containing the row types: `L` indicates a ⩽ constraint; `E` indicates an = constraint; `G` indicates a ⩾ constraint; `R` indicates a range constraint; `N` indicates a nonbinding constraint.
`rhs`	Double array of length `nrows` containing the right hand side coefficients. The right hand side value for a range row gives the **upper** bound on the row.
`rng`	Double array of length `nrows` containing the range values for range rows. Values for all other rows will be ignored. May be `NULL` if not required. The lower bound on a range row is the right hand side value minus the range value. The sign of the range value is ignored - the absolute value is used in all cases.
`objcoef`	Double array of length `ncols` containing the objective function coefficients. This can be `NULL` to set all objective coefficients to 0 (zero).
`start`	Integer array containing the offsets in the `rowind` and `rowcoef` arrays of the start of the elements for each column. This array is of length `ncols` or, if `collen` is `NULL`, length `ncols+1`. If `collen` is `NULL`, the extra entry of `start`, `start[ncols]`, contains the

	position in the `rowind` and `rowcoef` arrays at which an extra column would start, if it were present. In C, this value is also the length of the `rowind` and `rowcoef` arrays.
`collen`	Integer array of length `ncols` containing the number of nonzero elements in each column. May be `NULL` if not required. This array is not required if the non-zero coefficients in the `rowind` and `rowcoef` arrays are continuous, and the `start` array has `ncols+1` entries as described above. It may be `NULL` if not required.
`rowind`	Integer arrays containing the row indices for the nonzero elements in each column. If the indices are input contiguously, with the columns in ascending order, then the length of `rowind` is `start[ncols-1]+collen[ncols-1]` or, if `collen` is `NULL`, `start[ncols]`.
`rowcoef`	Double array containing the nonzero element values length as for `rowind`.
`lb`	Double array of length `ncols` containing the lower bounds on the columns. Use `XPRS_MINUSINFINITY` to represent a lower bound of minus infinity. If this is `NULL` then all lower bounds are 0 (zero).
`ub`	Double array of length `ncols` containing the upper bounds on the columns. Use `XPRS_PLUSINFINITY` to represent an upper bound of plus infinity. It this is `NULL` then all upper bounds are infinite.
`nentities`	Number of binary, integer, semi-continuous, semi-continuous integer and partial integer entities.
`nsets`	Number of SOS1 and SOS2 sets.
`coltype`	Character array of length `nentities` containing the entity types: `B` binary variables; `I` integer variables; `P` partial integer variables; `S` semi-continuous variables; `R` semi-continuous integer variables.
`entind`	Integer array of length `nentities` containing the column indices of the MIP entities.
`limit`	Double array of length `nentities` containing the integer limits for the partial integer variables and lower bounds for semi-continuous and semi-continuous integer variables (any entries in the positions corresponding to binary and integer variables will be ignored). May be `NULL` if not required.
`settype`	Character array of length `nsets` containing the set types: `1` SOS1 type sets; `2` SOS2 type sets. May be `NULL` if not required.
`setstart`	Integer array containing the offsets in the `setind` and `refval` arrays indicating the start of the sets. This array is of length `nsets+1`, the last member containing the offset where set `nsets+1` would start. May be `NULL` if not required.
`setind`	Integer array of length `setstart[nsets]-1` containing the columns in each set. May be `NULL` if not required.
`refval`	Double array of length `setstart[nsets]-1` containing the reference row entries for each member of the sets. These define the order for SOS2 constraints and may be used in branching for both types. May be `NULL` if not required.

Related controls

Integer

`EXTRACOLS`	Number of extra columns to be allowed for.
`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAMIPENTS`	Number of extra MIP entities to be allowed for.
`EXTRAROWS`	Number of extra rows to be allowed for.
`KEEPNROWS`	Status for nonbinding rows.

SCALING	Type of scaling.
Double	
MATRIXTOL	Tolerance on matrix elements.
SOSREFTOL	Minimum gap between reference row entries.

Example

The following specifies an integer problem, `MIPexample`, corresponding to:

minimize:	x + 2y		
subject to:	3x + 2y	$\leqslant$	400
	x + 3y	$\leqslant$	200

with both *x* and *y* integral:

```
char probname[]   = "MIPexample";
int ncols = 2, nrows = 2;
char rowtype[]     = {'L','L'};
double rhs[]      = {400.0, 200.0};
int start[]       = {0, 2, 4};
int rowind[]      = {0, 1, 0, 1};
double rowcoef[]  = {3.0, 1.0, 2.0, 3.0};
double objcoefs[] = {1.0, 2.0};
double lb[]       = {0.0, 0.0};
double ub[]       = {200.0, 200.0};

int nentities = 2;
int nsets = 0;
char coltype[] = {"I","I"};
int entind[] = {0,1};
...
XPRSloadmip(prob, probname, ncols, nrows, rowtype, rhs, NULL,
               objcoefs, start, NULL, rowind,
               rowcoef, lb, ub, nentities, nsets, coltype, entind,
               NULL, NULL, NULL, NULL, NULL);
```

Further information

1. The row and column indices follow the usual C convention of going from `0` to `nrows-1` and `0` to `ncols-1` respectively.
2. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` are defined in the Optimizer library header file.
3. Semi-continuous lower bounds are taken from the `limit` array. If this is `NULL` then they are given a default value of `1.0`. If a semi-continuous variable has a positive lower bound then this will be used as the semi-continuous lower bound and the lower bound on the variable will be set to zero.

Related topics

`XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSreadprob`.

XPRSloadlp, XPRSloadlp64

Purpose

Enables the user to pass a matrix directly to the Optimizer, rather than reading the matrix from a file.

Topic areas

Problem Input, LP

Synopsis

```
int XPRS_CC XPRSloadlp(XPRSprob prob, const char *probname, int ncols, int
       nrows, const char rowtype[], const double rhs[], const double rng[],
       const double objcoef[], const int start[], const int collen[], const
       int rowind[], const double rowcoef[], const double lb[], const double
       ub[]);

int XPRS_CC XPRSloadlp64(XPRSprob prob, const char *probname, int ncols,
       int nrows, const char rowtype[], const double rhs[], const double
       rng[], const double objcoef[], const XPRSint64 start[], const int
       collen[], const int rowind[], const double rowcoef[], const double
       lb[], const double ub[]);
```

Arguments

`prob`	The current problem.
`probname`	A string of up to `MAXPROBNAMELENGTH` characters containing a names for the problem.
`ncols`	Number of structural columns in the matrix.
`nrows`	Number of rows in the matrix (not including the objective). Objective coefficients must be supplied in the `objcoef` array, and the objective function should not be included in any of the other arrays.
`rowtype`	Character array of length `nrows` containing the row types: `L` indicates a ⩽ constraint; `E` indicates an = constraint; `G` indicates a ⩾ constraint; `R` indicates a range constraint; `N` indicates a nonbinding constraint.
`rhs`	Double array of length `nrows` containing the right hand side coefficients of the rows. The right hand side value for a range row gives the **upper** bound on the row.
`rng`	Double array of length `nrows` containing the range values for range rows. Values for all other rows will be ignored. May be `NULL` if not required. The lower bound on a range row is the right hand side value minus the range value. The sign of the range value is ignored - the absolute value is used in all cases.
`objcoef`	Double array of length `ncols` containing the objective function coefficients. This can be `NULL` to set all objective coefficients to 0 (zero).
`start`	Integer array containing the offsets in the `rowind` and `rowcoef` arrays of the start of the elements for each column. This array is of length `ncols` or, if `collen` is `NULL`, length `ncols+1`. If `collen` is `NULL`, the extra entry of `start`, `start[ncols]`, contains the position in the `rowind` and `rowcoef` arrays at which an extra column would start, if it were present. In C, this value is also the length of the `rowind` and `rowcoef` arrays.
`collen`	Integer array of length `ncols` containing the number of nonzero elements in each column. May be `NULL` if not required. This array is not required if the non-zero coefficients in the `rowind` and `rowcoef` arrays are continuous, and the `start` array has `ncols+1` entries as described above.
`rowind`	Integer array containing the row indices for the nonzero elements in each column. If the indices are input contiguously, with the columns in ascending order, the length of the

	`rowind` is `start[ncols-1]+collen[ncols-1]` or, if `collen` is `NULL`, `start[ncols]`.
`rowcoef`	Double array containing the nonzero element values; length as for `rowind`.
`lb`	Double array of length `ncols` containing the lower bounds on the columns. Use `XPRS_MINUSINFINITY` to represent a lower bound of minus infinity. If this is `NULL` then all lower bounds are 0 (zero).
`ub`	Double array of length `ncols` containing the upper bounds on the columns. Use `XPRS_PLUSINFINITY` to represent an upper bound of plus infinity. It this is `NULL` then all upper bounds are infinite.

Related controls

Integer

`EXTRACOLS`	Number of extra columns to be allowed for.
`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAROWS`	Number of extra rows to be allowed for.
`KEEPNROWS`	Status for nonbinding rows.
`SCALING`	Type of scaling.

Double

`MATRIXTOL`	Tolerance on matrix elements.

Example

Given an LP problem:

minimize:	x + y		
subject to:	2x	$\geqslant$	3
	x + 2y	$\geqslant$	3
	x + y	$\geqslant$	1

the following shows how this may be loaded into the Optimizer using `XPRSloadlp`:

```
char probname[] = "small";
int ncols = 2, nrows = 3;
char rowtype[]    = {'G','G','G'};
double rhs[]     = { 3 , 3 , 1 };
double objcoef[]     = { 1 , 1 };
int start[]     = { 0 , 3 , 5 };
int rowind[]     = { 0 , 1 , 2 , 1 , 2 };
double rowcoef[] = { 2 , 1 , 1 , 2 , 1 };
double lb[]    = { 0 , 0 };
double ub[]    = {XPRS_PLUSINFINITY,XPRS_PLUSINFINITY};

XPRSloadlp(prob, probname, ncols, nrows, rowtype, rhs, NULL,
           objcoef, start, NULL, rowind, rowcoef, lb, ub)
```

See also examples `ComplexUserFunctions.c`, `els_managedcuts.c`, `els_usercuts.c`, `loadlp.c`, `Polygon_initialvalue.c`, `Polygon_textformula.c`, `Polygon_tokens.c`, `Polygon_userfunc.c`, `Polygon_userfunc_map.c`, `Polygon_userfunc_mapdelta.c`, `Polygon_userfunc_multimap.c`, `Polygon_userfunc_multimapdelta.c`, `Polygon_userfunc_vecmap.c`, `Polygon_userfunc_vecmapdelta.c` .

Further information

1. The row and column indices follow the usual C convention of going from `0` to `nrows-1` and `0` to `ncols-1` respectively.
2. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` are defined in the Optimizer library header file.
3. For a range constraint, the value in the `rhs` array specifies the upper bound on the constraint, while the value in the `rng` array specifies the range on the constraint. So a range constraint *j* is interpreted as:

$$rhs_j - |rng_j| \leqslant \sum_i a_{ij} x_i \leqslant rhs_j$$

Related topics

`XPRSloadmip`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSreadprob`.

XPRSloadlpsol

Purpose

Loads an LP solution for the problem into the Optimizer.

Topic areas

Solution, LP, Data Input

Synopsis

```
int XPRS_CC XPRSloadlpsol(XPRSprob prob, const double x[], const double
        slack[], const double duals[], const double djs[], int *p_status);
```

Arguments

`prob`	The current problem.
`x`	Optional: Double array of length `COLS` (for the original problem and not the presolve problem) containing the values of the variables.
`slack`	Optional: double array of length `ROWS` containing the values of slack variables.
`duals`	Optional: double array of length `ROWS` containing the values of dual variables.
`djs`	Optional: double array of length `COLS` containing the values of reduced costs.
`p_status`	Pointer to an `int` where the status will be returned. The status is one of: `0` Solution is loaded. `1` Solution is not loaded because the problem is in presolved status.

Example

This example loads a problem, loads a solution for the problem and then uses `XPRScrossoverlpsol` to find a basic optimal solution.

```
XPRSreadprob(prob, "problem", "");
XPRSloadlpsol(prob, x, NULL, duals, NULL, &status);
XPRScrossoverlpsol(prob, &status);
```

Further information

1. If neither variable values nor dual values are provided, the current solution in memory will be unloaded.
2. When variables `x` is `NULL`, the variables will be set to their bounds.
3. When slack variables `slack` is `NULL`, it will be computed from variables `x`. If slacks are provided, variables cannot be omitted.
4. When dual variables `duals` is `NULL`, both dual variables and reduced costs will be set to zero.
5. When reduced costs `djs` is `NULL`, it will be computed from dual variables `duals`. If reduced costs are provided, dual variables cannot be omitted.

Related topics

`XPRSgetsolution`, `XPRScrossoverlpsol`.

XPRSloadmipsol

Purpose

Loads a starting MIP solution for the problem into the Optimizer.

Topic areas

Solution, MIP Entities, Data Input

Synopsis

```
int XPRS_CC XPRSloadmipsol(XPRSprob prob, const double x[], int *p_status);
```

Arguments

`prob`	The current problem.
`x`	Double array of length `COLS` (for the original problem and not the presolve problem) containing the values of the variables.
`p_status`	Pointer to an `int` where the status will be returned. The status is one of: `-1` Solution rejected because an error occurred; `0` Solution accepted.

Example

This example loads a problem and then loads a solution found previously for the problem to help speed up the MIP search:

```
XPRSreadprob(prob,"problem","");
XPRSloadmipsol(prob,x,&status);
XPRSmipoptimize(prob,"");
```

Further information

1. When a solution is loaded, the solution is placed in temporary storage until the MIP solve is started. Only after the MIP solve has commenced and any presolve has been applied, will the loaded solution be checked and possibly accepted as a new incumbent integer solution. There are no checks performed on the solution before the MIP solve and the returned status in `XPRSloadmipsol` will always be `0` for accepted.
2. Loaded solution values will automatically be adjusted to fit within the current problem bounds.
3. It is recommended to use `XPRSaddmipsol` instead of `XPRSloadmipsol`. `XPRSaddmipsol` can be called both before a solve, to load a starting solution, and during a MIP solve, to load new solutions within callbacks. `XPRSaddmipsol` also allows for loading of infeasible or partial solutions and comes with a callback to check the status of loaded solutions.
4. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSaddmipsol`, `XPRSgetsolution`.

XPRSloadmodelcuts

Purpose

Specifies that a set of rows in the matrix will be treated as model cuts.

Topic areas

Data Input, Cuts

Synopsis

```
int XPRS_CC XPRSloadmodelcuts(XPRSprob prob, int nrows, const int
       rowind[]);
```

Arguments

`prob`	The current problem.
`nrows`	The number of model cuts.
`rowind`	An array of row indices to be treated as cuts.

Error value

`268`	Cannot perform operation on presolved matrix.

Example

This sets the first six matrix rows as model cuts in the MIP problem `myprob`.

```
int rowind[] = {0,1,2,3,4,5}
...
XPRSloadmodelcuts(prob,6,rowind);
XPRSmipoptimize(prob,"");
```

Further information

1. During presolve the model cuts are removed from the matrix and added to an internal cut pool. During the tree search, the Optimizer will regularly check this cut pool for any violated model cuts and add those that cuts off a node LP solution.
2. The model cuts must be "true" model cuts, in the sense that they are redundant at the optimal MIP solution. The Optimizer does not guarantee to add all violated model cuts, so they must not be required to define the optimal MIP solution.
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

5.9.

XPRSloadpresolvebasis

Purpose

Loads a presolved basis from the user's areas.

Topic areas

Data Input, LP

Synopsis

```
int XPRS_CC XPRSloadpresolvebasis(XPRSprob prob, const int rowstat[], const
       int colstat[]);
```

Arguments

`prob` The current problem.

`rowstat` Integer array of length `ROWS` containing the basis status of the slack, surplus or artificial variable associated with each row. The status must be one of:
- `XPRS_NONBASIC_LOWER (0)` slack, surplus or artificial is non-basic at lower bound;
- `XPRS_BASIC (1)` slack, surplus or artificial is basic;
- `XPRS_NONBASIC_UPPER (2)` slack or surplus is non-basic at upper bound.

`colstat` Integer array of length `COLS` containing the basis status of each of the columns in the matrix. The status must be one of:
- `XPRS_NONBASIC_LOWER (0)` variable is non-basic at lower bound or superbasic at zero if the variable has no lower bound;
- `XPRS_BASIC (1)` variable is basic;
- `XPRS_NONBASIC_UPPER (2)` variable is at upper bound;
- `XPRS_SUPERBASIC (3)` variable is super-basic.

Example

The following example saves the presolved basis for one problem, loading it into another:

```
int rows, cols, *rowstat, *colstat;
...
XPRSreadprob(prob,"myprob","");
XPRSmipoptimize(prob,"l");
XPRSgetintattrib(prob,XPRS_ROWS,&rows);
XPRSgetintattrib(prob,XPRS_COLS,&cols);
rowstat = malloc(rows*sizeof(int));
colstat = malloc(cols*sizeof(int));
XPRSgetpresolvebasis(prob,rowstat,colstat);
XPRSreadprob(prob2,"myotherprob","");
XPRSmipoptimize(prob2,"l");
XPRSloadpresolvebasis(prob2,rowstat,colstat);
```

Related topics

`XPRSgetbasis`, `XPRSgetpresolvebasis`, `XPRSloadbasis`.

XPRSloadpresolvedirs

Purpose

Loads directives into the presolved matrix.

Topic areas

Data Input, Branching

Synopsis

```
int XPRS_CC XPRSloadpresolvedirs(XPRSprob prob, int ndirs, const int
      colind[], const int priority[], const char dir[], const double
      uppseudo[], const double downpseudo[]);
```

Arguments

`prob`	The current problem.
`ndirs`	Number of directives.
`colind`	Integer array of length `ndirs` containing the column numbers. A negative value indicates a set number (−1 being the first set, −2 the second, and so on).
`priority`	Integer array of length `ndirs` containing the priorities for the columns or sets. May be `NULL` if not required.
`dir`	Character array of length `ndirs` specifying the branching direction for each column or set: `U` the entity is to be forced up; `D` the entity is to be forced down; `N` not specified. May be `NULL` if not required.
`uppseudo`	Double array of length `ndirs` containing the up pseudo costs for the columns or sets. May be `NULL` if not required.
`downpseudo`	Double array of length `ndirs` containing the down pseudo costs for the columns or sets. May be `NULL` if not required.

Example

The following loads priority directives for column `0` in the matrix:

```
int colind[] = {0}, priority[] = {1};
...
XPRSmipoptimize(prob,"l");
XPRSloadpresolvedirs(prob,1,colind,priority,NULL,NULL,NULL);
XPRSmipoptimize(prob,"");
```

Related topics

`XPRSgetdirs`, `XPRSloaddirs`.

XPRSloadqcqp, XPRSloadqcqp64

Purpose

Used to load a quadratic problem with quadratic side constraints into the Optimizer data structure. Such a problem may have quadratic terms in its objective function as well as in its constraints.

Topic areas

Problem Input, Quadratic

Synopsis

```
int XPRS_CC XPRSloadqcqp(XPRSprob prob, const char * probname, int ncols,
        int nrows, const char rowtype[], const double rhs[], const double
        rng[], const double objcoef[], const int start[], const int collen[],
        const int rowind[], const double rowcoef[], const double lb[], const
        double ub[], int nobjqcoefs, const int objqcol1[], const int
        objqcol2[], const double objqcoef[], int nqrows, const int qrowind[],
        const int nrowqcoef[], const int rowqcol1[], const int rowqcol2[],
        const double rowqcoef[]);

int XPRS_CC XPRSloadqcqp64(XPRSprob prob, const char * probname, int ncols,
        int nrows, const char rowtype[], const double rhs[], const double
        rng[], const double objcoef[], const XPRSint64 start[], const int
        collen[], const int rowind[], const double rowcoef[], const double
        lb[], const double ub[], XPRSint64 nobjqcoefs, const int objqcol1[],
        const int objqcol2[], const double objqcoef[], int nqrows, const int
        qrowind[], const XPRSint64 nrowqcoef[], const int rowqcol1[], const
        int rowqcol2[], const double rowqcoef[]);
```

Arguments

`prob`	The current problem.
`probname`	A string of up to `MAXPROBNAMELENGTH` characters containing a name for the problem.
`ncols`	Number of structural columns in the matrix.
`nrows`	Number of rows in the matrix (not including the objective row). Objective coefficients must be supplied in the `objcoef` array, and the objective function should not be included in any of the other arrays.
`rowtype`	Character array of length `nrows` containing the row types: `L` indicates a <= constraint (use this one for quadratic constraints as well); `E` indicates an = constraint; `G` indicates a >= constraint; `R` indicates a range constraint; `N` indicates a nonbinding constraint.
`rhs`	Double array of length `nrows` containing the right hand side coefficients of the rows. The right hand side value for a range row gives the **upper** bound on the row.
`rng`	Double array of length `nrows` containing the range values for range rows. Values for all other rows will be ignored. May be NULL if there are no ranged constraints. The lower bound on a range row is the right hand side value minus the range value. The sign of the range value is ignored - the absolute value is used in all cases.
`objcoef`	Double array of length `ncols` containing the objective function coefficients. This can be `NULL` to set all objective coefficients to 0 (zero).
`start`	Integer array containing the offsets in the `rowind` and `rowcoef` arrays of the start of the elements for each column. This array is of length `ncols` or, if `collen` is NULL, `length ncols+1`. If `collen` is NULL the extra entry of `start`, `start[ncols]`, contains the position in the `rowind` and `rowcoef` arrays at which an extra column would

	start, if it were present. In C, this value is also the length of the `rowind` and `rowcoef` arrays.
`collen`	Integer array of length `ncols` containing the number of nonzero elements in each column. May be NULL if all elements are contiguous and `start[ncols]` contains the offset where the elements for column `ncols+1` would start. This array is not required if the non-zero coefficients in the `rowind` and `rowcoef` arrays are continuous, and the `start` array has `ncols+1` entries as described above. It may be NULL if not required.
`rowind`	Integer array containing the row indices for the nonzero elements in each column. If the indices are input contiguously, with the columns in ascending order, the length of the `rowind` is `start[ncols-1]+collen[ncols-1]` or, if `collen` is NULL, `start[ncols]`.
`rowcoef`	Double array containing the nonzero element values; length as for `rowind`.
`lb`	Double array of length `ncols` containing the lower bounds on the columns. Use `XPRS_MINUSINFINITY` to represent a lower bound of minus infinity. If this is `NULL` then all lower bounds are 0 (zero).
`ub`	Double array of length `ncols` containing the upper bounds on the columns. Use `XPRS_PLUSINFINITY` to represent an upper bound of plus infinity. It this is `NULL` then all upper bounds are infinite.
`nobjqcoefs`	Number of quadratic terms.
`objqcol1`	Integer array of size `nobjqcoefs` containing the column index of the first variable in each quadratic term.
`objqcol2`	Integer array of size `nobjqcoefs` containing the column index of the second variable in each quadratic term.
`objqcoef`	Double array of size `nobjqcoefs` containing the quadratic coefficients.
`nqrows`	Number of rows containing quadratic matrices.
`qrowind`	Integer array of size `nqrows`, containing the indices of rows with quadratic matrices in them. Note that the rows are expected to be defined in `rowtype` as type `L`.
`nrowqcoef`	Integer array of size `nqrows`, containing the number of nonzeros in each quadratic constraint matrix.
`rowqcol1`	Integer array of size `nqcelem`, where `nqcelem` equals the sum of the elements in `nrowqcoef` (i.e. the total number of quadratic matrix elements in all the constraints). It contains the first column indices of the quadratic matrices. Indices for the first matrix are listed from `0` to `nrowqcoef[0]-1`, for the second matrix from `nrowqcoef[0]` to `nrowqcoef[0]+ nrowqcoef[1]-1`, etc.
`rowqcol2`	Integer array of size `nqcelem`, containing the second index for the quadratic constraint matrices.
`rowqcoef`	Integer array of size `nqcelem`, containing the coefficients for the quadratic constraint matrices.

Related controls

Integer

`EXTRACOLS`	Number of extra columns to be allowed for.
`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAMIPENTS`	Number of extra MIP entities to be allowed for.
`EXTRAROWS`	Number of extra rows to be allowed for.
`KEEPNROWS`	Status for nonbinding rows.
`SCALING`	Type of scaling.

Double

`MATRIXTOL`	Tolerance on matrix elements.

Example

To load the following problem presented in LP format:

```
minimize [ x^2 ]
s.t.
4 x + y <= 4
x + y + [z^2] <= 5
 [ x^2 + 2 x*y + y^2 + 4 y*z + z^2 ] <= 10
x + 2 y >= 8
 [ 3 y^2 ] <= 20
end
```

the following code may be used:

```
{
    int ncols = 3;
    int nrows = 5;
    char rowtypes[] = {'L','L','L','G','L'};
    double rhs[] = {4,5,10,8,20};
    double rng[] = {0,0,0,0,0};
    double objcoef[] = {0,0,0,0,0};
    int start[] = {0,3,6,6};
    int* collen = NULL;
    int mrind[] = {0,1,3,0,1,3};
    double rowcoef[] = {4,1,1,1,1,2};
    double lb[] = {0,0,0};
    double ub[] = {XPRS_PLUSINFINITY,XPRS_PLUSINFINITY,
    XPRS_PLUSINFINITY};

    int nobjqcoefs = 1;
    int objqcol1[] = {0};
    int objqcol2[] = {0};
    double objqcoef[] = {1};

    int nqrows = 3;
    int qrowind[] = {1,2,4};
    int nrowqcoef[] = {1,5,1};
    int qcmcol1[] = {2,0,0,1,1,2,1};
    int qcmcol2[] = {2,0,1,1,2,2,1};
    // ! to have 2xy define 1xy (1yx will be assumed to be implicitly present)
    double rowqcoef[] = {1,9,1,8,2,7,3};
}

XPRSloadqcqp(xprob,"qcqp",ncols,nrows,rowtypes,rhs,rng,objcoef,start,
collen,mrind,rowcoef,lb,ub,nobjqcoefs,objqcol1,objqcol2,objqcoef,nqrows,qrowind,nrowqcoef,
qcmcol1,qcmcol2,rowqcoef);

}
```

Further information

1. The objective function is of the form $c^T x + 0.5\ x^T Q x$. Note that only the upper or lower triangular part of the `Q` matrix is specified, both for the objective and constraints.
2. The row and column indices follow the usual C convention of going from `0` to `nrows-1` and `0` to `ncols-1` respectively.
3. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` are defined in the Optimizer library header file.

Related topics

`XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSreadprob`.

XPRSloadqcqpglobal, XPRSloadqcqpglobal64

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSloadmiqcqp` *instead.*

Topic areas

Problem Input, Quadratic

Synopsis

```
int XPRS_CC XPRSloadqcqpglobal(XPRSprob prob, const char * probname, int
      ncols, int nrows, const char qrtypes[], const double rhs[], const
      double rng[], const double objcoef[], const int start[], const int
      collen[], const int rowind[], const double rowcoef[], const double
      lb[], const double ub[], int nobjqcoefs, const int mqcol1[], const
      int mqcol2[], const double objqcoef[], int nqrows, const int
      qrowind[], const int nrowqcoefs[], const int rowqcol1[], const int
      rowqcol2[], const double rowqcoef[], const int nentities, const int
      nsets, const char coltype[], const int entind[], const double
      limit[], const char settype[], const int setstart[], const int
      setind[], const double refval[]);

int XPRS_CC XPRSloadqcqpglobal64(XPRSprob prob, const char * probname, int
      ncols, int nrows, const char qrtypes[], const double rhs[], const
      double rng[], const double objcoef[], const XPRSint64 start[], const
      int collen[], const int rowind[], const double rowcoef[], const
      double lb[], const double ub[], XPRSint64 nobjqcoefs, const int
      mqcol1[], const int mqcol2[], const double objqcoef[], int nqrows,
      const int qrowind[], const XPRSint64 nrowqcoefs[], const int
      rowqcol1[], const int rowqcol2[], const double rowqcoef[], const int
      nentities, const int nsets, const char coltype[], const int entind[],
      const double limit[], const char settype[], const XPRSint64
      setstart[], const int setind[], const double refval[]);
```

XPRSloadqglobal, XPRSloadqglobal64

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSloadmiqp` *instead.*

Topic areas

Problem Input, Quadratic

Synopsis

```
int XPRS_CC XPRSloadqglobal(XPRSprob prob, const char *probname, int ncols,
      int nrows, const char rowtype[], const double rhs[], const double
      rng[], const double objcoef[], const int start[], const int collen[],
      const int rowind[], const double rowcoef[], const double lb[], const
      double ub[], int nobjqcoefs, const int objqcol1[], const int
      objqcol2[], const double objqcoef[], const int nentities, const int
      nsets, const char coltype[], const int entind[], const double
      limit[], const char settype[], const int setstart[], const int
      setind[], const double refval[]);

int XPRS_CC XPRSloadqglobal64(XPRSprob prob, const char *probname, int
      ncols, int nrows, const char rowtype[], const double rhs[], const
      double rng[], const double objcoef[], const XPRSint64 start[], const
      int collen[], const int rowind[], const double rowcoef[], const
      double lb[], const double ub[], XPRSint64 nobjqcoefs, const int
      objqcol1[], const int objqcol2[], const double objqcoef[], const int
      nentities, const int nsets, const char coltype[], const int entind[],
      const double limit[], const char settype[], const XPRSint64
      setstart[], const int setind[], const double refval[]);
```

XPRSloadmiqcqp, XPRSloadmiqcqp64

Purpose

Used to load a mixed integer quadratic problem with quadratic constraints into the Optimizer data structure. Such a problem may have quadratic terms in its objective function as well as in its constraints. Integer, binary, partial integer, semi-continuous and semi-continuous integer variables can be defined, together with sets of type 1 and 2. The reference row values for the set members are passed as an array rather than specifying a reference row.

Topic areas

Problem Input, Quadratic

Synopsis

```
int XPRS_CC XPRSloadmiqcqp(XPRSprob prob, const char * probname, int ncols,
      int nrows, const char rowtype[], const double rhs[], const double
      rng[], const double objcoef[], const int start[], const int collen[],
      const int rowind[], const double rowcoef[], const double lb[], const
      double ub[], int nobjqcoefs, const int objqcol1[], const int
      objqcol2[], const double objqcoef[], int nqrows, const int qrowind[],
      const int nrowqcoefs[], const int rowqcol1[], const int rowqcol2[],
      const double rowqcoef[], const int nentities, const int nsets, const
      char coltype[], const int entind[], const double limit[], const char
      settype[], const int setstart[], const int setind[], const double
      refval[]);

int XPRS_CC XPRSloadmiqcqp64(XPRSprob prob, const char * probname, int
      ncols, int nrows, const char rowtype[], const double rhs[], const
      double rng[], const double objcoef[], const XPRSint64 start[], const
      int collen[], const int rowind[], const double rowcoef[], const
      double lb[], const double ub[], XPRSint64 nobjqcoefs, const int
      objqcol1[], const int objqcol2[], const double objqcoef[], int
      nqrows, const int qrowind[], const XPRSint64 nrowqcoefs[], const int
      rowqcol1[], const int rowqcol2[], const double rowqcoef[], const int
      nentities, const int nsets, const char coltype[], const int entind[],
      const double limit[], const char settype[], const XPRSint64
      setstart[], const int setind[], const double refval[]);
```

Arguments

`prob`	The current problem.
`probname`	A string of up to `MAXPROBNAMELENGTH` characters containing a name for the problem.
`ncols`	Number of structural columns in the matrix.
`nrows`	Number of rows in the matrix (not including the objective row). Objective coefficients must be supplied in the `objcoef` array, and the objective function should not be included in any of the other arrays.
`rowtype`	Character array of length `nrows` containing the row types: `L` indicates a <= constraint (use this one for quadratic constraints as well); `E` indicates an = constraint; `G` indicates a >= constraint; `R` indicates a range constraint; `N` indicates a nonbinding constraint.
`rhs`	Double array of length `nrows` containing the right hand side coefficients of the rows. The right hand side value for a range row gives the **upper** bound on the row.
`rng`	Double array of length `nrows` containing the range values for range rows. Values for all other rows will be ignored. May be NULL if there are no ranged constraints. The lower

	bound on a range row is the right hand side value minus the range value. The sign of the range value is ignored - the absolute value is used in all cases.
`objcoef`	Double array of length `ncols` containing the objective function coefficients. This can be `NULL` to set all objective coefficients to 0 (zero).
`start`	Integer array containing the offsets in the `rowind` and `rowcoef` arrays of the start of the elements for each column. This array is of length `ncols` or, if `collen` is NULL, `length ncols+1`. If `collen` is NULL the extra entry of `start`, `start[ncols]`, contains the position in the `rowind` and `rowcoef` arrays at which an extra column would start, if it were present. In C, this value is also the length of the `rowind` and `rowcoef` arrays.
`collen`	Integer array of length `ncols` containing the number of nonzero elements in each column. May be NULL if all elements are contiguous and `start[ncols]` contains the offset where the elements for column `ncols+1` would start. This array is not required if the non-zero coefficients in the `rowind` and `rowcoef` arrays are continuous, and the `start` array has `ncols+1` entries as described above. It may be NULL if not required.
`rowind`	Integer array containing the row indices for the nonzero elements in each column. If the indices are input contiguously, with the columns in ascending order, the length of the `rowind` is `start[ncols-1]+collen[ncols-1]` or, if `collen` is NULL, `start[ncols]`.
`rowcoef`	Double array containing the nonzero element values; length as for `rowind`.
`lb`	Double array of length `ncols` containing the lower bounds on the columns. Use `XPRS_MINUSINFINITY` to represent a lower bound of minus infinity. If this is `NULL` then all lower bounds are 0 (zero).
`ub`	Double array of length `ncols` containing the upper bounds on the columns. Use `XPRS_PLUSINFINITY` to represent an upper bound of plus infinity. It this is `NULL` then all upper bounds are infinite.
`nobjqcoefs`	Number of quadratic terms.
`objqcol1`	Integer array of size `nobjqcoefs` containing the column index of the first variable in each quadratic term.
`objqcol2`	Integer array of size `nobjqcoefs` containing the column index of the second variable in each quadratic term.
`objqcoef`	Double array of size `nobjqcoefs` containing the quadratic coefficients.
`nqrows`	Number of rows containing quadratic matrices.
`qrowind`	Integer array of size `nqrows`, containing the indices of rows with quadratic matrices in them. Note that the rows are expected to be defined in `rowtype` as type `L`.
`nrowqcoefs`	Integer array of size `nqrows`, containing the number of nonzeros in each quadratic constraint matrix.
`rowqcol1`	Integer array of size `nqcelem`, where `nqcelem` equals the sum of the elements in `nrowqcoefs` (i.e. the total number of quadratic matrix elements in all the constraints). It contains the first column indices of the quadratic matrices. Indices for the first matrix are listed from `0` to `nrowqcoefs[0]-1`, for the second matrix from `nrowqcoefs[0]` to `nrowqcoefs[0]+ nrowqcoefs[1]-1`, etc.
`rowqcol2`	Integer array of size `nqcelem`, containing the second index for the quadratic constraint matrices.
`rowqcoef`	Integer array of size `nqcelem`, containing the coefficients for the quadratic constraint matrices.
`nentities`	Number of binary, integer, semi-continuous, semi-continuous integer and partial integer entities.
`nsets`	Number of SOS1 and SOS2 sets.

`coltype`	Character array of length `nentities` containing the entity types: `B` binary variables; `I` integer variables; `P` partial integer variables; `S` semi-continuous variables; `R` semi-continuous integer variables.
`entind`	Integer array of length `nentities` containing the column indices of the MIP entities.
`limit`	Double array of length `nentities` containing the integer limits for the partial integer variables and lower bounds for semi-continuous and semi-continuous integer variables (any entries in the positions corresponding to binary and integer variables will be ignored). May be `NULL` if not required.
`settype`	Character array of length `nsets` containing the set types: `1` SOS1 type sets; `2` SOS2 type sets. May be `NULL` if not required.
`setstart`	Integer array containing the offsets in the `setind` and `refval` arrays indicating the start of the sets. This array is of length `nsets+1`, the last member containing the offset where set `nsets+1` would start. May be `NULL` if not required.
`setind`	Integer array of length `setstart[nsets]-1` containing the columns in each set. May be `NULL` if not required.
`refval`	Double array of length `setstart[nsets]-1` containing the reference row entries for each member of the sets. These define the order for SOS2 constraints and may be used in branching for both types. May be `NULL` if not required.

Related controls

Integer

`EXTRACOLS`	Number of extra columns to be allowed for.
`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAMIPENTS`	Number of extra MIP entities to be allowed for.
`EXTRAROWS`	Number of extra rows to be allowed for.
`KEEPNROWS`	Status for nonbinding rows.
`SCALING`	Type of scaling.

Double

`MATRIXTOL`	Tolerance on matrix elements.

Further information

1. The objective function is of the form $c^T x + 0.5\ x^T Q x$. Note that only the upper or lower triangular part of the `Q` matrix is specified, both for the objective and constraints.

2. The row and column indices follow the usual C convention of going from `0` to `nrows-1` and `0` to `ncols-1` respectively.

3. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` are defined in the Optimizer library header file.

4. The row and column indices follow the usual C convention of going from `0` to `nrows-1` and `0` to `ncols-1` respectively.

5. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` are defined in the Optimizer library header file.

6. Semi-continuous lower bounds are taken from the limit array. If this is `NULL` then they are given a default value of `1.0`. If a semi-continuous variable has a positive lower bound then this will be used as the semi-continuous lower bound and the lower bound on the variable will be set to zero.

Related topics

XPRSloadmip, XPRSloadlp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadqp, XPRSreadprob.

XPRSloadmiqp, XPRSloadmiqp64

Purpose

Used to load a MIQP problem, hence a MIP with quadratic objective coefficients, into the Optimizer data structures. Integer, binary, partial integer, semi-continuous and semi-continuous integer variables can be defined, together with sets of type 1 and 2. The reference row values for the set members are passed as an array rather than specifying a reference row.

Topic areas

Problem Input, Quadratic

Synopsis

```
int XPRS_CC XPRSloadmiqp(XPRSprob prob, const char *probname, int ncols,
      int nrows, const char rowtype[], const double rhs[], const double
      rng[], const double objcoef[], const int start[], const int collen[],
      const int rowind[], const double rowcoef[], const double lb[], const
      double ub[], int nobjqcoefs, const int objqcol1[], const int
      objqcol2[], const double objqcoef[], const int nentities, const int
      nsets, const char coltype[], const int entind[], const double
      limit[], const char settype[], const int setstart[], const int
      setind[], const double refval[]);

int XPRS_CC XPRSloadmiqp64(XPRSprob prob, const char *probname, int ncols,
      int nrows, const char rowtype[], const double rhs[], const double
      rng[], const double objcoef[], const XPRSint64 start[], const int
      collen[], const int rowind[], const double rowcoef[], const double
      lb[], const double ub[], XPRSint64 nobjqcoefs, const int objqcol1[],
      const int objqcol2[], const double objqcoef[], const int nentities,
      const int nsets, const char coltype[], const int entind[], const
      double limit[], const char settype[], const XPRSint64 setstart[],
      const int setind[], const double refval[]);
```

Arguments

`prob`	The current problem.
`probname`	A string of up to `MAXPROBNAMELENGTH` characters containing a name for the problem.
`ncols`	Number of structural columns in the matrix.
`nrows`	Number of rows in the matrix (not including the objective). Objective coefficients must be supplied in the `objcoef` array, and the objective function should not be included in any of the other arrays.
`rowtype`	Character array of length `nrows` containing the row type: `L` indicates a $\leqslant$ constraint; `E` indicates an = constraint; `G` indicates a $\geqslant$ constraint; `R` indicates a range constraint; `N` indicates a nonbinding constraint.
`rhs`	Double array of length `nrows` containing the right hand side coefficients. The right hand side value for a range row gives the **upper** bound on the row.
`rng`	Double array of length `nrows` containing the range values for range rows. The values in the range array will only be read for `R` type rows. The entries for other type rows will be ignored. May be `NULL` if not required. The lower bound on a range row is the right hand side value minus the range value. The sign of the range value is ignored - the absolute value is used in all cases.
`objcoef`	Double array of length `ncols` containing the objective function coefficients. This can be `NULL` to set all objective coefficients to 0 (zero). This can be `NULL` to set all objective

	coefficients to 0 (zero).
`start`	Integer array containing the offsets in the `rowind` and `rowcoef` arrays of the start of the elements for each column. This array is of length `ncols` or, if `collen` is `NULL`, length `ncols+1`.
`collen`	Integer array of length `ncols` containing the number of nonzero elements in each column. May be `NULL` if not required. This array is not required if the non-zero coefficients in the `rowind` and `rowcoef` arrays are continuous, and the `start` array has `ncols+1` entries as described above. It may be `NULL` if not required.
`rowind`	Integer arrays containing the row indices for the nonzero elements in each column. If the indices are input contiguously, with the columns in ascending order, then the length of `rowind` is `start[ncols-1]+collen[ncols-1]` or, if `collen` is `NULL`, `start[ncols]`.
`rowcoef`	Double array containing the nonzero element values length as for `rowind`.
`lb`	Double array of length `ncols` containing the lower bounds on the columns. Use `XPRS_MINUSINFINITY` to represent a lower bound of minus infinity. If this is `NULL` then all lower bounds are 0 (zero).
`ub`	Double array of length `ncols` containing the upper bounds on the columns. Use `XPRS_PLUSINFINITY` to represent an upper bound of plus infinity. It this is `NULL` then all upper bounds are infinite.
`nobjqcoefs`	Number of quadratic terms.
`objqcol1`	Integer array of size `nobjqcoefs` containing the column index of the first variable in each quadratic term.
`objqcol2`	Integer array of size `nobjqcoefs` containing the column index of the second variable in each quadratic term.
`objqcoef`	Double array of size `nobjqcoefs` containing the quadratic coefficients.
`nentities`	Number of binary, integer, semi-continuous, semi-continuous integer and partial integer entities.
`nsets`	Number of SOS1 and SOS2 sets.
`coltype`	Character array of length `nentities` containing the entity types: `B` binary variables; `I` integer variables; `P` partial integer variables; `S` semi-continuous variables; `R` semi-continuous integers.
`entind`	Integer array of length `nentities` containing the column indices of the MIP entities.
`limit`	Double array of length `nentities` containing the integer limits for the partial integer variables and lower bounds for semi-continuous and semi-continuous integer variables (any entries in the positions corresponding to binary and integer variables will be ignored). May be `NULL` if not required.
`settype`	Character array of length `nsets` containing: `1` SOS1 type sets; `2` SOS2 type sets. May be `NULL` if not required.
`setstart`	Integer array containing the offsets in the `setind` and `refval` arrays indicating the start of the sets. This array is of length `nsets+1`, the last member containing the offset where set `nsets+1` would start. May be `NULL` if not required.
`setind`	Integer array of length `setstart[nsets]-1` containing the columns in each set. May be `NULL` if not required.
`refval`	Double array of length `setstart[nsets]-1` containing the reference row entries for each member of the sets. These define the order for SOS2 constraints and may be used

in branching for both types. May be NULL if not required.

Related controls

Integer

EXTRACOLS	Number of extra columns to be allowed for.
EXTRAELEMS	Number of extra matrix elements to be allowed for.
EXTRAMIPENTS	Number of extra MIP entities to be allowed for.
EXTRAROWS	Number of extra rows to be allowed for.
KEEPNROWS	Status for nonbinding rows.
SCALING	Type of scaling.

Double

MATRIXTOL	Tolerance on matrix elements.
SOSREFTOL	Minimum gap between reference row entries.

Example

Minimize $-6x_1 + 2x_1^2 - 2x_1x_2 + 2x_2^2$ subject to $x_1 + x_2 \leqslant 1.9$, where x_1 must be an integer:

```
int nrows = 1, ncols = 2, nquad = 3;
int start[] = {0, 1, 2};
int rowind[] = {0, 0};
double rowcoef[] = {1, 1};
double rhs[] = {1.9};
char rowtype[] = {'L'};
double lbound[] = {0, 0};
double ubound[] = {XPRS_PLUSINFINITY, XPRS_PLUSINFINITY};

double objcoef[] = {-6, 0};
int objqcol1[] = {0, 0, 1};
int objqcol2[] = {0, 1, 1};
double dquad[] = {4, -2, 4};

int nentities = 1, nsets = 0;
int entind[] = {0};
char coltype[]={'I'};

double *primal, *dual;

primal = malloc(ncols*sizeof(double));
dual = malloc(nrows*sizeof(double));
...
XPRSloadmiqp(prob, "myprob", ncols, nrows, rowtype, rhs,
                NULL, objcoef, start, NULL, rowind,
                rowcoef, lbound, ubound, nquad, objqcol1, objqcol2,
                dquad, nentities, nsets, coltype, entind, NULL,
                NULL, NULL, NULL, NULL)
```

Further information

1. The objective function is of the form `c'x+ 0.5 x'Qx`. Note that only the upper or lower triangular part of the `Q` matrix is specified.
2. The row and column indices follow the usual C convention of going from `0` to `nrows-1` and `0` to `ncols-1` respectively.
3. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` are defined in the Optimizer library header file.

Related topics

`XPRSloadmip`, `XPRSloadlp`, `XPRSloadqp`, `XPRSreadprob`.

XPRSloadqp, XPRSloadqp64

Purpose

Used to load a quadratic problem into the Optimizer data structure. Such a problem may have quadratic terms in its objective function, although not in its constraints.

Topic areas

Problem Input, Quadratic

Synopsis

```
int XPRS_CC XPRSloadqp(XPRSprob prob, const char *probname, int ncols, int
      nrows, const char rowtype[], const double rhs[], const double rng[],
      const double objcoef[], const int start[], const int collen[], const
      int rowind[], const double rowcoef[], const double lb[], const double
      ub[], int nobjqcoefs, const int objqcol1[], const int objqcol2[],
      const double objqcoef[]);

int XPRS_CC XPRSloadqp64(XPRSprob prob, const char *probname, int ncols,
      int nrows, const char rowtype[], const double rhs[], const double
      rng[], const double objcoef[], const XPRSint64 start[], const int
      collen[], const int rowind[], const double rowcoef[], const double
      lb[], const double ub[], XPRSint64 nobjqcoefs, const int objqcol1[],
      const int objqcol2[], const double objqcoef[]);
```

Arguments

`prob`	The current problem.
`probname`	A string of up to `MAXPROBNAMELENGTH` characters containing a names for the problem.
`ncols`	Number of structural columns in the matrix.
`nrows`	Number of rows in the matrix (not including the objective row). Objective coefficients must be supplied in the `objcoef` array, and the objective function should not be included in any of the other arrays.
`rowtype`	Character array of length `nrows` containing the row types: `L` indicates a ⩽ constraint; `E` indicates an = constraint; `G` indicates a ⩾ constraint; `R` indicates a range constraint; `N` indicates a nonbinding constraint.
`rhs`	Double array of length `nrows` containing the right hand side coefficients of the rows. The right hand side value for a range row gives the **upper** bound on the row.
`rng`	Double array of length `nrows` containing the range values for range rows. Values for all other rows will be ignored. May be `NULL` if there are no ranged constraints. The lower bound on a range row is the right hand side value minus the range value. The sign of the range value is ignored - the absolute value is used in all cases.
`objcoef`	Double array of length `ncols` containing the objective function coefficients. This can be `NULL` to set all objective coefficients to 0 (zero).
`start`	Integer array containing the offsets in the `rowind` and `rowcoef` arrays of the start of the elements for each column. This array is of length `ncols` or, if `collen` is `NULL`, length `ncols+1`. If `collen` is `NULL` the extra entry of `start`, `start[ncols]`, contains the position in the `rowind` and `rowcoef` arrays at which an extra column would start, if it were present. In C, this value is also the length of the `rowind` and `rowcoef` arrays.
`collen`	Integer array of length `ncols` containing the number of nonzero elements in each column. May be `NULL` if all elements are contiguous and `start[ncols]` contains the offset where the elements for column `ncols+1` would start. This array is not required if

	the non-zero coefficients in the `rowind` and `rowcoef` arrays are continuous, and the `start` array has `ncols+1` entries as described above. It may be `NULL` if not required.
`rowind`	Integer array containing the row indices for the nonzero elements in each column. If the indices are input contiguously, with the columns in ascending order, the length of the `rowind` is `start[ncols-1]+collen[ncols-1]` or, if `collen` is `NULL`, `start[ncols]`.
`rowcoef`	Double array containing the nonzero element values; length as for `rowind`.
`lb`	Double array of length `ncols` containing the lower bounds on the columns. Use `XPRS_MINUSINFINITY` to represent a lower bound of minus infinity. If this is `NULL` then all lower bounds are 0 (zero).
`ub`	Double array of length `ncols` containing the upper bounds on the columns. Use `XPRS_PLUSINFINITY` to represent an upper bound of plus infinity. It this is `NULL` then all upper bounds are infinite.
`nobjqcoefs`	Number of quadratic terms.
`objqcol1`	Integer array of size `nobjqcoefs` containing the column index of the first variable in each quadratic term.
`objqcol2`	Integer array of size `nobjqcoefs` containing the column index of the second variable in each quadratic term.
`objqcoef`	Double array of size `nobjqcoefs` containing the quadratic coefficients.

Related controls

Integer

`EXTRACOLS`	Number of extra columns to be allowed for.
`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAROWS`	Number of extra rows to be allowed for.
`KEEPNROWS`	Status for nonbinding rows.
`SCALING`	Type of scaling.

Double

`MATRIXTOL`	Tolerance on matrix elements.

Example

Minimize $-6x_1 + 2x_1^2 - 2x_1x_2 + 2x_2^2$ subject to $x_1 + x_2 \leqslant 1.9$:

```
int nrows = 1, ncols = 2, nquad = 3;
int start[] = {0, 1, 2};
int rowind[] = {0, 0};
double rowcoef[] = {1, 1};
double rhs[] = {1.9};
char rowtype[] = {'L'};
double lbound[] = {0, 0};
double ubound[] = {XPRS_PLUSINFINITY,XPRS_PLUSINFINITY};

double objcoef[] = {-6, 0};
int objqcol1[] = {0, 0, 1};
int objqcol2[] = {0, 1, 1};
double dquad[] = {4, -2, 4};

double *primal, *dual;

primal = malloc(ncols*sizeof(double));
dual = malloc(nrows*sizeof(double));
...
XPRSloadqp(prob, "example", ncols, nrows, rowtype, rhs,
           NULL, objcoef, start, NULL, rowind, rowcoef,
           lbound, ubound, nquad, objqcol1, objqcol2, dquad)
```

Further information

1. The objective function is of the form `c'x+ 0.5 x'Qx`. Note that only the upper or lower triangular part of the Q matrix is specified.
2. The row and column indices follow the usual C convention of going from `0` to `nrows-1` and `0` to `ncols-1` respectively.
3. The double constants `XPRS_PLUSINFINITY` and `XPRS_MINUSINFINITY` are defined in the Optimizer library header file.

Related topics

`XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSreadprob`.

XPRSloadsecurevecs

Purpose

Allows the user to mark rows and columns in order to prevent the presolve removing these rows and columns from the matrix.

Topic areas

Data Input, Presolve

Synopsis

```
int XPRS_CC XPRSloadsecurevecs(XPRSprob prob, int nrows, int ncols, const
       int rowind[], const int colind[]);
```

Arguments

`prob`	The current problem.
`nrows`	Number of rows to be marked.
`ncols`	Number of columns to be marked.
`rowind`	Integer array of length `nrows` containing the rows to be marked. May be `NULL` if not required.
`colind`	Integer array of length `ncols` containing the columns to be marked. May be `NULL` if not required.

Example

This sets the first six rows and the first four columns to not be removed during presolve.

```
int rowind[] = {0,1,2,3,4,5};
int colind[] = {0,1,2,3};
...
XPRSreadprob(prob,"myprob","");
XPRSloadsecurevecs(prob,6,4,rowind,colind);
XPRSmipoptimize(prob,"");
```

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

5.3.

XPRSlpoptimize LPOPTIMIZE

Purpose

This function begins a search for the optimal continuous (LP) solution. The direction of optimization is given by OBJSENSE. The status of the problem when the function completes can be checked using LPSTATUS. Any MIP entities in the problem will be ignored.

Topic areas

Solution Process, LP

Synopsis

```
int XPRS_CC XPRSlpoptimize(XPRSprob prob, const char *flags);
LPOPTIMIZE [-flags]
```

Arguments

prob	The current problem.
flags	Flags to pass to XPRSlpoptimize (LPOPTIMIZE). The default is "" or NULL, in which case the algorithm used is determined by the DEFAULTALG control. If the argument includes: b the problem will be solved using the Newton barrier method, or the Hybrid gradient method if BARALG is set to 4; p the problem will be solved using the primal simplex algorithm; d the problem will be solved using the dual simplex algorithm; n (lower case N), the network part of the problem will be identified and solved using the network simplex algorithm;

Further information

1. The algorithm used to optimize is determined by the DEFAULTALG control if no flags are provided. By default, the dual simplex is used for linear problems and the barrier is used for nonlinear problems.

2. The d and p flags can be used with the n flag to complete the solution of the problem with either the dual or primal algorithms once the network algorithm has solved the network part of the problem.

3. The b flag cannot be used with the n flag.

Related topics

XPRSmipoptimize (MIPOPTIMIZE), 4.

Example

See also examples loadlp.c, roundint.c, tableau.c .

XPRSmaxim, XPRSminim — MAXIM, MINIM

Purpose

This subroutine is deprecated and will be removed in a future release. `XPRSlpoptimize` or `XPRSmipoptimize` should be used instead.
Begins a search for the optimal LP solution.

Topic area

Solution Process

Synopsis

```
int XPRS_CC XPRSmaxim(XPRSprob prob, const char *flags);
int XPRS_CC XPRSminim(XPRSprob prob, const char *flags);
MAXIM [-flags]
MINIM [-flags]
```

Arguments

prob	The current problem.
flags	Flags to pass to XPRSmaxim (MAXIM) or XPRSminim (MINIM). The default is "" or NULL, in which case the algorithm used is determined by the DEFAULTALG control. If the argument includes: b the problem will be solved using the Newton barrier method, or the Hybrid gradient method if BARALG is set to 4; p the problem will be solved using the primal simplex algorithm; d the problem will be solved using the dual simplex algorithm; l (lower case L), the problem will be solved as a linear problem ignoring the discreteness of MIP entities, unless the g flag is also provided (ref. note below); n (lower case N), the network part of the problem will be identified and solved using the network simplex algorithm; g the MIP problem will be solved by calling XPRSmipoptimize (MIPOPTIMIZE). Certain combinations of options may be used where this makes sense so, for example, pg will solve the LP with the primal algorithm and then go on to perform the tree search.

Related controls

Integer

AUTOPERTURB	Whether automatic perturbation is performed.
BARHGMAXRESTARTS	Maximum number of restarts for the hybrid gradient method.
BARHGOPS	Control options for the Hybrid gradient method.
BARITERLIMIT	Maximum number of Newton Barrier iterations.
BARORDER	Ordering algorithm for the Cholesky factorization.
BARORDERTHREADS	Maximum number of threads for the ordering algorithm.
BAROUTPUT	Newton barrier and hybrid gradient: level of solution output.
BARTHREADS	Newton barrier and hybrid gradient maximum number of threads to run on.
BIGMMETHOD	Specifies "Big M" method, or phaseI/phaseII.
CACHESIZE	Cache size in Kbytes for the Newton barrier.
CPUTIME	1 for CPU time; 0 for elapsed time.
CRASH	Type of crash.
CROSSOVER	Newton barrier and hybrid gradient crossover control.
DEFAULTALG	Algorithm to use with the tree search.
DENSECOLLIMIT	Columns with this many elements are considered dense.
DUALGRADIENT	Pricing method for the dual algorithm.
INVERTFREQ	Invert frequency.

INVERTMIN	Minimum number of iterations between inverts.
KEEPBASIS	Whether to use previously loaded basis.
LPITERLIMIT	Iteration limit for the simplex algorithms.
LPLOG	Frequency and type of simplex algorithm log.
MAXTIME	Maximum time allowed.
PRESOLVE	Degree of presolving to perform.
PRESOLVEOPS	Specifies the operations performed during presolve.
PRICINGALG	Type of pricing to be used.
REFACTOR	Indicates whether to re-factorize the optimal basis.
TRACE	Control of the infeasibility diagnosis during presolve.

Double

BARDUALSTOP	Newton barrier and hybrid gradient tolerance for dual infeasibilities.
BARGAPSTOP	Newton barrier and hybrid gradient tolerance for relative duality gap.
BARHGEXTRAPOLATE	Extrapolation parameter for the hybrid gradient method.
BARPRIMALSTOP	Newton barrier and hybrid gradient tolerance for primal infeasibilities.
BARSTEPSTOP	Newton barrier and hybrid gradient minimal step size.
BIGM	Infeasibility penalty.
CHOLESKYTOL	Tolerance in the Cholesky decomposition.
DUALPERTURB	Dual perturbation value.
ELIMTOL	Markowitz tolerance for elimination phase of presolve.
ETATOL	Tolerance on eta elements.
FEASTOL	Tolerance on RHS.
MARKOWITZTOL	Markowitz tolerance for the factorization.
MIPABSCUTOFF	Cutoff set after an LP Optimizer command. (Dual only)
OPTIMALITYTOL	Reduced cost tolerance.
PENALTY	Maximum absolute penalty variable coefficient.
PIVOTTOL	Pivot tolerance.
PPFACTOR	Partial pricing candidate list sizing parameter.
PRIMALPERTURB	Primal perturbation value.
RELPIVOTTOL	Relative pivot tolerance.

Example 1 (Library)

```
XPRSmaxim(prob,"b");
```

This maximizes the current problem using the Newton barrier method.

Example 2 (Console)

```
MINIM -g
```

This minimizes the current problem and commences the tree search.

Further information

1. This function does not apply to general nonlinear problems and will solve the problem without the general nonlinearities.
2. The algorithm used to optimize is determined by the `DEFAULTALG` control. By default, the dual simplex is used for LP and MIP problems and the barrier is used for QP problems.
3. The `d` and `p` flags can be used with the `n` flag to complete the solution of the model with either the dual or primal algorithms once the network algorithm has solved the network part of the model.
4. The `b` flag cannot be used with the `n` flag.
5. The dual simplex algorithm is a two phase algorithm which can remove dual infeasibilities.
6. If both the `g` and `l` flags are used, the problem will be solved as a MIP and the `l` flag will be passed to `XPRSmipoptimize`, which will stop the solve after the initial LP relaxation has been solved.
7. (*Console*) If the user prematurely terminates the solution process by typing CTRL-C, the iterative procedure will terminate at the first "safe" point.

Related topics

`XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSreadbasis` (`READBASIS`), 4, A.7.

XPRSmipoptimize — MIPOPTIMIZE

Purpose

This function begins a tree search for the optimal MIP solution. The direction of optimization is given by OBJSENSE. The status of the problem when the function completes can be checked using MIPSTATUS.

Topic areas

Solution Process, MIP Entities

Synopsis

```
int XPRS_CC XPRSmipoptimize(XPRSprob prob, const char *flags);
MIPOPTIMIZE [-flags]
```

Arguments

prob	The current problem.
flags	Flags to pass to XPRSmipoptimize (MIPOPTIMIZE), which specifies how to solve the initial continuous problem where the MIP entities are relaxed. If the argument includes: b — the initial continuous relaxation will be solved using the Newton barrier method (or the hybrid gradient method if BARALG is set to 4); p — the initial continuous relaxation will be solved using the primal simplex algorithm; d — the initial continuous relaxation will be solved using the dual simplex algorithm; n — the network part of the initial continuous relaxation will be identified and solved using the network simplex algorithm; l — stop after having solved the initial continous relaxation.

Further information

1. If the l flag is used, the Optimizer will stop immediately after solving the initial continuous relaxation. The status of the continuous solve can be checked with LPSTATUS and standard LP results are available, such as the objective value (LPOBJVAL) and solution (use XPRSgetlpsol), depending on LPSTATUS.

2. It is possible for the Optimizer to find integer solutions before solving the initial continuous relaxation, either through heuristics or by having the user load an initial integer solution. This can potentially result in the tree search finishing before solving the continuous relaxation to optimality.

3. If the function returns without having completed the search for an optimal solution, the search can be resumed from where it stopped by calling XPRSmipoptimize again.

4. The algorithm used to reoptimize the continuous relaxations during the tree search is given by DEFAULTALG. The default is to use the dual simplex algorithm.

Related topics

XPRSlpoptimize (LPOPTIMIZE), 4.

Example

See also examples addmipsol.c, fixbv.c, globjpar.c, glrhspar.c, knapsack.c, mipsolpool.c, mostviolated.c, roundint.c, savesol.c, trimloss.c, tsp.c .

XPRSobjsa

Purpose

Returns upper and lower sensitivity ranges for specified objective function coefficients. If the objective coefficients are varied within these ranges the current basis remains optimal and the reduced costs remain valid.

Topic areas

Sensitivity Analysis, LP

Synopsis

```
int XPRS_CC XPRSobjsa(XPRSprob prob, int ncols, const int colind[], double
       lower[], double upper[]);
```

Arguments

`prob`	The current problem.
`ncols`	Number of objective function coefficients whose sensitivity is sought.
`colind`	Integer array of length `ncols` containing the indices of the columns whose objective function coefficients sensitivity ranges are required.
`lower`	Double array of length `ncols` where the objective function lower range values are to be returned.
`upper`	Double array of length `ncols` where the objective function upper range values are to be returned.

Example

Here we obtain the objective function ranges for the three columns: `2`, `6` and `8`:

```
colind[0] = 2; colind[1] = 8; colind[2] = 6;
XPRSobjsa(prob,3,colind,lower,upper);
```

After which lower and upper contain:

```
lower[0] = 5.0; upper[0] = 7.0;
lower[1] = 3.8; upper[1] = 5.2;
lower[2] = 5.7; upper[2] = 1e+20;
```

Meaning that the current basis remains optimal when $5.0 \leqslant C_2 \leqslant 7.0$, $3.8 \leqslant C_8 \leqslant 5.2$ and $5.7 \leqslant C_6$, C_i being the objective coefficient of column `i`.

Further information

1. This function does not currently apply to general nonlinear problems.
2. `XPRSobjsa` can only be called when an optimal solution to the current LP has been found. It cannot be used when the problem is MIP presolved.

Related topics

`XPRSrhssa`, `XPRSbndsa`.

XPRSoptimize OPTIMIZE

Purpose

This function begins a search for the optimal solution of the problem. The direction of optimization is given by OBJSENSE.

Topic area

Solution Process

Synopsis

```
int XPRS_CC XPRSoptimize(XPRSprob prob, const char *flags, int
       *solvestatus, int *solstatus);
OPTIMIZE [-flags]
```

Arguments

prob	The current problem.
flags	Flags to pass to XPRSoptimize (OPTIMIZE). The default is "" or NULL. If the argument includes: s solve the problem to local optimality; x solve the problem to global optimality; l if a branch and bound search is necessary to solve the problem, stop after solving the root node.
solvestatus	The solve status after termination. Takes the same values as SOLVESTATUS
solstatus	The solution status after termination. Takes the same values as SOLSTATUS

Further information

1. A problem is identified as a MIP if it contains any of the following: MIP entities, piecewise linear constraints or general constraints, or any of the callbacks set by XPRSaddcboptnode, XPRSaddcbpreintsol or XPRSaddcbchgbranchobject.

2. XPRSoptimize automatically selects the optimization method based on the problem type. If the problem contains nonlinear formulas or a non-convex quadratic constraint or objective function, it is equivalent to calling XPRSnlpoptimize, except that the g flag will be applied by default. Otherwise, for a problem identified as a MIP, it is equivalent to calling XPRSmipoptimize and for all other problems it is equivalent to calling XPRSlpoptimize. To determine which method was selected check the OPTIMIZETYPEUSED attribute. If the value is XPRS_OPTIMIZETYPE_LOCAL, then you can check the LOCALSOLVERSELECTED attribute to determine which local solver was selected.

3. Passing the s flag is equivalent to setting the NLPSOLVER control to XSLP_NLPSOLVER_LOCAL. Passing the x flag is equivalent to setting the NLPSOLVER control to XSLP_NLPSOLVER_GLOBAL. If neither flag is passed, and if the NLPSOLVER control is XSLP_NLPSOLVER_AUTOMATIC, then this decision will be made based on problem attributes, including convexity and the presence of user functions, and which features are authorized by the license file.

4. Any additional flags not listed above will be treated in the same way as for XPRSlpoptimize, XPRSmipoptimize and XPRSnlpoptimize, depending on the type of optimization performed. The DEFAULTALG control will also behave in the same way as for these functions.

5. If there is a solve in progress, XPRSoptimize will always try to continue that solve (similar to XPRSmipoptimize but unlike XPRSnlpoptimize, which would only do so if the -c flag was given).

6. The method used to solve the problem is stored in the OPTIMIZETYPEUSED attribute.

7. Regarding the l flag, a branch and bound search is required when a problem is identified as a MIP, or when the global solver is selected. This includes a problem with non-linear formulas that can be reformulated as a MIP (for example, if the formulas contain only piecewise linear functions, min/max functions or convex quadratic constraints). In these cases, passing the l flag will cause the Optimizer to stop after solving the initial LP relaxation.

Related topics

XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRSnlpoptimize (NLPOPTIMIZE), 4.

Example

See also examples els_managedcuts.c, els_usercuts.c, goalprog.c, Polygon_initialvalue.c, Polygon_textformula.c, Polygon_tokens.c, Polygon_userfunc.c, Polygon_userfunc_map.c, Polygon_userfunc_mapdelta.c, Polygon_userfunc_multimap.c, Polygon_userfunc_multimapdelta.c, Polygon_userfunc_vecmap.c, Polygon_userfunc_vecmapdelta.c .

XPRSgetsolution

Purpose

Returns the incumbent solution during or after optimization with XPRSoptimize, XPRSmipoptimize, XPRSlpoptimize or XPRSnlpoptimize.

Topic area

Solution

Synopsis

```
int XPRS_CC XPRSgetsolution(XPRSprob prob, int *status, double x[], int
        first, int last);
```

Arguments

prob	The current problem.
status	Information about the solution returned.
x	Double array of length last-first+1 where the value of the primal variables will be returned. May be NULL if not required.
first	First column in the solution.
last	Last column in the solution.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Column indices should be in the range 0 to INPUTCOLS-1.
2. The status will be set to XPRS_SOLAVAILABLE_NOTFOUND (0), XPRS_SOLAVAILABLE_OPTIMAL (1), or XPRS_SOLAVAILABLE_FEASIBLE (2).
3. During a deterministic MIP solve, the incumbent solution is only available once it has been deterministically communicated to the current MIP worker problem.
4. During an SLP solve, the incumbent solution is not necessarily feasible. It is generally the best solution seen so far in terms of feasibility and the objective function, depending on the values of the SLPFILTER and NLPMERITLAMBDA controls.

Related topics

XPRSoptimize, XPRSgetduals, XPRSgetslacks, XPRSgetredcosts.

Example

See also examples addmipsol.c, ComplexUserFunctions.c, goalprog.c, repair.c, roundint.c, tableau.c, trimloss.c, tsp.c .

XPRSgetslacks

Purpose

Returns the slack values

from the incumbent solution during or after optimization with XPRSoptimize, XPRSmipoptimize, XPRSlpoptimize or XPRSnlpoptimize.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetslacks(XPRSprob prob, int *status, double slacks[], int
        first, int last);
```

Arguments

prob	The current problem.
status	Information about the slacks returned.
slacks	Double array of length last-first+1 where the value of the slack variables will be returned. May be NULL if not required.
first	First row in the slacks.
last	Last row in the slacks.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Row indices should be in the range 0 to INPUTROWS-1.
2. The status will be set to XPRS_SOLAVAILABLE_NOTFOUND (0), XPRS_SOLAVAILABLE_OPTIMAL (1), or XPRS_SOLAVAILABLE_FEASIBLE (2).
3. During a deterministic MIP solve, the incumbent solution is only available once it has been deterministically communicated to the current MIP worker problem.

Related topics

XPRSoptimize, XPRSgetduals, XPRSgetsolution, XPRSgetredcosts.

Example

See also examples repair.c, tableau.c .

XPRSgetduals

Purpose

Returns the dual values

from the incumbent solution during or after optimization of a continuous problem with `XPRSoptimize`, `XPRSlpoptimize` or `XPRSnlpoptimize`.

Topic areas

Solution, LP

Synopsis

```
int XPRS_CC XPRSgetduals(XPRSprob prob, int *status, double duals[], int
        first, int last);
```

Arguments

`prob`	The current problem.
`status`	Information about the dual solution returned.
`duals`	Double array of length `last-first+1` where the values of the dual variables will be returned. May be `NULL` if not required.
`first`	First row in the dual solution.
`last`	Last row in the dual solution.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Row indices should be in the range `0` to `INPUTROWS-1`.
2. The `status` will be set to `XPRS_SOLAVAILABLE_NOTFOUND` (0), `XPRS_SOLAVAILABLE_OPTIMAL` (1), or `XPRS_SOLAVAILABLE_FEASIBLE` (2).
3. Dual values are available when the problem is solved with SLP, Knitro or any LP or convex QP algorithm: the `status` returned will be the same as for `XPRSgetsolution`. Dual values are not available during or after MIP or global solves: `status` will be `XPRS_SOLAVAILABLE_NOTFOUND`.

Related topics

`XPRSoptimize`, `XPRSgetsolution`, `XPRSgetslacks`, `XPRSgetredcosts`.

XPRSgetredcosts

Purpose

Returns the reduced costs

from the incumbent solution during or after optimization of a continuous problem with `XPRSoptimize`, `XPRSlpoptimize` or `XPRSnlpoptimize`.

Topic areas

Linear Algebra, LP

Synopsis

```
int XPRS_CC XPRSgetredcosts(XPRSprob prob, int *status, double djs[], int
        first, int last);
```

Arguments

`prob`	The current problem.
`status`	Information about the reduced costs returned.
`djs`	Double array of length `last-first+1` where the reduced costs for the variables will be returned. May be `NULL` if not required.
`first`	First column in the reduced costs.
`last`	Last column in the reduced costs.

Further information

1. The function always returns data in the input space, even if the problem is currently in presolved state. Column indices should be in the range `0` to `INPUTCOLS-1`.
2. The `status` will be set to `XPRS_SOLAVAILABLE_NOTFOUND` (0), `XPRS_SOLAVAILABLE_OPTIMAL` (1), or `XPRS_SOLAVAILABLE_FEASIBLE` (2).
3. Reduced costs are available when the problem is solved with SLP, Knitro or any LP or convex QP algorithm: the `status` returned will be the same as for `XPRSgetsolution`. Reduced costs are not available during or after MIP or global solves: `status` will be `XPRS_SOLAVAILABLE_NOTFOUND`.

Related topics

`XPRSoptimize`, `XPRSgetduals`, `XPRSgetslacks`, `XPRSgetsolution`.

XPRSpivot

Purpose

Performs a simplex pivot by bringing variable `enter` into the basis and removing `leave`.

Topic areas

Linear Algebra, Simplex

Synopsis

```
int XPRS_CC XPRSpivot(XPRSprob prob, int enter, int leave);
```

Arguments

`prob`	The current problem.
`enter`	Index of row or column to enter basis.
`leave`	Index of row or column to leave basis.

Error values

425	`enter` is invalid (out of range or already basic).
426	`leave` is invalid (out of range or not eligible, e.g. nonbasic, zero pivot, etc.).

Related controls

Double

`PIVOTTOL`	Pivot tolerance.
`RELPIVOTTOL`	Relative pivot tolerance.

Example

The following brings the 7th variable into the basis and removes the 5th:

```
XPRSpivot(prob,6,4)
```

Further information

Row indices are in the range `0` to `ROWS-1`, whilst columns are in the range `ROWS+SPAREROWS` to `ROWS+SPAREROWS+COLS-1`.

Related topics

`XPRSgetpivotorder`, `XPRSgetpivots`.

XPRSpostsolve POSTSOLVE

Purpose

Postsolve the current matrix when it is in a presolved state.

Topic areas

Solution Process, Presolve

Synopsis

```
int XPRS_CC XPRSpostsolve(XPRSprob prob);
POSTSOLVE
```

Argument

`prob`	The current problem.

Further information

1. This function currently does not apply to general nonlinear problems, for general nonlinear please use `XPRSnlppostsolve`.
2. A problem is left in a presolved state whenever a LP or MIP optimization does not complete. In these cases `XPRSpostsolve` (`POSTSOLVE`) can be called to get the problem back into its original state.

Related topics

`XPRSlpoptimize`, `XPRSmipoptimize`

Example

See also example `addmipsol.c`.

XPRSpostsolvesol

Purpose

Postsolves a primal solution formulated in the presolved space into the corresponding solution formulated in the original space. The problem itself is unchanged.

Topic areas

Solution, Presolve

Synopsis

```
int XPRS_CC XPRSpostsolvesol(XPRSprob prob, const double prex[], double
      origx[]);
```

Arguments

`prob`	The current problem.
`prex`	Double array of length `COLS` with the values of the primal variables in the presolved space. Cannot be `NULL`.
`origx`	Double array of length `ORIGINALCOLS` where the values of the primal variables will be returned. Cannot be `NULL`.

Further information

This function is useful for advanced users implementing heuristics directly in the presolved space. Once a new solution is obtained, it needs to be mapped to the original space in order to be passed back into the solver.

Related topics

`XPRSgetsolution`, `XPRSgetcallbacksolution`, `XPRSgetcallbackpresolvesolution`.

XPRSpresolverow

Purpose

Presolves a row formulated in terms of the original variables such that it can be added to a presolved matrix.

Topic areas

Misc, Presolve

Synopsis

```
int XPRS_CC XPRSpresolverow(XPRSprob prob, char rowtype, int norigcoefs,
        const int origcolind[], const double origrowcoef[], double origrhs,
        int maxcoefs, int * p_ncoefs, int colind[], double rowcoef[], double
        * p_rhs, int * p_status);
```

Arguments

`prob`	The current problem.
`rowtype`	The type of the row: `L` indicates a $\leqslant$ row; `G` indicates a $\geqslant$ row.
`norigcoefs`	Number of elements in the `origcolind` and `origrowcoef` arrays.
`origcolind`	Integer array of length `norigcoefs` containing the column indices of the row to presolve.
`origrowcoef`	Double array of length `norigcoefs` containing the non-zero coefficients of the row to presolve.
`origrhs`	The right-hand side constant of the row to presolve.
`maxcoefs`	Maximum number of elements to return in the `colind` and `rowcoef` arrays.
`p_ncoefs`	Pointer to the integer where the number of elements in the `colind` and `rowcoef` arrays will be returned.
`colind`	Integer array which will be filled with the column indices of the presolved row. It must be allocated to hold at least `COLS` elements.
`rowcoef`	Double array which will be filled with the coefficients of the presolved row. It must be allocated to hold at least `COLS` elements.
`p_rhs`	Pointer to the double where the presolved right-hand side will be returned.
`p_status`	Status of the presolved row: `-3` Failed to presolve the row due to presolve dual reductions; `-2` Failed to presolve the row due to presolve duplicate column reductions; `-1` Failed to presolve the row due to an error. Check the Optimizer error code for the cause; `0` The row was successfully presolved; `1` The row was presolved, but may be relaxed.

Related controls

Integer

`PRESOLVE`	Turns presolve on or off.
`PRESOLVEOPS`	Selects the presolve operations.

Example

Suppose we want to add the row $2x_1 + x_2 \leqslant 1$ to our presolved matrix. This could be done in the following way:

```
int mindo[] = { 1, 2 };
int origrowcoef[] = { 2.0, 1.0 };
char rowtype = 'L';
double origrhs = 1.0;
```

```
int ncols, ncoefs, status, mtype, mstart[2], *mindp;
double rhs, *rowcoef;
...
XPRSgetintattrib(prob, XPRS_COLS, &ncols);
mindp = (int*) malloc(ncols*sizeof(int));
rowcoef = (double*) malloc(ncols*sizeof(double));
XPRSpresolverow(prob, rowtype, 2, mindo, origrowcoef, origrhs, ncols,
                &ncoefs, mindp, rowcoef, &rhs, &status);
if (status >= 0) {
  mtype = 0;
  mstart[0] = 0; mstart[1] = ncoefs;
  XPRSaddcuts(prob, 1, &mtype, &rowtype, &rhs, mstart, mindp,
              rowcoef);
}
```

See also examples `els_usercuts.c`, `tsp.c` .

Further information

1. This function applies to linear and convex quadratic (including second order conic) problems only.
2. There are certain presolve operations that can prevent a row from being presolved exactly. If the row contains a coefficient for a column that was eliminated due to duplicate column reductions or singleton column reductions, the row might have to be relaxed to remain valid for the presolved problem. The relaxation will be done automatically by the `XPRSpresolverow` function, but a return status of `+1` will be returned. If it is not possible to relax the row, a status of `-2` will be returned instead. Likewise, it is possible that certain dual reductions prevents the row from being presolved. In such a case a status of `-3` will be returned instead.
3. If `XPRSpresolverow` will be used for presolving hard constraints, such as e.g. branching bounds or delayed constraints, then dual reductions should be disabled, by setting `MIPDUALREDUCTIONS` to 0.
4. If the user knows in advance which columns will have non-zero coefficients in rows that will be presolved, it is possible to protect these individual columns through the `XPRSloadsecurevecs` function. This way the Optimizer is left free to apply all possible reductions to the remaining columns.

Related topics

`XPRSaddcuts`, `XPRSloadsecurevecs`, `XPRSstorecuts`.

PRINTSOL

Purpose

Writes the current solution to the screen.

Topic area

Solution

Synopsis

```
PRINTSOL [-flags]
```

Argument

`flags`	Flags to pass to `PRINTSOL`:	
	`s`	include classical sensitivity analysis.

Related controls

Integer

`MAXPAGELINES`	Number of lines between page breaks.

Double

`OUTPUTTOL`	Tolerance on print values.

Further information

See `WRITEPRTSOL` for more information.

Related topics

`XPRSgetsolution`, `XPRSwriteprtsol`.

QUIT

Purpose

Terminates the Console Optimizer, returning a zero exit code to the operating system. Alias for EXIT.

Topic area

Misc

Synopsis

`QUIT`

Example

The command is called simply as:

```
QUIT
```

Further information

1. Fatal error conditions return nonzero exit values which may be of use to the host operating system. These are described in 11.
2. If you wish to return an exit code reflecting the final solution status, then use the `STOP` command instead.

Related topics

`STOP`, `XPRSsave` (`SAVE`).

XPRSreadbasis READBASIS

Purpose
Instructs the Optimizer to read in a previously saved basis from a file.

Topic areas
File IO, LP

Synopsis
```
int XPRS_CC XPRSreadbasis(XPRSprob prob, const char *filename, const char
      *flags);
READBASIS [-flags] [filename]
```

Arguments
`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name from which the basis is to be read. If omitted, the default *problem_name* is used with a `.bss` extension.
`flags`	Flags to pass to `XPRSreadbasis` (`READBASIS`): `n` input basis file containing basic solution values; `t` input a compact advanced form of the basis; `v` use the provided filename verbatim, without appending the `.bss` extension; `z` read a compressed input file.

Example 1 (Library)
If an advanced basis is available for the current problem the Optimizer input might be:
```
XPRSreadprob(prob,"filename","");
XPRSreadbasis(prob,"","");
XPRSmipoptimize(prob,"");
```
This reads in a matrix file, inputs an advanced starting basis and maximizes the MIP.

See also example `addmipsol.c`.

Example 2 (Console)
An equivalent set of commands for the Console user may look like:
```
READPROB
READBASIS
MIPOPTIMIZE
```

Further information
1. The only check done when reading compact basis is that the number of rows and columns in the basis agrees with the current number of rows and columns.
2. `XPRSreadbasis` (`READBASIS`) will read the basis for the original problem even if the matrix has been presolved. The Optimizer will read the basis, checking that it is valid, and will display error messages if it detects inconsistencies.
3. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics
`XPRSloadbasis`, `XPRSwritebasis` (`WRITEBASIS`).

XPRSreadbinsol — READBINSOL

Purpose

Reads a solution from a binary solution file.

Topic areas

File IO, Solution

Synopsis

```
int XPRS_CC XPRSreadbinsol(XPRSprob prob, const char *filename, const char
       *flags);
READBINSOL [-flags] [filename]
```

Arguments

prob	The current problem.
filename	A string of up to MAXPROBNAMELENGTH characters containing the file name from which the solution is to be read. If omitted, the default *problem_name* is used with a .sol extension.
flags	Flags to pass to XPRSreadbinsol (READBINSOL): m load the solution as a solution for the MIP; x load the solution as a solution for the LP; v use the provided filename verbatim, without appending the .sol extension; z read a compressed input file.

Example 1 (Library)

A previously saved solution can be loaded into memory and a print file created from it with the following commands:

```
XPRSreadprob(prob, "myprob", "");
XPRSreadbinsol(prob, "", "");
XPRSwriteprtsol(prob, "", "");
```

Example 2 (Console)

An equivalent set of commands to the above for console users would be:

```
READPROB
READBINSOL
WRITEPRTSOL
```

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use XPRSpostsolve to restore the problem to a non-solving state.

Related topics

XPRSgetsolution, XPRSwritebinsol (WRITEBINSOL), XPRSwritesol (WRITESOL), XPRSwriteprtsol (WRITEPRTSOL).

XPRSreaddirs READDIRS

Purpose

Reads a directives file to help direct the tree search.

Topic areas

File IO, Branching

Synopsis

```
int XPRS_CC XPRSreaddirs(XPRSprob prob, const char *filename);
READDIRS [filename]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name from which the directives are to be read. If omitted (or `NULL`), the default *problem_name* is used with a `.dir` extension.

Related controls

Double

`PSEUDOCOST`	Default pseudo cost in node degradation estimation.

Example 1 (Library)

The following example reads in directives from the file `sue.dir` for use with the problem, `steve`:

```
XPRSreadprob(prob,"steve","");
XPRSreaddirs(prob,"sue");
XPRSmipoptimize(prob,"");
```

Example 2 (Console)

```
READPROB
READDIRS
MIPOPTIMIZE
```

This is the most usual form at the console. It will attempt to read in a directives file with the current problem name and an extension of `.dir`.

Further information

1. Directives can be given relating to priorities, forced branching directions, pseudo costs and model cuts. There is a priority value associated with each MIP entity. The *lower* the number, the *more* likely the entity is to be selected for branching; the *higher*, the *less* likely. By default, all MIP entities have a priority value of `500` which can be altered with a priority entry in the directives file. In general, it is advantageous for the entity's priority to reflect its relative importance in the model. Priority entries with values in excess of `1000` are illegal and are ignored. A full description of the directives file format may be found in A.5.
2. By default, `XPRSmipoptimize` (`MIPOPTIMIZE`) will explore the branch expected to yield the best integer solution from each node, irrespective of whether this forces the MIP entity up or down. This can be overridden with an `UP` or `DN` entry in the directives file, which forces `XPRSmipoptimize` (`MIPOPTIMIZE`) to branch up first or down first on the specified entity.
3. Pseudo-costs are estimates of the unit cost of forcing an entity up or down. By default `XPRSmipoptimize` (`MIPOPTIMIZE`) uses dual information to calculate estimates of the unit up and down costs and these are added to the default pseudo costs which are set to the `PSEUDOCOST` control. The default pseudo costs can be overridden by a `PU` or `PD` entry in the directives file.
4. If model cuts are used, then the specified constraints are removed from the matrix and added to the Optimizer cut pool, and only put back in the matrix when they are violated by an LP solution at one of the nodes in the tree search.
5. If creating a directives file by hand, wild cards can be used to specify several vectors at once, for example `PR x1* 2` will give all MIP entities whose names start with `x1` a priority of `2`.
6. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSloaddirs`, A.5.

XPRSreadprob READPROB

Purpose

Reads an (X)MPS or LP format matrix from file.

Topic areas

File IO, Problem Creation

Synopsis

```
int XPRS_CC XPRSreadprob(XPRSprob prob, const char *filename, const char
      *flags);
READPROB [-flags] [filename]
```

Arguments

`prob`	The current problem.
`filename`	The path and file name from which the problem is to be read. Limited to `MAXPROBNAMELENGTH` characters. If omitted (console users only), the default *problem_name* is used with various extensions - see below.
`flags`	Flags to be passed: `l` only `filename.lp` is searched for; `v` use the provided filename verbatim, without appending the `.mps`, `.mat` or `.lp` extension; `z` read a compressed input file.

Related controls

Integer

`EXTRACOLS`	Number of extra columns to be allowed for.
`EXTRAELEMS`	Number of extra matrix elements to be allowed for.
`EXTRAMIPENTS`	Number of extra MIP entities to be allowed for.
`EXTRAROWS`	Number of extra rows to be allowed for.
`KEEPNROWS`	Status for nonbinding rows.
`MPSECHO`	Whether MPS comments are to be echoed.
`MPSFORMAT`	Specifies format of MPS files.
`SCALING`	Type of scaling.

Double

`MATRIXTOL`	Tolerance on matrix elements.
`SOSREFTOL`	Minimum gap between reference row entries.

String

`MPSBOUNDNAME`	The active bound name.
`MPSOBJNAME`	Name of objective function row.
`MPSRANGENAME`	Name of range.
`MPSRHSNAME`	Name of right hand side.

Example 1 (Library)

```
XPRSreadprob(prob,"myprob","");
```

This instructs the Optimizer to read an MPS format matrix from the first file found out of `myprob.mat`, `myprob.mps` or (in LP format) `myprob.lp`.

See also examples `addmipsol.c`, `fixbv.c`, `globjpar.c`, `glrhspar.c`, `knapsack.c`, `mipsolenum.c`, `mipsolpool.c`, `mostviolated.c`, `repair.c`, `roundint.c`, `savesol.c`, `tableau.c`, `trimloss.c` .

Example 2 (Console)

```
READPROB -l
```

This instructs the Optimizer to read an LP format matrix from the file *problem_name* `.lp`.

Further information

1. If the function fails and the error is anything but "Xpress has not been initialized" or "out of memory" then the problem will be reset to empty and the objective sense will be reset to "minimize". If the function returns with "out of memory" then the only safe thing to do with the problem is to call `XPRSdestroyprob()`.
2. If no flags are given, file types are searched for in the order: `.mps`, `.mat`, `.lp`. Matrix files are assumed to be in XMPS or MPS format unless their file extension is `.lp` in which case they must be LP files. Files with compressed extensions are also searched for, e.g., `.mps.gz`.
3. The Optimizer can read matrix files that have been compressed using one of the following formats, as long as the command-line tool necessary to decompress the file can be located in the path: `bzip2`, `xz`, `lzma`, `Z`, `zip`, `tar`, `tgz`. The Optimizer has built-in support for the `gz` format: no external tools are necessary to read gzipped matrix files.
4. If `filename` has been specified, the problem name is changed to `filename`, ignoring any extension.
5. `XPRSreadprob` (`READPROB`) will take as the objective function the first `N` type row in the matrix, unless the string parameter `MPSOBJNAME` has been set, in which case the objective row sought will be the one named by `MPSOBJNAME`. Similarly, if non-blank, the string parameters `MPSRHSNAME`, `MPSBOUNDNAME` and `MPSRANGENAME` specify the right hand side, bound and range sets to be taken. For example:
   ```
   MPSOBJNAME="Cost"
   MPSRHSNAME="RHS 1"
   READPROB
   ```
   The treatment of `N` type rows other than the objective function depends on the `KEEPNROWS` control. If `KEEPNROWS` is `1` the rows and their elements are kept in memory; if it is `0` the rows are retained, but their elements are removed; and if it is `-1` the rows are deleted entirely. The performance impact of retaining such `N` type rows will be small unless the presolve has been disabled by setting `PRESOLVE` to `0` prior to optimization.
6. The Optimizer checks that the matrix file is in a legal format and displays error messages if it detects errors. When the Optimizer has read and verified the problem, it will display summary problem statistics.
7. By default, the `MPSFORMAT` control is set to `-1` and `XPRSreadprob` (`READPROB`) determines automatically whether the MPS files are in free or fixed format. If `MPSFORMAT` is set to `0`, fixed format is assumed and if it is set to `1`, free format is assumed. Fields in free format MPS files are delimited by one or more blank characters. The keywords NAME, ROWS, COLUMNS, QUADOBJ / QMATRIX, QCMATRIX, DELAYEDROWS, MODELCUTS, SETS, RHS, RANGES, BOUNDS and ENDATAmust start in column one and no vector name may contain blanks. If a special ordered set is specified with a reference row, its name may not be the same as that of a column. Note that numeric values which contain embedded spaces (for example after unary minus sign) will not be read correctly unless `MPSFORMAT` is set to `0`.
8. If the problem is not to be scaled automatically, set the parameter `SCALING` to `0` before issuing the `XPRSreadprob` (`READPROB`) command.

Related topics

`XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSwriteprob`.

XPRSreadslxsol READSLXSOL

Purpose

Reads an ASCII solution file (`.slx`) created by the `XPRSwriteslxsol` function.

Topic areas

File IO, Solution

Synopsis

```
int XPRS_CC XPRSreadslxsol(XPRSprob prob, const char *filename, const char
        *flags);
READSLXSOL -[flags] [filename]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name to which the solution is to be read. If omitted, the default *problem_name* is used with a `.slx` extension.
`flags`	Flags to pass to `XPRSreadslxsol` (`READSLXSOL`): `l` read the solution as an LP solution in case of a MIP problem; `m` read the solution as a solution for the MIP problem; `a` read multiple MIP solutions from the `.slx` file and add them to the MIP problem; `v` use the provided filename verbatim, without appending the `.slx` extension; `z` read a compressed input file.

Example 1 (Library)

```
XPRSreadslxsol(prob,"lpsolution","");
```

This loads the solution to the MIP problem if the problem contains MIP entities, or otherwise loads it as an LP (barrier in case of quadratic problems) solution into the problem.

Example 2 (Console)

```
READSLXSOL lpsolution
```

Which is equivalent to the library example above.

Further information

1. When `XPRSreadslxsol` is called before a MIP solve, the loaded solutions will not be checked before calling `XPRSmipoptimize`. By default, only the last MIP solution read from the `.slx` file will be stored. Use the `a` flag to store all MIP solutions read from the file.
2. When using the `a` flag, read solutions will be queued similarly to the user of the `XPRSaddmipsol` function. Each name string given by the `NAME` field in the `.slx` file will be associated with the corresponding solution. Any registered `usersolnotify` callback will be fired when the solution has been checked, and will include the read name string as one of its arguments.
3. Refer to the Appendix on Log and File Formats for a description of the ASCII Solution (.slx) File format A.4.4.

Related topics

`XPRSreadbinsol` (`READBINSOL`), `XPRSwriteslxsol` (`WRITESLXSOL`), `XPRSwritebinsol` `WRITEBINSOL`, `XPRSreadbinsol` (`READBINSOL`), `XPRSaddmipsol`, `XPRSaddcbusersolnotify`.

XPRSrefinemipsol REFINEMIPSOL

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `REFINEOPS` instead.
Executes the MIP solution refiner.

Topic areas

Numerics, Solution Refinement

Synopsis

```
int XPRS_CC XPRSrefinemipsol(XPRSprob prob, int options, const char* flags,
      const double solution[], double refined[], int* p_status);
REFINEMIPSOL
```

Arguments

`prob`	The current problem.
`options`	Refinement options: `0` Reducing MIP fractionality is priority (If bit 10 of `REFINEOPS` is set, will switch to other mode if unsuccessful). `1` Reducing LP infeasibility is priority.
`flags`	Flags passed to any optimization calls during refinement.
`solution`	The MIP solution to refine. Must be a valid MIP solution.
`refined`	The refined MIP solution in case of success
`p_status`	Refinement results: `0` An error has occurred `1` The solution has been refined `2` Current solution meets target criteria `3` Solution cannot be refined `5` The solution has been refined, but MIP fractionality could not be reduced.

Further information

1. This function currently does not apply to general nonlinear problems.
2. The function provides a mechanism to refine the MIP solution by attempting to round any fractional MIP entity and by attempting to reduce LP infeasibility.

Related topics

`REFINEOPS`.

XPRSremovecbbariteration

Purpose

Removes a barrier iteration callback function previously added by `XPRSaddcbbariteration`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Barrier

Synopsis

```
int XPRS_CC XPRSremovecbbariteration(XPRSprob prob, void (XPRS_CC
        *bariteration)(XPRSprob prob, void* vContext, int* barrier_action),
        void* data);
```

Arguments

`prob`	The current problem.
`bariteration`	The callback function to remove. If NULL then all bariteration callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all barrier iteration callbacks with the function pointer `bariteration` will be removed.

Related topics

`XPRSaddcbbariteration`.

XPRSremovecbcomputerestart

Purpose

Removes a computerestart callback function previously added by `XPRSaddcbcomputerestart`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Compute Interface

Synopsis

```
int XPRS_CC XPRSremovecbpcomputerestart(XPRSprob prob, void (XPRS_CC
       *computerestart)(XPRSprob prob, void* vContext), void* data);
```

Arguments

`prob`	The current problem.
`computerestart`	The callback function to remove. If NULL then all computerestart callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all computerestart callbacks with the function pointer `computerestart` will be removed.

Related topics

`XPRSaddcbcomputerestart`.

XPRSremovecbpresolve

Purpose

Removes a presolve callback function previously added by `XPRSaddcbpresolve`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Presolve

Synopsis

```
int XPRS_CC XPRSremovecbpresolve(XPRSprob prob, void (XPRS_CC
       *presolve)(XPRSprob prob, void* vContext), void* data);
```

Arguments

`prob`	The current problem.
`presolve`	The callback function to remove. If NULL then all presolve callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and presolve callbacks with the function pointer `presolve` will be removed.

Related topics

`XPRSaddcbpresolve`.

XPRSremovecbbarlog

Purpose

Removes a Newton barrier log callback function previously added by `XPRSaddcbbarlog`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Barrier

Synopsis

```
int XPRS_CC XPRSremovecbbarlog(XPRSprob prob, int (XPRS_CC
       *barlog)(XPRSprob prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`barlog`	The callback function to remove. If NULL then all barrier log callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all barrier log callbacks with the function pointer `barlog` will be removed.

Related topics

`XPRSaddcbbarlog`.

XPRSremovecbafterobjective

Purpose

Removes a callback function previously added by `XPRSaddcbafterobjective`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Multiobjective

Synopsis

```
int XPRS_CC XPRSremovecbafterobjective(XPRSprob prob, void (XPRS_CC
       *afterobjective)(XPRSprob prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`afterobjective`	The callback function to remove. If NULL then all after-objective callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all after-objective callbacks with the function pointer `afterobjective` will be removed.

Related topics

`XPRSaddcbafterobjective`.

XPRSremovecbbeforeobjective

Purpose

Removes a callback function previously added by `XPRSaddcbbeforeobjective`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Multiobjective

Synopsis

```
int XPRS_CC XPRSremovecbbeforeobjective(XPRSprob prob, void (XPRS_CC
       *beforeobjective)(XPRSprob prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`beforeobjective`	The callback function to remove. If NULL then all before-objective callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all before-objective callbacks with the function pointer `beforeobjective` will be removed.

Related topics

`XPRSaddcbbeforeobjective`.

XPRSremovecbnodelpsolved

Purpose

Removes a node callback function previously added by `XPRSaddcbnodelpsolved`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSremovecbnodelpsolved(XPRSprob prob, void (XPRS_CC
      *nodelpsolved)(XPRSprob my_prob, void *my_object), void* data);
```

Arguments

`prob`	The current problem.
`nodelpsolved`	The callback function to remove. If NULL then all node-optimal callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all node-optimal callbacks with the function pointer `optnode` will be removed.

Further information

1. This callback is fired only for the first continuous relaxation solve of each node (including the root) in a MIP branch-and-bound search, and only if the continuous relaxation is feasible and solved to optimality.
2. It is allowed to tighten the local node bounds and add or remove cutting planes from within this callback, but the callback will not be fired again in case such modifications triggers a resolve of the continuous relaxation of the node.

Related topics

`XPRSaddcbnodelpsolved`

XPRSremovecbchgbranch

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSremovecbchgbranchobject` instead.
Removes a variable branching callback function previously added by `XPRSaddcbchgbranch`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branching

Synopsis

```
int XPRS_CC XPRSremovecbchgbranch(XPRSprob prob, void (XPRS_CC
      *chgbranch)(XPRSprob prob, void* vContext, int* entity, int* up,
      double* estdeg), void* data);
```

Arguments

`prob`	The current problem.
`chgbranch`	The callback function to remove. If NULL then all variable branching callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all variable branching callbacks with the function pointer `chgbranch` will be removed.

Related topics

`XPRSaddcbchgbranch`.

XPRSremovecbchgbranchobject

Purpose

Removes a callback function previously added by `XPRSaddcbchgbranchobject`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branching

Synopsis

```
int XPRS_CC XPRSremovecbchgbranchobject(XPRSprob prob, void (XPRS_CC
       *chgbranchobject)(XPRSprob my_prob, void* my_object, XPRSbranchobject
       obranch, XPRSbranchobject* p_newobject), void* data);
```

Arguments

`prob`	The current problem.
`chgbranchobject`	The callback function to remove. If NULL then all branch data callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all branch data callbacks with the function pointer `chgbranchobject` will be removed.

Related topics

`XPRSaddcbchgbranchobject`

XPRSremovecbchecktime

Purpose

Removes a callback function previously added by `XPRSaddcbchecktime`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Limits

Synopsis

```
int XPRS_CC XPRSremovecbchecktime(XPRSprob prob, int (XPRS_CC
       *checktime)(XPRSprob prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`checktime`	The callback function to remove. If NULL then all checktime callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all checktime callbacks with the function pointer `checktime` will be removed.

Related topics

`XPRSaddcbchecktime`

XPRSremovecbchgnode

Purpose

This subroutine is deprecated and will be removed in a future release. Please use branching objects instead.

Removes a node selection callback function previously added by `XPRSaddcbchgnode`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSremovecbchgnode(XPRSprob prob, void (XPRS_CC
      *chgnode)(XPRSprob prob, void* data, int* nodnum), void* data);
```

Arguments

`prob`	The current problem.
`chgnode`	The callback function to remove. If NULL then all node selection callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all node selection callbacks with the function pointer `chgnode` will be removed.

Related topics

`XPRSaddcbchgnode`

XPRSremovecbcutlog

Purpose

Removes a cut log callback function previously added by `XPRSaddcbcutlog`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSremovecbcutlog(XPRSprob prob, int (XPRS_CC
      *cutlog)(XPRSprob prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`cutlog`	The callback function to remove. If NULL then all cut log callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all cut log callbacks with the function pointer `cutlog` will be removed.

Related topics

`XPRSaddcbcutlog`

XPRSremovecbcutmgr

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSremovecboptnode` instead.

Removes a cut manager callback function previously added by `XPRSaddcbcutmgr`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Cuts

Synopsis

```
int XPRS_CC XPRSremovecbcutmgr(XPRSprob prob, int (XPRS_CC
      *cutmgr)(XPRSprob prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`cutmgr`	The callback function to remove. If NULL then all cut manager callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all cut manager callbacks with the function pointer `cutmgr` will be removed.

Related topics

`XPRSaddcbcutmgr`

XPRSremovecbcutround

Purpose

Removes a cut round callback function previously added by `XPRSaddcbcutround`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Cuts

Synopsis

```
int XPRS_CC XPRSremovecbcutround(XPRSprob prob, void (XPRS_CC
       *cutround)(XPRSprob cbprob, void* cbdata, int ifxpresscuts, int
       *p_action), void* data);
```

Arguments

`prob`	The current problem.
`cutround`	The callback function to remove. If NULL then all cut round callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all cut manager callbacks with the function pointer `cutround` will be removed.

Related topics

`XPRSaddcbcutround`

XPRSremovecbdestroymt

Purpose

Removes a slave thread destruction callback function previously added by `XPRSaddcbdestroymt`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Parallel

Synopsis

```
int XPRS_CC XPRSremovecbdestroymt(XPRSprob prob, void (XPRS_CC
       *destroymt)(XPRSprob prob, void* vContext), void* data);
```

Arguments

`prob`	The current problem.
`destroymt`	The callback function to remove. If NULL then all thread destruction callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all thread destruction callbacks with the function pointer `destroymt` will be removed.

Related topics

`XPRSaddcbdestroymt`

XPRSremovecbgapnotify

Purpose

Removes a callback function previously added by `XPRSaddcbgapnotify`. The specified callback function will no longer be removed after it has been returned.

Topic areas

Callback, Limits

Synopsis

```
int XPRS_CC XPRSremovecbgapnotify(XPRSprob prob, void (XPRS_CC
      *gapnotify)(XPRSprob prob, void* vContext, double*
      newRelGapNotifyTarget, double* newAbsGapNotifyTarget, double*
      newAbsGapNotifyObjTarget, double* newAbsGapNotifyBoundTarget), void*
      data);
```

Arguments

`prob`	The current problem.
`gapnotify`	The callback function to remove. If `NULL` then all `gapnotify` callback functions added with the given user-defined pointer value will be removed.
`data`	The user-defined pointer value that the callback was added with. If `NULL` then the pointer value will not be checked and all the `gapnotify` callbacks with the function pointer `gapnotify` will be removed.

Related topics

`XPRSaddcbgapnotify`.

XPRSremovecbgloballog

Purpose

This subroutine is deprecated and will be removed in a future release. Please use `XPRSremovecbmiplog` instead.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSremovecbgloballog(XPRSprob prob, int (XPRS_CC
       *globallog)(XPRSprob prob, void* vContext), void* data);
```

XPRSremovecbmiplog

Purpose

Removes a MIP log callback function previously added by `XPRSaddcbmiplog`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSremovecbmiplog(XPRSprob prob, int (XPRS_CC
       *miplog)(XPRSprob prob, void* vContext), void* data);
```

Arguments

`prob`	The current problem.
`miplog`	The callback function to remove. If NULL then all MIP log callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all MIP log callbacks with the function pointer `miplog` will be removed.

Example

The following code sets and removes a callback function:

```
XPRSsetintcontrol(prob, XPRS_MIPLOG, 3);
XPRSaddcbmiplog(prob, mipLog, NULL, 0);
XPRSmipoptimize(prob,"");
XPRSremovecbmiplog(prob,mipLog,NULL);
}
```

Related topics

`XPRSaddcbmiplog`

XPRSremovecbinfnode

Purpose

Removes a user infeasible node callback function previously added by `XPRSaddcbinfnode`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSremovecbinfnode(XPRSprob prob, void (XPRS_CC
       *infnode)(XPRSprob prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`infnode`	The callback function to remove. If NULL then all user infeasible node callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all user infeasible node callbacks with the function pointer `infnode` will be removed.

Related topics

`XPRSaddcbinfnode`

XPRSremovecbintsol

Purpose

Removes an integer solution callback function previously added by `XPRSaddcbintsol`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Solution

Synopsis

```
int XPRS_CC XPRSremovecbintsol(XPRSprob prob, void (XPRS_CC
      *intsol)(XPRSprob prob, void* my_object), void* data);
```

Arguments

`prob`	The current problem.
`intsol`	The callback function to remove. If NULL then all integer solution callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all integer solution callbacks with the function pointer `intsol` will be removed.

Related topics

`XPRSaddcbintsol`

XPRSremovecblplog

Purpose

Removes a simplex log callback function previously added by `XPRSaddcblplog`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSremovecblplog(XPRSprob prob, int (XPRS_CC *lplog)(XPRSprob
        prob, void* data), void* data);
```

Arguments

`prob`	The current problem.
`lplog`	The callback function to remove. If NULL then all lplog callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all lplog callbacks with the function pointer `lplog` will be removed.

Example

The following code sets and removes a callback function:

```
XPRSsetintcontrol(prob,XPRS_LPLOG,10);
XPRSaddcblplog(prob,lpLog,NULL,0);
XPRSreadprob(prob,"problem","");
XPRSlpoptimize(prob,"");
XPRSremovecblplog(prob,lpLog,NULL);
}
```

Related topics

`XPRSaddcblplog`

XPRSremovecbmessage

Purpose

Removes a message callback function previously added by `XPRSaddcbmessage`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XPRSremovecbmessage(XPRSprob prob, void (XPRS_CC
      *message)(XPRSprob prob, void* vContext, const char* msg, int len,
      int msgtype), void* data);
```

Arguments

`prob`	The current problem.
`message`	The callback function to remove. If NULL then all message callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all message callbacks with the function pointer `message` will be removed.

Related topics

`XPRSaddcbmessage`

XPRSremovecbmipthread

Purpose

Removes a callback function previously added by `XPRSaddcbmipthread`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Parallel

Synopsis

```
int XPRS_CC XPRSremovecbmipthread(XPRSprob prob, void (XPRS_CC
       *mipthread)(XPRSprob master_prob, void* vContext, XPRSprob prob),
       void* data);
```

Arguments

`prob`	The current problem.
`mipthread`	The callback function to remove. If NULL then all variable branching callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all variable branching callbacks with the function pointer `mipthread` will be removed.

Related topics

`XPRSaddcbmipthread`

XPRSremovecbnewnode

Purpose

Removes a new-node callback function previously added by `XPRSaddcbnewnode`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSremovecbnewnode(XPRSprob prob, void (XPRS_CC
        *newnode)(XPRSprob my_prob, void* my_object, int parentnode, int
        newnode, int branch), void* data);
```

Arguments

`prob`	The current problem.
`newnode`	The callback function to remove. If NULL then all separation callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all separation callbacks with the function pointer `newnode` will be removed.

Related topics

`XPRSaddcbnewnode`

XPRSremovecbnodecutoff

Purpose

Removes a node-cutoff callback function previously added by `XPRSaddcbnodecutoff`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSremovecbnodecutoff(XPRSprob prob, void (XPRS_CC
      *nodecutoff)(XPRSprob my_prob, void *my_object, int nodnum), void*
      data);
```

Arguments

`prob`	The current problem.
`nodecutoff`	The callback function to remove. If NULL then all node-cutoff callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all node-cutoff callbacks with the function pointer `nodecutoff` will be removed.

Related topics

`XPRSaddcbnodecutoff`

XPRSremovecboptnode

Purpose

Removes a node-optimal callback function previously added by `XPRSaddcboptnode`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSremovecboptnode(XPRSprob prob, void (XPRS_CC
      *optnode)(XPRSprob my_prob, void *my_object, int *feas), void* data);
```

Arguments

`prob`	The current problem.
`optnode`	The callback function to remove. If NULL then all node-optimal callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all node-optimal callbacks with the function pointer `optnode` will be removed.

Related topics

`XPRSaddcboptnode`

XPRSremovecbpreintsol

Purpose

Removes a pre-integer solution callback function previously added by `XPRSaddcbpreintsol`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Solution

Synopsis

```
int XPRS_CC XPRSremovecbpreintsol(XPRSprob prob, void (XPRS_CC
       *preintsol)(XPRSprob my_prob, void *my_object, int soltype, int
       *ifreject, double *cutoff), void* data);
```

Arguments

`prob`	The current problem.
`preintsol`	The callback function to remove. If NULL then all user pre-integer solution callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all user pre-integer solution callbacks with the function pointer `preintsol` will be removed.

Related topics

`XPRSaddcbpreintsol`

XPRSremovecbprenode

Purpose

Removes a preprocess node callback function previously added by `XPRSaddcbprenode`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Branch and Bound Search

Synopsis

```
int XPRS_CC XPRSremovecbprenode(XPRSprob prob, void (XPRS_CC
      *prenode)(XPRSprob prob, void* my_object, int* nodinfeas), void*
      data);
```

Arguments

`prob`	The current problem.
`prenode`	The callback function to remove. If NULL then all preprocess node callback functions added with the given user-defined data value will be removed.
`data`	The data value that the callback was added with. If NULL, then the data value will not be checked and all preprocess node callbacks with the function pointer `prenode` will be removed.

Related topics

`XPRSaddcbprenode`

XPRSremovecbusersolnotify

Purpose

Removes a user solution notification callback previously added by `XPRSaddcbusersolnotify`. The specified callback function will no longer be called after it has been removed.

Topic areas

Callback, Solution

Synopsis

```
int XPRS_CC XPRSremovecbusersolnotify(XPRSprob prob, void (XPRS_CC
       *usersolnotify)(XPRSprob my_prob, void* my_object, const char*
       solname, int status), void* data);
```

Arguments

`prob`	The current problem.
`usersolnotify`	The callback function to remove. If `NULL` then all user solution notification callback functions added with the given user defined data value will be removed.
`data`	The data value that the callback was added with. If `NULL`, then the data value will not be checked and all integer solution callbacks with the function pointer `usersolnotify` will be removed.

Related topics

`XPRSaddcbusersolnotify`.

XPRSrepairinfeas

Purpose

Provides a simplified interface for `XPRSrepairweightedinfeas`.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSrepairinfeas (XPRSprob prob, int *p_status, char penalty,
        char phase2, char flags, double lepref, double gepref, double lbpref,
        double ubpref, double delta);
```

Arguments

`prob`	The current problem.
`p_status`	The status after the relaxation: `0` relaxed optimum found; `1` relaxed problem is infeasible; `2` relaxed problem is unbounded; `3` solution of the relaxed problem regarding the original objective is nonoptimal; `4` error (when return code is nonzero); `5` numerical instability; `6` analysis of an infeasible relaxation was performed, but the relaxation is feasible.
`penalty`	The type of penalties created from the preferences: `c` each penalty is the reciprocal of the preference (default); `s` the penalties are placed in the scaled problem.
`phase2`	Controls the second phase of optimization: `o` use the objective sense of the original problem (default); `x` maximize the relaxed problem using the original objective; `f` skip optimization regarding the original objective; `n` minimize the relaxed problem using the original objective; `i` if the relaxation is infeasible, generate an irreducible infeasible subset for the analys of the problem; `a` if the relaxation is infeasible, generate all irreducible infeasible subsets for the analys of the problem.
`flags`	Specifies flags to be passed to XPRSoptimize.
`lepref`	Preference for relaxing the less or equal side of row.
`gepref`	Preference for relaxing the greater or equal side of a row.
`lbpref`	Preferences for relaxing lower bounds.
`ubpref`	Preferences for relaxing upper bounds.
`delta`	The relaxation multiplier in the second phase -1. For console use `-d` value. A positive value means a relative relaxation by multiplying the first phase objective with (`delta-1`), while a negative value means an absolute relaxation, by adding `abs(delta)` to the first phase objective.

Related controls

Integer

`DEFAULTALG`	Forced algorithm selection (default for repairinfeas is primal).

Further information

1. This function currently does not apply to general nonlinear problems.
2. A row or bound is relaxed by introducing a new nonnegative variable that will contain the infeasibility of the row or bound. Suppose for example that row $a^T x = b$ is relaxed from below. Then a new variable (infeasibility breaker) $s \geqslant 0$ is added to the row, which becomes $a^T x + s = b$. Observe that $a^T x$ may now take smaller values than b. To minimize such violations, the weighted sum of these new variables is minimized.
3. A preference of 0 results in the row or bound not being relaxed.
4. A negative preference indicates that a quadratic penalty cost should be applied. This can specified on a per constraint side or bound basis.
5. Note that the set of preferences are scaling independent.
6. If a feasible solution is identified for the relaxed problem, with a sum of violations p, then the sum of violations is restricted to be no greater than `(1+delta)p`, and the problem is optimized with respect to the original objective function. A nonzero delta increases the freedom of the original problem.
7. Note that on some problems, slight modifications of delta may affect the value of the original objective drastically.
8. The default value for delta in the console is 0.001.
9. Note that because of their special associated modeling properties, binary and semi-continuous variables are not relaxed.
10. The default algorithm for the first phase is the simplex algorithm, since the primal problem can be efficiently warm started in case of the extended problem. These may be altered by setting the value of control `DEFAULTALG`.
11. If `penalty` is set such that each penalty is the reciprocal of the preference, the following rules are applied while introducing the auxiliary variables:

Preference	Affects	Relaxation	Cost if pref.>0	Cost if pref.<0
lepref	= rows	$a^T x$ - aux_var = b	1/lepref*aux_var	1/lepref*aux_var^2
lepref	<= rows	$a^T x$ - aux_var <= b	1/lepref*aux_var	1/lepref*aux_var^2
gepref	= rows	$a^T x$ + aux_var = b	1/gepref*aux_var	1/gepref*aux_var^2
gepref	>= rows	$a^T x$ + aux_var >= b	1/gepref*aux_var	1/gepref*aux_var^2
ubpref	upper bounds	x_i - aux_var <= u	1/ubpref*aux_var	1/ubpref*aux_var^2
lbpref	lower bounds	x_i + aux_var >= l	1/lbpref*aux_var	1/lbpref*aux_var^2

12. If an irreducible infeasible set (IIS) has been identified, the generated IIS(s) are accessible through the IIS retrieval functions, see `NUMIIS` and `XPRSgetiisdata`.

Related topics

`XPRSrepairweightedinfeas`, `XPRSoptimize`, 6.1.4.

XPRSrepairweightedinfeas

Purpose

By relaxing a set of selected constraints and bounds of an infeasible problem, it attempts to identify a 'solution' that violates the selected set of constraints and bounds minimally, while satisfying all other constraints and bounds. Among such solution candidates, it selects one that is optimal regarding the original objective function. For the console version, see `REPAIRINFEAS`.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSrepairweightedinfeas(XPRSprob prob, int * p_status, const
        double lepref[], const double gepref[], const double lbpref[], const
        double ubpref[], char phase2, double delta, const char *flags);
```

Arguments

`prob`	The current problem.
`p_status`	The status after the relaxation: `0` relaxed optimum found; `1` relaxed problem is infeasible; `2` relaxed problem is unbounded; `3` solution of the relaxed problem regarding the original objective is nonoptimal; `4` error (when return code is nonzero); `5` numerical instability; `6` analysis of an infeasible relaxation was performed, but the relaxation is feasible.
`lepref`	Array of size `ROWS` containing the preferences for relaxing the less or equal side of row.
`gepref`	Array of size `ROWS` containing the preferences for relaxing the greater or equal side of a row.
`lbpref`	Array of size `COLS` containing the preferences for relaxing lower bounds.
`ubpref`	Array of size `COLS` containing preferences for relaxing upper bounds.
`phase2`	Controls the second phase of optimization: `o` use the objective sense of the original problem (default); `x` maximize the relaxed problem using the original objective; `f` skip optimization regarding the original objective; `n` minimize the relaxed problem using the original objective; `i` if the relaxation is infeasible, generate an irreducible infeasible subset for the analys of the problem; `a` if the relaxation is infeasible, generate all irreducible infeasible subsets for the analys of the problem.
`delta`	The relaxation multiplier in the second phase -1.
`flags`	Specifies flags to be passed to XPRSoptimize.

Related controls

Double

`PENALTYVALUE`	The weighted sum of violations if a solution is identified to the relaxed problem.

Further information

1. This function applies to all problem types.
2. A row or bound is relaxed by introducing a new nonnegative variable that will contain the infeasibility of the row or bound. Suppose for example that row $a^T x = b$ is relaxed from below. Then a new variable ('infeasibility breaker') $s \geqslant 0$ is added to the row, which becomes $a^T x + s = b$. Observe that $a^T x$ may now take smaller values than b. To minimize such violations, the weighted sum of these new variables is minimized.
3. A preference of 0 results in the row or bound not being relaxed. The higher the preference, the more willing the modeller is to relax a given row or bound.
4. The weight of each infeasibility breaker in the objective minimizing the violations is 1/p, where `p` is the preference associated with the infeasibility breaker. Thus the higher the preference is, the lower a penalty is associated with the infeasibility breaker while minimizing the violations.
5. If a feasible solution is identified for the relaxed problem, with a sum of violations `p`, then the sum of violations is restricted to be no greater than `(1+delta)p`, and the problem is optimized with respect to the original objective function. A nonzero delta increases the freedom of the original problem.
6. Note that on some problems, slight modifications of delta may affect the value of the original objective drastically.
7. The default value for delta in the console is 0.001.
8. Note that because of their special associated modeling properties, binary and semi-continuous variables are not relaxed.
9. If `pflags` is set such that each penalty is the reciprocal of the preference, the following rules are applied while introducing the auxiliary variables:

Pref. array	Affects	Relaxation	Cost if pref.>0	Cost if pref.<0
lepref	= rows	$a^T x$ - aux_var = b	1/lrp*aux_var	1/lrp*aux_var^2
lepref	<= rows	$a^T x$ - aux_var <= b	1/lrp*aux_var	1/lrp*aux_var^2
gepref	= rows	$a^T x$ + aux_var = b	1/grp*aux_var	1/grp*aux_var^2
gepref	>= rows	$a^T x$ + aux_var >= b	1/grp*aux_var	1/grp*aux_var^2
ubpref	upper bounds	x_i - aux_var <= u	1/ubp*aux_var	1/ubp*aux_var^2
lbpref	lower bounds	x_i + aux_var >= l	1/lbp*aux_var	1/lbp*aux_var^2

10. If an irreducible infeasible set (IIS) has been identified, the generated IIS(s) are accessible through the IIS retrieval functions, see `NUMIIS` and `XPRSgetiisdata`.

Related topics

`XPRSrepairinfeas` (`REPAIRINFEAS`), `XPRSrepairweightedinfeasbounds`, `XPRSoptimize`, 6.1.4.

Example

See also example `repair.c`.

XPRSrepairweightedinfeasbounds — REPAIRINFEAS

Purpose

An extended version of `XPRSrepairweightedinfeas` that allows for bounding the level of relaxation allowed.

Topic area

Infeasibility

Synopsis

```
int XPRS_CC XPRSrepairweightedinfeasbounds(XPRSprob prob, int * p_status,
      const double lepref[], const double gepref[], const double lbpref[],
      const double ubpref[], const double lerelax[], const double
      gerelax[], const double lbrelax[], const double ubrelax[], char
      phase2, double delta, const char *flags);
REPAIRINFEAS -[pflags] -[oflags] -[gflags] -[lrp value] -[grp value] -[lbp
      value] -[ubp value] -[lrb value] -[grb value] -[lbb value] -[ubb
      value] -[d value] -[r]
```

Arguments

`prob`	The current problem.
`p_status`	The status after the relaxation: 0 relaxed optimum found; 1 relaxed problem is infeasible; 2 relaxed problem is unbounded; 3 solution of the relaxed problem regarding the original objective is nonoptimal; 4 error (when return code is nonzero); 5 numerical instability; 6 analysis of an infeasible relaxation was performed, but the relaxation is feasible.
`lepref`	Array of size `ROWS` containing the preferences for relaxing the less or equal side of row. For the console use `-lrp` value.
`gepref`	Array of size `ROWS` containing the preferences for relaxing the greater or equal side of a row. For the console use `-grp` value.
`lbpref`	Array of size `COLS` containing the preferences for relaxing lower bounds. For the console use `-lbp` value.
`ubpref`	Array of size `COLS` containing preferences for relaxing upper bounds. For the console use `-ubp` value.
`lerelax`	Array of size `ROWS` containing the upper bounds on the amount the less or equal side of a row can be relaxed. For the console use `-lrb` value.
`gerelax`	Array of size `ROWS` containing the upper bounds on the amount the greater or equal side of a row can be relaxed. For the console use `-grb` value.
`lbrelax`	Array of size `COLS` containing the upper bounds on the amount the lower bounds can be relaxed. For the console use `-lbb` value.
`ubrelax`	Array of size `COLS` containing the upper bounds on the amount the upper bounds can be relaxed. For the console use `-ubb` value.

`phase2`	Controls the second phase of optimization: `o` use the objective sense of the original problem (default); `x` maximize the relaxed problem using the original objective; `f` skip optimization regarding the original objective; `n` minimize the relaxed problem using the original objective; `i` if the relaxation is infeasible, generate an irreducible infeasible subset for the analys of the problem; `a` if the relaxation is infeasible, generate all irreducible infeasible subsets for the analys of the problem.
`delta`	The relaxation multiplier in the second phase -1.
`flags`	Specifies flags to be passed to XPRSoptimize.
`r`	If a summary of the violated variables and constraints should be printed after the relaxed solution is determined.
`y`	Call the global solver (console command only).

Example

```
READPROB MYPROB.LP
REPAIRINFEAS -a -d 0.002
```

This example reads in an infeasible problem and identifies a solution which violates the rows and bounds the least. It then fixes the measure of violations to this value multiplied by 1.002 to slightly increase the freedom in the problem and optimizes the relaxed problem using the original objective function.

Related controls

Double

`PENALTYVALUE` The weighted sum of violations if a solution is identified to the relaxed problem.

Further information

1. This function applies to all problem types.
2. The console command `REPAIRINFEAS` assumes that all preferences are 1 by default. Use the options `-lrp`, `-grp`, `-lbp` or `-ubp` to change them. The default limit on the maximum allowed relaxation per row or bound in plus infinity.
3. A row or bound is relaxed by introducing a new nonnegative variable that will contain the infeasibility of the row or bound. Suppose for example that row $a^T x = b$ is relaxed from below. Then a new variable ('infeasibility breaker') $s \geqslant 0$ is added to the row, which becomes $a^T x + s = b$. Observe that $a^T x$ may now take smaller values than b. To minimize such violations, the weighted sum of these new variables is minimized.
4. A preference of 0 results in the row or bound not being relaxed. The higher the preference, the more willing the modeller is to relax a given row or bound.
5. A negative preference indicates that a quadratic penalty cost should be applied. This can specified on a per constraint side or bound basis.
6. If a feasible solution is identified for the relaxed problem, with a sum of violations `p`, then the sum of violations is restricted to be no greater than `(1+delta)p`, and the problem is optimized with respect to the original objective function. A nonzero delta increases the freedom of the original problem.
7. Note that on some problems, slight modifications of delta may affect the value of the original objective drastically.
8. The default value for delta in the console is 0.001.
9. Note that because of their special associated modeling properties, binary and semi-continuous variables are not relaxed.
10. Given any row `j` with preferences `lrp=lepref[j]` and `grp=gepref[j]`, or variable `i` with bound preferences `ubp=ubpref[i]` and `lbp=lbpref[i]`, the following rules are applied while introducing the auxiliary variables:

Preference	Affects	Relaxation	Cost if pref.>0	Cost if pref.<0
lrp	= rows	a^Tx - aux_var = b	1/lrp*aux_var	1/lrp*aux_var^2
lrp	<= rows	a^Tx - aux_var <= b	1/lrp*aux_var	1/lrp*aux_var^2
grp	= rows	a^Tx + aux_var = b	1/grp*aux_var	1/grp*aux_var^2
grp	>= rows	a^Tx + aux_var >= b	1/grp*aux_var	1/grp*aux_var^2
ubp	upper bounds	x_i - aux_var <= u	1/ubp*aux_var	1/ubp*aux_var^2
lbp	lower bounds	x_i + aux_var >= l	1/lbp*aux_var	1/lbp*aux_var^2

11. Only positive bounds are applied; a zero or negative bound is ignored and the amount of relaxation allowed for the corresponding row or bound is not limited. The effect of a zero bound on a row or bound would be equivalent with not relaxing it, and can be achieved by setting its preference array value to zero instead, or not including it in the preference arrays.
12. If an irreducible infeasible set (IIS) has been identified, the generated IIS(s) are accessible through the IIS retrieval functions, see `NUMIIS` and `XPRSgetiisdata`.

Related topics

`XPRSrepairinfeas` (`REPAIRINFEAS`), `XPRSoptimize`, 6.1.4.

XPRSrestore RESTORE

Purpose

Restores the Optimizer's data structures from a file created by XPRSsave (SAVE). Optimization may then recommence from the point at which the file was created.

Topic areas

File IO, Save Restore

Synopsis

```
int XPRS_CC XPRSrestore(XPRSprob prob, const char *probname, const char
      *flags);
RESTORE [probname] [flags]
```

Arguments

prob	The current problem.
probname	A string of up to MAXPROBNAMELENGTH characters containing the problem name.
flags	Additional flags h Do not restore hardware information from the file; v use the provided filename verbatim, without appending the .svf extension.

Example 1 (Library)

```
XPRSrestore(prob,"","");
```

Example 2 (Console)

```
RESTORE
```

Further information

1. This routine restores the data structures from the file *problem_name*.svf that was created by a previous execution of XPRSsave (SAVE). Note that .svf files are particular to the release of the Optimizer used to create them. They can only be read using the same major release of the Optimizer that was used to create them.
2. (*Console*) The main use for XPRSsave (SAVE) and XPRSrestore (RESTORE) is to enable the user to interrupt a long optimization run using CTRL-C, and save the Optimizer status with the ability to restart it later from where it left off. It might also be used to save the optimal status of a problem when the user then intends to implement several uses of XPRSalter (ALTER) on the problem, re-optimizing each time from the saved status.

Related topics

XPRSalter (ALTER), XPRSsave (SAVE).

XPRSrhssa

Purpose

Returns upper and lower sensitivity ranges for specified right hand side (RHS) function coefficients. If the RHS coefficients are varied within these ranges the current basis remains optimal and the reduced costs remain valid.

Topic areas

Sensitivity Analysis, LP

Synopsis

```
int XPRS_CC XPRSrhssa(XPRSprob prob, int nrows, const int rowind[], double
       lower[], double upper[]);
```

Arguments

`prob`	The current problem.
`nrows`	The number of RHS coefficients for which sensitivity ranges are required.
`rowind`	Integer array of length `nrows` containing the indices of the rows whose RHS coefficients sensitivity ranges are required.
`lower`	Double array of length `nrows` where the RHS lower range values are to be returned.
`upper`	Double array of length `nrows` where the RHS upper range values are to be returned.

Example

Here we obtain the RHS ranges for the three columns: `2`, `6` and `8`:

```
rowind[0] = 2; rowind[1] = 8; rowind[2] = 6;
XPRSrhssa(prob,3,rowind,lower,upper);
```

After which lower and upper contain:

```
lower[0] = 5.0; upper[0] = 7.0;
lower[1] = 3.8; upper[1] = 5.2;
lower[2] = 5.7; upper[2] = 1e+20;
```

Meaning that the current basis remains optimal when $5.0 \leqslant \text{rhs}_2$, $3.8 \leqslant \text{rhs}_8 \leqslant 5.2$ and $5.7 \leqslant \text{rhs}_6$, rhs_i being the RHS coefficient of row `i`.

Further information

1. This function currently does not apply to general nonlinear problems.
2. `XPRSrhssa` can only be called when an optimal solution to the current LP has been found. It cannot be used when the problem is MIP presolved.

Related topics

`XPRSobjsa`, `XPRSbndsa`.

XPRSsave, XPRSsaveas — SAVE

Purpose

Saves the current data structures, i.e. matrices, control settings and problem attribute settings to file and terminates the run so that optimization can be resumed later.

Topic areas

File IO, Save Restore

Synopsis

```
int XPRS_CC XPRSsave(XPRSprob prob);
int XPRS_CC XPRSsaveas(XPRSprob prob, const char *sSaveFileName);
SAVE
```

Arguments

`prob`	The current problem.
`sSaveFileName`	The name of the file (without .svf) to save to.

Example 1 (Library)

```
XPRSsave(prob);
```

Example 2 (Console)

```
SAVE
```

Further information

1. The data structures are written to the file *problem_name*`.svf`. Optimization may recommence from the same point when the data structures are restored by a call to `XPRSrestore` (`RESTORE`). Note that the `.svf` files created are particular to the release of the Optimizer used to create them. They can only be read using the same release Optimizer as used to create them.
2. The function `XPRSsaveas` is equivalent to `XPRSsave` with the exception of allowing to adjust the name of the file created. The name of the file must not be greater than `MAXPROBNAMELENGTH`.

Related topics

`XPRSrestore` (`RESTORE`).

XPRSscale — SCALE

Purpose

Re-scales the current matrix.

Topic area

Numerics

Synopsis

```
int XPRS_CC XPRSscale(XPRSprob prob, const int rowscale[], const int
        colscale[]);
SCALE
```

Arguments

prob	The current problem.
rowscale	Integer array of size ROWS containing the powers of 2 with which to scale the rows, or NULL if not required.
colscale	Integer array of size COLS containing the powers of 2 with which to scale the columns, or NULL if not required.

Related controls

Integer

SCALING	Type of scaling.

Example 1 (Library)

```
XPRSreadprob(prob,"jovial","");
XPRSalter(prob,"serious");
XPRSscale(prob,NULL,NULL);
XPRSlpoptimize(prob,"");
```

This reads the MPS file `jovial.mat`, modifies it according to instructions in the file `serious.alt`, rescales the matrix and seeks the minimum objective value.

Example 2 (Console)

The equivalent set of commands for the Console user would be:

```
READPROB jovial
ALTER serious
SCALE
LPOPTIMIZE
```

Further information

1. This function currently does not apply to general nonlinear problems.
2. If `rowscale` and `colscale` are both non-`NULL` then they will be used to scale the matrix. Otherwise the matrix will be scaled according to the control `SCALING`. This routine may be useful when the current matrix has been modified by calls to routines such as `XPRSalter` (`ALTER`), `XPRSchgmcoef` and `XPRSaddrows`.
3. Any existing scaling will be replaced by the scaling specified by `rowscale` and `colscale`.
4. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSalter` (`ALTER`), `XPRSreadprob` (`READPROB`).

XPRSsetcheckedmode

Purpose

You can use this function to disable some of the checking and validation of function calls and function call parameters for calls to the Xpress Optimizer API. This checking is relatively lightweight but disabling it can improve performance in cases where non-intensive Xpress Optimizer functions are called repeatedly in a short space of time.

Please note: after disabling function call checking and validation, invalid usage of Xpress Optimizer functions may not be detected and may cause the Xpress Optimizer process to behave unexpectedly or crash. It is not recommended that you disable function call checking and validation during application development.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSsetcheckedmode(int checkedmode);
```

Argument

`checkedmode` Pass as 0 to disable much of the validation for all Xpress function calls from the current process. Pass 1 to re-enable validation. By default, validation is enabled.

Related topics

`XPRSgetcheckedmode`.

XPRSsetdblcontrol

Purpose

Sets the value of a given double control parameter.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSsetdblcontrol(XPRSprob prob, int control, double value);
```

Arguments

`prob`	The current problem.
`control`	Control parameter whose value is to be set. A full list of all controls may be found in 9, or from the list in the `xprs.h` header file.
`value`	Value to which the control parameter is to be set.

Related topics

`XPRSgetdblcontrol`, `XPRSsetintcontrol`, `XPRSsetstrcontrol`.

Example

See also examples `addmipsol.c`, `knapsack.c`, `roundint.c` .

XPRSsetdefaultcontrol SETDEFAULTCONTROL

Purpose

Sets a single control to its default value.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSsetdefaultcontrol(XPRSprob prob, int control);
SETDEFAULTCONTROL controlname
```

Arguments

`prob`	The current problem.
`control`	Integer, double or string control parameter whose default value is to be set.
`controlname`	Integer, double or string control parameter whose default value is to be set.

Example

The following turns off presolve to solve a problem, before resetting it to its default value and solving it again:

```
XPRSsetintcontrol(prob, XPRS_PRESOLVE, 0);
XPRSmipoptimize(prob, "");
XPRSwriteprtsol(prob);
XPRSsetdefaultcontrol(prob, XPRS_PRESOLVE);
XPRSmipoptimize(prob, "");
```

Further information

A full list of all controls may be found in Chapter 9, or from the list in the `xprs.h` header file.

Related topics

`XPRSsetdefaults`, `XPRSsetintcontrol`, `XPRSsetdblcontrol`, `XPRSsetstrcontrol`.

XPRSsetdefaults SETDEFAULTS

Purpose

Sets all controls to their default values. Must be called before the problem is read or loaded by `XPRSreadprob`, `XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSsetdefaults(XPRSprob prob);
SETDEFAULTS
```

Argument

`prob`	The current problem.

Example

The following turns off presolve to solve a problem, before resetting the control defaults, reading it and solving it again:

```
XPRSsetintcontrol(prob, XPRS_PRESOLVE, 0);
XPRSmipoptimize(prob, "");
XPRSwriteprtsol(prob);
XPRSsetdefaults(prob);
XPRSreadprob(prob);
XPRSmipoptimize(prob, "");
```

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSsetdefaultcontrol`, `XPRSsetintcontrol`, `XPRSsetdblcontrol`, `XPRSsetstrcontrol`.

XPRSsetindicators

Purpose

Specifies that a set of rows in the matrix will be treated as indicator constraints during a tree search. An indicator constraint is made of a `condition` and a `constraint`. The `condition` is of the type "bin = value", where `bin` is a binary variable and `value` is either 0 or 1. The `constraint` is any matrix row (may be linear, quadratic or general nonlinear). During tree search, a row configured as an indicator constraint is enforced only when condition holds, that is only if the indicator variable `bin` has the specified value. Note that every row may only get assigned a single indicator variable and term. If a row needs to be activated by multiple different terms, the row needs to be duplicated so that each term can be assigned to a distinct row. If the indicator variable should be changed, the old term needs to be deleted first (by calling `XPRSdelindicators` or by calling this function with a comps argument of 0) before assigning a new one.

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSsetindicators(XPRSprob prob, int nrows, const int rowind[],
       const int colind[], const int complement[]);
```

Arguments

`prob`	The current problem.
`nrows`	The number of indicator constraints.
`rowind`	Integer array of length `nrows` containing the indices of the rows that define the constraint part for the indicator constraints.
`colind`	Integer array of length `nrows` containing the column indices of the indicator variables.
`complement`	Integer array of length `nrows` with the complement flags: `0` not an indicator constraint (in this case the corresponding entry in the `colind` array is ignored); `1` for indicator constraints with condition "`bin = 1`"; `-1` for indicator constraints with condition "`bin = 0`".

Example

This sets the first two matrix rows as indicator rows in the MIP problem prob; the first row controlled by condition `x4=1` and the second row controlled by condition `x5=0` (assuming `x4` and `x5` correspond to columns indices 4 and 5).

```
int rowind[] = {0,1};
int colind[] = {4,5};
int complement[] = {1,-1};

...
XPRSsetindicators(prob,2,rowind,colind,complement);
XPRSmipoptimize(prob,"");
```

Further information

1. Indicator rows must be set up before solving the problem. Any indicator row will be removed from the matrix after presolve and added to a special pool. An indicator row will be added back into the active matrix only when its associated condition holds. An indicator variable can be used in multiple indicator rows and can also appear in normal rows and in the objective function.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSgetindicators`, `XPRSdelindicators`.

XPRSsetintcontrol, XPRSsetintcontrol64

Purpose
Sets the value of a given integer control parameter.

Topic area
Controls and Attributes

Synopsis
```
int XPRS_CC XPRSsetintcontrol(XPRSprob prob, int control, int value);

int XPRS_CC XPRSsetintcontrol64(XPRSprob prob, int control, XPRSint64
      value);
```

Arguments

prob	The current problem.
control	Control parameter whose value is to be set. A full list of all controls may be found in 9, or from the list in the xprs.h header file.
value	Value to which the control parameter is to be set.

Example
The following sets the control PRESOLVE to 0, turning off the presolve facility prior to optimization:
```
XPRSsetintcontrol(prob, XPRS_PRESOLVE, 0);
XPRSlpoptimize(prob, "");
```

See also examples addmipsol.c, ComplexUserFunctions.c, fixbv.c, knapsack.c, mipsolenum.c, mostviolated.c, Polygon_initialvalue.c, Polygon_textformula.c, Polygon_tokens.c, Polygon_userfunc.c, roundint.c, savesol.c, tableau.c, trimloss.c, tsp.c .

Further information
Some of the integer control parameters, such as SCALING, are bitmaps, with each bit controlling different behavior. Bit 0 has value 1, bit 1 has value 2, bit 2 has value 4, and so on.

Related topics
XPRSgetintcontrol, XPRSsetdblcontrol, XPRSsetstrcontrol.

XPRSsetlogfile — SETLOGFILE

Purpose
This directs all Optimizer output to a log file.

Topic areas
File IO, Logging

Synopsis
```
int XPRS_CC XPRSsetlogfile(XPRSprob prob, const char *filename);
SETLOGFILE filename
```

Arguments
`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name to which all logging output should be written. If set to `NULL`, redirection of the output will stop and all screen output will be turned back on (except for library users where screen output is always turned off).

Example
The following directs output to the file `logfile.log`:

```
XPRSinit(NULL);
XPRScreateprob(&prob);
XPRSsetlogfile(prob,"logfile.log");
```

See also examples `fixbv.c`, `globjpar.c`, `glrhspar.c`, `knapsack.c`, `loadlp.c`, `tableau.c`, `trimloss.c` .

Further information
1. It is recommended that a log file be set up for each problem being worked on, since it provides a means for obtaining any errors or warnings output by the Optimizer during the solution process.
2. If output is redirected with `XPRSsetlogfile` all screen output will be turned off.
3. Alternatively, an output callback can be defined using `XPRSaddcbmessage`, which will be called every time a line of text is output. To discard all output messages the `OUTPUTLOG` integer control can be set to `0`.

Related topics
`XPRSaddcbmessage`.

XPRSsetmessagestatus

Purpose

Manages suppression of messages.

Topic area

Misc

Synopsis

```
int XPRS_CC XPRSsetmessagestatus(XPRSprob prob, int msgcode, int status);
```

Arguments

`prob`	The problem for which message `msgcode` is to have its suppression status changed; pass `NULL` if the message should have the status apply globally to all problems.
`msgcode`	The id number of the message. Refer to the section 11 for a list of possible message numbers.
`status`	Non-zero if the message is not suppressed; `0` otherwise. If a value for `status` is not supplied in the command-line call then the console Optimizer prints the value of the suppression status to screen i.e., non-zero if the message is not suppressed; `0` otherwise.

Example

Attempting to optimize a problem that has no matrix loaded gives error 91. The following code uses `XPRSsetmessagestatus` to suppress the error message:

```
XPRScreateprob(&prob);
XPRSsetmessagestatus(prob,91,0);
XPRSlpoptimize(prob,"");
```

Further information

If a message is suppressed globally then the message can only be enabled for any problem once the global suppression is removed with a call to `XPRSsetmessagestatus` with `prob` passed as `NULL`.

Related topics

`XPRSgetmessagestatus`.

XPRSsetobjdblcontrol

Purpose

Sets the value of a given double control parameter associated with an objective. These parameters control how the objective is treated during multi-objective optimization.

Topic areas

Controls and Attributes, Multiobjective

Synopsis

```
int XPRS_CC XPRSsetobjdblcontrol(XPRSprob prob, int objidx, int control,
        double value);
```

Arguments

`prob`	The current problem.
`objidx`	Index of the objective to modify.
`control`	Control parameter whose value is to be modified. Must be one of: `XPRS_OBJECTIVE_WEIGHT` set the weight of the given objective; `XPRS_OBJECTIVE_ABSTOL` set the absolute tolerance of the given objective; `XPRS_OBJECTIVE_RELTOL` set the relative tolerance of the given objective; `XPRS_OBJECTIVE_RHS` set the constant term of the given objective.
`value`	Value to which the control parameter is to be set.

Example

The following sets the weight of the second objective to `0.5`:

```
XPRSsetobjdblcontrol(prob, 1, XPRS_OBJECTIVE_WEIGHT, 0.5);
```

See also example `goalprog.c`.

Further information

1. When `XPRS_OBJECTIVE_WEIGHT` is zero, the objective is disabled and will be ignored when solving. If all objectives with a given priority have zero weight, no solve will take place for that priority value.
2. This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSgetobjdblcontrol`, `XPRSgetobjintcontrol`, `XPRSsetobjintcontrol`.

XPRSsetobjintcontrol

Purpose

Sets the value of a given integer control parameter associated with an objective. These parameters control how the objective is treated during multi-objective optimization.

Topic areas

Controls and Attributes, Multiobjective

Synopsis

```
int XPRS_CC XPRSsetobjintcontrol(XPRSprob prob, int objidx, int control,
        int value);
```

Arguments

`prob`	The current problem.
`objidx`	Index of the objective to modify.
`control`	Control parameter whose value is to be modified. Must be one of: `XPRS_OBJECTIVE_PRIORITY` set the priority of the given objective.
`value`	Value to which the control parameter is to be set.

Example

The following sets the priority of the second objective to `1`:

```
XPRSsetobjintcontrol(prob, 1, XPRS_OBJECTIVE_PRIORITY, 1);
```

See also example `goalprog.c`.

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSgetobjintcontrol`, `XPRSgetobjdblcontrol`, `XPRSsetobjdblcontrol`.

XPRSsetprobname SETPROBNAME

Purpose

Sets the current default problem name. This command is rarely used.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSsetprobname(XPRSprob prob, const char *probname);
SETPROBNAME probname
```

Arguments

`prob`	The current problem.
`probname`	A string of up to `MAXPROBNAMELENGTH` characters containing the problem name.

Example

```
READPROB bob
LPOPTIMIZE
SETPROBNAME jim
READPROB
```

The above will read the problem `bob` and then read the problem `jim`.

See also example `addmipsol.c`.

Further information

This function cannot be called while the current problem is being solved or when it is in an interrupted state. Use `XPRSpostsolve` to restore the problem to a non-solving state.

Related topics

`XPRSreadprob` (`READPROB`), `XPRSgetprobname`, `MAXPROBNAMELENGTH`.

XPRSsetstrcontrol

Purpose

Used to set the value of a given string control parameter.

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XPRSsetstrcontrol(XPRSprob prob, int control, const char
      *value);
```

Arguments

`prob`	The current problem.
`control`	Control parameter whose value is to be set. A full list of all controls may be found in 9, or from the list in the `xprs.h` header file.
`value`	A string containing the value to which the control is to be set (plus a null terminator).

Example

The following sets the control `MPSOBJNAME` to `"Profit"`:

```
XPRSsetstrcontrol(prob, XPRS_MPSOBJNAME, "Profit");
```

Related topics

`XPRSgetstrcontrol`, `XPRSsetdblcontrol`, `XPRSsetintcontrol`.

XPRSsparsebtran

Purpose

Post-multiplies a (row) vector provided by the user by the inverse of the current matrix. Sparse version of `XPRSbtran`.

Topic areas

Linear Algebra, Simplex

Synopsis

```
int XPRS_CC XPRSsparsebtran(XPRSprob prob, double val[], int ind[], int
       *p_ncoefs);
```

Arguments

`prob`	The current problem.
`val`	Double array of length `ROWS` containing the values which are to be post-multiplied by the basis inverse. The transformed values are returned in this array.
`ind`	Integer array of indices identifying the non-zero entries of `val`. The indices of the non-zero entries of the transformed vector in `val` will be returned here. This array must be allocated to hold at least `ROWS` indices.
`p_ncoefs`	Memory location where the number of non-zero entries is given. The number of non-zero entries of the transformed vector will be returned here.

Related controls

Double

`ETATOL` Tolerance on eta elements.

Example

Get the (unscaled) tableau row `val` of the variable basic in constraint number `irow`, assuming that all arrays have been dimensioned.

```
/* Minimum size of arrays:
 * val: nrow + ncol;
 * start: 2;
 * ind: nrow;
 * colind, colcoef: ncol.
 */

/* set up the unit vector y to pick out row irow */
for(i = 0; i < nrow; i++) val[i] = 0.0;
val[irow] = 1.0;
ind[0] = irow;
nzcoefs = 1;

rc = XPRSsparsebtran(prob, val, ind, &nzcoefs);       /* y = e*B^{-1} */

/* Form z = y * A */
for (j = 0; j < ncol; j++) val[nrow + j] = 0.0;
for(i = 0; i < nzcoefs; i++) {
  r = ind[i];
  rc = XPRSgetrows(prob, start, colind, colcoef, ncol, &ncoefs, r, r);
  for (j = 0; j < ncoefs; j++)
    val[nrow+colind[j]] += val[r]*colcoef[j];
}
```

Further information

If the matrix is in a presolved state, the function will work with the basis for the presolved problem.

Related topics

`XPRSbtran`, `XPRSftran`, `XPRSsparseftran`.

XPRSsparseftran

Purpose

Pre-multiplies a (column) vector provided by the user by the inverse of the current matrix. Sparse version of `XPRSftran`.

Topic areas

Linear Algebra, Simplex

Synopsis

```
int XPRS_CC XPRSsparseftran(XPRSprob prob, double val[], int ind[], int
       *p_ncoefs);
```

Arguments

`prob`	The current problem.
`val`	Double array of length `ROWS` containing the values which are to be multiplied by the basis inverse. The transformed values are returned in this array.
`ind`	Integer array of indices identifying the non-zero entries of `val`. The indices of the non-zero entries of the transformed vector in `val` will be returned here. This array must be allocated to hold at least `ROWS` indices.
`p_ncoefs`	Memory location where the number of non-zero entries is given. The number of non-zero entries of the transformed vector will be returned here.

Related controls

Double

`ETATOL`	Tolerance on eta elements.

Example

To get the (unscaled) tableau column of structural variable number `jcol`, assuming that all arrays have been dimensioned, do the following:

```
/* Min size of arrays: start: 2; rowind, rowcoef, val & ind: nrow. */
/* Get column as loaded originally, in sparse format */
rc = XPRSgetcols(prob, start, rowind, rowcoef, nrow, &ncoefs,
                jcol, jcol);

/* Unpack into the zeroed array */
for(i = 0; i < nrow; i++) val[i] = 0.0;
nzcoefs = 0;
for(ielt = 0; ielt < ncoefs; ielt++)
  val[ind[nzcoefs++] = rowind[ielt]] = rowcoef[ielt];

rc = XPRSsparseftran(prob, val, ind, &nzcoefs);
```

Get the (unscaled) tableau column of the slack variable for row number `irow` , assuming that all arrays have been dimensioned.

```
/* Min size of arrays: val & ind: nrow */
/* Set up the original slack column in full format */
for(i = 0; i < nrow; i++) val[i] = 0.0;
val[irow] = 1.0;
ind[0] = irow;
nzcoefs = 1;

rc = XPRSsparseftran(prob, val, ind &nzcoefs);
```

Further information

If the matrix is in a presolved state, the function will work with the basis for the presolved problem.

Related topics

`XPRSftran`, `XPRSbtran`, `XPRSsparsebtran`.

STOP

Purpose

Terminates the Console Optimizer, returning an exit code to the operating system. This is useful for batch operations.

Topic area

Misc

Synopsis

`STOP`

Example

The following example inputs a matrix file, `lama.mat`, runs optimization on it and then exits:

```
READPROB lama
MIPOPTIMIZE
STOP
```

Further information

This command may be used to terminate the Optimizer as with the `QUIT` command. It sets an exit value which may be inspected by the host operating system or invoking program.

Related topics

`QUIT`.

XPRSstorecuts, XPRSstorecuts64

Purpose

Stores cuts into the cut pool, but does not apply them to the current node. These cuts must be explicitly loaded into the matrix using `XPRSloadcuts` before they become active.

Topic area

Cuts

Synopsis

```
int XPRS_CC XPRSstorecuts(XPRSprob prob, int ncuts, int nodups, const int
      cuttype[], const char rowtype[], const double rhs[], const int
      start[], XPRScut cutind[], const int colind[], const double
      cutcoef[]);

int XPRS_CC XPRSstorecuts64(XPRSprob prob, int ncuts, int nodups, const int
      cuttype[], const char rowtype[], const double rhs[], const XPRSint64
      start[], XPRScut cutind[], const int colind[], const double
      cutcoef[]);
```

Arguments

`prob`	The current problem.
`ncuts`	Number of cuts to add.
`nodups`	`0` do not exclude duplicates from the cut pool; `1` duplicates are to be excluded from the cut pool; `2` duplicates are to be excluded from the cut pool, ignoring cut type.
`cuttype`	Integer array of length `ncuts` containing the cut types. The cut types can be any integer and are used to identify the cuts.
`rowtype`	Character array of length `ncuts` containing the row types: `L` indicates a $\leqslant$ row; `E` indicates an = row; `G` indicates a $\geqslant$ row.
`rhs`	Double array of length `ncuts` containing the right hand side elements for the cuts.
`start`	Integer array containing offsets into the `colind` and `dmtval` arrays indicating the start of each cut. This array is of length `ncuts+1` with the last element `start[ncuts]` being where cut `ncuts+1` would start.
`cutind`	Array of length `ncuts` where the pointers to the cuts will be returned.
`colind`	Integer array of length `start[ncuts]` containing the column indices in the cuts.
`cutcoef`	Double array of length `start[ncuts]` containing the matrix values for the cuts.

Related controls

Double

`MATRIXTOL`	Tolerance on matrix elements.

Further information

1. XPRSstorecuts can be used to eliminate duplicate cuts. If the nodups parameter is set to 1, the cut pool will be checked for duplicate cuts with a cut type identical to the cuts being added. If a duplicate cut is found the new cut will only be added if its right hand side value makes the cut stronger. If the cut in the pool is weaker than the added cut it will be removed unless it has been applied to an active node of the tree. If nodups is set to 2 the same test is carried out on all cuts, ignoring the cut type.

2. XPRSstorecuts returns a list of the cuts added to the cut pool in the cutind array. If the cut is not added to the cut pool because a stronger cut exits a NULL will be returned. The cutind array can be passed directly to XPRSloadcuts to load the most recently stored cuts into the matrix.

3. The columns and elements of the cuts must be stored contiguously in the colind and dmtval arrays passed to XPRSstorecuts. The starting point of each cut must be stored in the start array. To determine the length of the final cut the start array must be of length ncuts+1 with the last element of this array containing where the cut ncuts+1 would start.

Related topics

XPRSloadcuts 5.9.

XPRSstrongbranch

Purpose

Performs strong branching iterations on all specified bound changes. For each candidate bound change, `XPRSstrongbranch` performs dual simplex iterations starting from the current optimal solution of the base LP, and returns both the status and objective value reached after these iterations.

Topic area

Branching

Synopsis

```
int XPRS_CC XPRSstrongbranch(XPRSprob prob, const int nbounds, const int
       colind[], const char bndtype[], const double bndval[], const int
       iterlim, double objval[], int status[]);
```

Arguments

`prob`	The current problem.
`nbounds`	Number of bound changes to try.
`colind`	Integer array of size `nbounds` containing the indices of the columns on which the bounds will change.
`bndtype`	Character array of length `nbounds` indicating the type of bound to change: `U` indicates change the upper bound; `L` indicates change the lower bound; `B` indicates change both bounds, i.e. fix the column.
`bndval`	Double array of length `nbounds` giving the new bound values.
`iterlim`	Maximum number of LP iterations to perform for each bound change.
`objval`	Objective value of each LP after performing the strong branching iterations.
`status`	Status of each LP after performing the strong branching iterations, as detailed for the `LPSTATUS` attribute.

Example

Suppose that the current LP relaxation has two integer columns (columns `0` and `1` which are fractionals at 0.3 and 1.5, respectively, and we want to perform strong branching in order to choose which to branch on. This could be done in the following way:

```
int colind[] = { 0, 0, 1, 1 };
char bndtype[] = "LULU";
double bndval[] = {1, 0, 2, 1};
double objval[4];
int status[4];
...
XPRSstrongbranch(prob, 4, colind, bndtype, bndval, 1000,
                 objval, status);
```

Further information

1. This function currently does not apply to general nonlinear problems.
2. Prior to calling XPRSstrongbranch, the current LP problem must have been solved to optimality and an optimal basis must be available.

XPRSstrongbranchcb

Purpose

Performs strong branching iterations on all specified bound changes. For each candidate bound change, `XPRSstrongbranchcb` performs dual simplex iterations starting from the current optimal solution of the base LP, and returns both the status and objective value reached after these iterations.

Topic area

Branching

Synopsis

```
int XPRS_CC XPRSstrongbranchcb(XPRSprob prob, const int nbounds, const int
       colind[], const char bndtype[], const double bndval[], const int
       iterlim, double objval[], int status[], int (XPRS_CC
       *callback)(XPRSprob cbprob, void* cbdata, int bndidx), void* data);
```

Arguments

`prob`	The current problem.
`nbounds`	Number of bound changes to try.
`colind`	Integer array of size `nbounds` containing the indices of the columns on which the bounds will change.
`bndtype`	Character array of length `nbounds` indicating the type of bound to change: `U` indicates change the upper bound; `L` indicates change the lower bound; `B` indicates change both bounds, i.e. fix the column.
`bndval`	Double array of length `nbounds` giving the new bound values.
`iterlim`	Maximum number of LP iterations to perform for each bound change.
`objval`	Objective value of each LP after performing the strong branching iterations.
`status`	Status of each LP after performing the strong branching iterations, as detailed for the `LPSTATUS` attribute.
`callback`	Function to be called after each strong branch has been reoptimized. This function returns an integer. Use `0` to indicate that everything went fine. Use a return value different from `0` to signal an error. This will terminate the function.
`cbprob`	The problem passed to the callback function, `callback`.
`cbdata`	The user-defined data passed as `data`.
`bndidx`	The index of bound for which `callback` is called.
`data`	User context to be provided for `callback`.

Further information

This function currently does not apply to general nonlinear problems.

Prior to calling `XPRSstrongbranchcb`, the current LP problem must have been solved to optimality and an optimal basis must be available.

`XPRSstrongbranchcb` is an extension to `XPRSstrongbranch`. If identical input arguments are provided both will return identical results, the difference being that for the case of `XPRSstrongbranchcb` the `callback` function is called at the end of each LP reoptimization.

For each branch optimized, the LP can be interrogated: the LP status of the branch is available through checking `LPSTATUS`, and the objective function value is available through `LPOBJVAL`. It is possible to access the full current LP solution by using `XPRSgetcallbacksolution`.

The return value of callback function `callback` is currently ignored.

TUNE

Purpose

This command can start a tuner session for the current problem. In this case, the tuner will solve the problem multiple times while evaluating a list of control settings and promising combinations of them. When finished, the tuner will select and set the best control setting on the problem. Note that the direction of optimization is given by `OBJSENSE`. This command can also handle the input and output of tuner method files.

Topic area

Tuner

Synopsis

`TUNE [-flags] [subcommand [filename]] [objsense]`

Arguments

`flags` Flags to pass to `TUNE`, which specify whether to tune the current problem as an LP or a MIP problem, and the algorithm for solving the LP problem or the initial LP relaxation of the MIP. The flags are optional. If the argument includes:

- `l` will tune the problem as an LP (mutually exclusive with flag `g`);
- `g` will tune the problem as a MIP (mutually exclusive with flag `l`);
- `x` will tune the problem as a Global Optimization problem with Xpress Global;
- `d` will use the dual simplex method;
- `p` will use the primal simplex method;
- `b` will use the barrier method;
- `n` will use the network simplex method.

`subcommand` Subcommand to pass to `TUNE` for handling tuner method files. It can be one of:

`pm / printmethod`	Print the tuner method on the console.
`wm / writemethod`	Write the tuner method to a file.
`rm / readmethod`	Read the tuner method from a file.
`probset`	Tune a set of problems.
`mipset`	Tune a set of MIP problems.
`lpset`	Tune a set of LP problems.

`filename` Tuner method file or problem set file. This is an optional argument of the subcommand.

`objsense` Objective sense to use when solving the problem. This is optional and is only permitted for subcommands which perform tuning.

Related controls

Integer

`TUNERHISTORY`	Whether to reuse and append to previous tuner result.
`TUNERMAXTIME`	Maximum total time allowed for the tuner.
`TUNERMETHOD`	Selects a factory tuner method.
`TUNERMODE`	Enable or disable the tuner.
`TUNEROUTPUT`	Whether to write tuner result and logs to file system.
`TUNERPERMUTE`	Number of permutations to solve with each control setting.
`TUNERTARGET`	Defines the criterion by which individual runs are compared.
`TUNERTHREADS`	Number of threads to be used by the tuner.

String

`TUNERMETHODFILE`	A file which contains a user-defined tuner method.
`TUNEROUTPUTPATH`	The root path for all tuner result output.
`TUNERSESSIONNAME`	When defined, will override the problem name within the tuner.

Example 1 (Console)

```
TUNE -l
```

This tunes the current problem as an LP problem.

Example 2 (Console)

```
TUNE pm
```

```
TUNE printmethod
```

Both commands print the tuner method to the console.

Example 3 (Console)

```
TUNE rm method
```

```
TUNE readmethod  method
```

Both commands read the tuner method from the `method.xtm` file.

Example 4 (Console)

```
TUNE wm method
```

```
TUNE writemethod method
```

Both commands write the tuner method to the `method.xtm` file.

Example 5 (Console)

```
TUNE probset problem.set
```

Tune a set of problems defined by the `problem.set` file.

Example 6 (Console)

```
TUNE lpset problem.set
```

Tune a set of LP problems defined by the `problem.set` file.

Further information

1. When both flags and subcommand are provided with the `TUNE` command, the subcommand will be ignored.
2. Please refer to Section 5.13 for a detailed guide of how to use the tuner.
3. Please refer to Section 5.13.8 for more information about tuning a set of problems.

XPRStune

Purpose

This function begins a tuner session for the current problem. The tuner will solve the problem multiple times while evaluating a list of control settings and promising combinations of them. When finished, the tuner will select and set the best control setting on the problem. Note that the direction of optimization is given by OBJSENSE.

Topic area

Tuner

Synopsis

```
int XPRS_CC XPRStune(XPRSprob prob, const char *flags);
```

Arguments

prob	The current problem.
flags	Flags to pass to XPRStune, which specify whether to tune the current problem as an LP or a MIP problem, and the algorithm for solving the LP problem or the initial LP relaxation of the MIP. The flags are optional. If the argument includes: l will tune the problem as an LP (mutually exclusive with flag g); g will tune the problem as a MIP (mutually exclusive with flag l); x will tune the problem as a Global Optimization problem with Xpress Global; d will use the dual simplex method; p will use the primal simplex method; b will use the barrier method; n will use the network simplex method.

Example

```
XPRStune(prob, "dp");
```

This tunes the current problem. The problem type is automatically determined. If it is an LP problem, it will be solved with a concurrent run of the dual and primal simplex method. If it is a MIP problem, the initial LP relaxation of the MIP will be solved with a concurrent run of primal and dual simplex.

Further information

1. Please refer to command TUNE for a list of related controls.
2. Please refer to Section 5.13 for a detailed guide of how to use the tuner.

Related topics

XPRStuneprobsetfile, XPRStunerreadmethod, XPRStunerwritemethod.

XPRStuneprobsetfile

Purpose

This function begins a tuner session for a set of problems. The tuner will solve the problems multiple times while evaluating a list of control settings and promising combinations of them. When finished, the tuner will select and set the best control setting on the problems.

Topic area

Tuner

Synopsis

```
int XPRS_CC XPRStuneprobsetfile(XPRSprob prob, const char *setfile, int
        ifmip, int sense);
```

Arguments

`prob`	The current problem.	
`setfile`	A plain text file which contains a list of problem filenames.	
`ifmip`	-1	to automatically determine whether to solve the problem set as LP or MIP;
	0	to force the tuner to tune the problem set as LP;
	1	to force the tuner to tune the problem set as MIP.
`sense`	0	to automatically determine the sense of each problem;
	1	to force the tuner to minimize each problem;
	-1	to force the tuner to maximize each problem.

Example

```
XPRStuneprobsetfile(prob, "problem.set", -1, 0);
```

This tunes the problems whose filenames are contained in `problem.set`. The problem types and objective senses are automatically determined.

Further information

Please refer to Section 5.13 for a detailed guide of how to use the tuner.

Related topics

`XPRStune`, `XPRStunerreadmethod`, `XPRStunerwritemethod`.

XPRStunerreadmethod

Purpose

This function loads a user defined tuner method from the given file.

Topic areas

Tuner, File IO

Synopsis

```
int XPRS_CC XPRStunerreadmethod(XPRSprob prob, const char* methodfile);
```

Arguments

`prob`	The current problem.
`methodfile`	The method file name, from which the tuner can load a user-defined tuner method.

Example

```
XPRStunerreadmethod(prob, "method.xtm");
```

This loads the tuner method from the `method.xtm` file.

Further information

Please refer to Section 5.13.2 for more information about the tuner method, and Appendix A.8 for the format of the tuner method file.

Related topics

`XPRStunerwritemethod`, `XPRStune`, `XPRStuneprobsetfile`.

XPRStunerwritemethod

Purpose
This function writes the current tuner method to a given file or prints it to the console.

Topic areas
Tuner, File IO

Synopsis
```
int XPRS_CC XPRStunerwritemethod(XPRSprob prob, const char* methodfile);
```

Arguments

`prob`	The current problem.
`methodfile`	The method file name, to which the tuner will write the current tuner method. If the input is `stdout` or `STDOUT`, then the tuner will print the method to the console instead.

Example 1 (Library)
```
XPRStunerwritemethod(prob, "method.xtm");
```

This writes the tuner method to the `method.xtm` file.

Example 2 (Library)
```
XPRStunerwritemethod(prob, "stdout");
```

This prints the tuner method to the console.

Further information
Please refer to Section 5.13.2 for more information about the tuner method, and Appendix A.8 for the format of the tuner method file.

Related topics
`XPRStunerreadmethod`, `XPRStune`, `XPRStuneprobsetfile`.

XPRSunloadprob

Purpose

This subroutine is deprecated and will be removed in a future release. Call `XPRSloadlp` with all `NULL` arguments to reset the problem to an empty problem.
Unloads and frees all memory associated with the current problem. It also invalidates the current problem (as opposed to reading in an empty problem).

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XPRSunloadprob(XPRSprob prob);
```

Argument

`prob`	The current problem.

Related topics

`XPRSreadprob`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadmiqp`, `XPRSloadqp`.

XPRSwritebasis WRITEBASIS

Purpose

Writes the current basis to a file for later input into the Optimizer.

Topic areas

File IO, LP

Synopsis

```
int XPRS_CC XPRSwritebasis(XPRSprob prob, const char *filename, const char
       *flags);
WRITEBASIS [-flags] [filename]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name from which the basis is to be written. If omitted, the default *problem_name* is used with a `.bss` extension.
`flags`	Flags to pass to `XPRSwritebasis` (`WRITEBASIS`): `i` output the internal presolved basis; `t` output a compact advanced form of the basis; `n` output basis file containing current solution values; `h` output values in single precision; `p` output values in full precision (obsolete as this is now default behavior); `v` use the provided filename verbatim, without appending the `.bss` extension; `z` compress the output file.

Example 1 (Library)

After an LP has been solved it may be desirable to save the basis for future input as an advanced starting point for other similar problems. This may save significant amounts of time if the LP is complex. The Optimizer input commands might then be:

```
XPRSreadprob(prob, "myprob", "");
XPRSlpoptimize(prob, "");
XPRSwritebasis(prob, "", "");
```

This reads in a matrix file, maximizes the LP and saves the basis. Loading a basis for a MIP problem can disable some MIP presolve operations which can result in a large increase in solution times so it is generally not recommended.

See also example `addmipsol.c`.

Example 2 (Console)

An equivalent set of commands to the above for console users would be:

```
READPROB
LPOPTIMIZE
WRITEBASIS
```

Further information

1. The `t` flag is only useful for later input to a similar problem using the `t` flag with `XPRSreadbasis` (`READBASIS`).
2. If the Newton barrier or hybrid gradient algorithms have been used for optimization then crossover must have been performed before there is a valid basis. This basis can then be used for restarting the simplex (primal or dual) or the hybrid gradient algorithms.
3. `XPRSwritebasis` (`WRITEBASIS`) will output the basis for the original problem even if the matrix has been presolved.

Related topics

`XPRSgetbasis`, `XPRSreadbasis` (`READBASIS`).

XPRSwritebinsol — WRITEBINSOL

Purpose
Writes the current MIP or LP solution to a binary solution file for later input into the Optimizer.

Topic areas
File IO, Solution

Synopsis
```
int XPRS_CC XPRSwritebinsol(XPRSprob prob, const char *filename, const char
       *flags);
WRITEBINSOL [-flags] [filename]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name to which the solution is to be written. If omitted, the default *problem_name* is used with a `.sol` extension.
`flags`	Flags to pass to `XPRSwritebinsol` (`WRITEBINSOL`): `m` output the MIP solution; `x` output the LP solution; `v` use the provided filename verbatim, without appending the `.sol` extension; `z` compress the output file.

Example 1 (Library)
After an LP has been solved or a MIP solution has been found the solution can be saved to file. If a MIP solution exists it will be written to file unless the -x flag is passed to `XPRSwritebinsol` (`WRITEBINSOL`) in which case the LP solution will be written. The Optimizer input commands might then be:

```
XPRSreadprob(prob, "myprob", "");
XPRSmipoptimize(prob, "");
XPRSwritebinsol(prob, "", "");
```

This reads in a matrix file, maximizes the MIP and saves the last found MIP solution.

Example 2 (Console)
An equivalent set of commands to the above for console users would be:

```
READPROB
MIPOPTIMIZE
WRITEBINSOL
```

Further information
This function currently does not apply to general nonlinear problems.

Related topics
`XPRSgetsolution`, `XPRSreadbinsol` (`READBINSOL`), `XPRSwritesol` (`WRITESOL`), `XPRSwriteprtsol` (`WRITEPRTSOL`).

XPRSwritedirs WRITEDIRS

Purpose

Writes the tree search directives from the current problem to a directives file.

Topic areas

File IO, Branching

Synopsis

```
int XPRS_CC XPRSwritedirs(XPRSprob prob, const char *filename);
WRITEDIRS [filename]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name to which the directives should be written. If omitted (or NULL), the default *problem_name* is used with a `.dir` extension.

Further information

If the problem has been presolved, only the directives for columns in the presolved problem will be written to file.

Related topics

`XPRSloaddirs`, A.5.

XPRSwriteprob WRITEPROB

Purpose

Writes the current problem to an MPS or LP file.

Topic area

File IO

Synopsis

```
int XPRS_CC XPRSwriteprob(XPRSprob prob, const char *filename, const char
      *flags);
WRITEPROB [-flags] [filename]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters to contain the file name to which the problem is to be written. If omitted, the default *problem_name* is used with a `.mps` extension, unless the `l` flag is used in which case the extension is `.lp`.
`flags`	Flags, which can be one or more of the following: `o` one element per line; `n` output the scaled problem; `s` scrambled vector names; `l` output in LP format; `p` output values in full precision (obsolete as this is now default behavior); `t` omit the Xpress header in LP or MPS format; `v` use the provided filename verbatim, without appending the `.mps` or `.lp` extension; `z` compress the output file.

Example

The following example outputs the current problem in LP format with scrambled vector names to the file *problem_name*`.lp`.

```
XPRSwriteprob(prob, "", "ls");
```

See also examples `els_managedcuts.c`, `els_usercuts.c`, `loadlp.c`, `repair.c`, `trimloss.c`, `tsp.c` .

Further information

1. If `XPRSloadlp`, `XPRSloadmip`, `XPRSloadmiqp` or `XPRSloadqp` is used to obtain a matrix then there is no association between the objective function and the `N` rows in the matrix and so a separate `N` row (called `__OBJ___`) is created when you do an `XPRSwriteprob` (`WRITEPROB`). Also if you do an `XPRSreadprob` (`READPROB`) and then change either the objective row or the `N` row in the matrix corresponding to the objective row, you lose the association between the two and the `__OBJ___` row is created when you do an `XPRSwriteprob` (`WRITEPROB`). To remove the objective row from the matrix when doing an `XPRSreadprob` (`READPROB`), set `KEEPNROWS` to `-1` before `XPRSreadprob` (`READPROB`).
2. **Warning:** If `XPRSreadprob` (`READPROB`) is used to input a problem, then the input file will be overwritten by `XPRSwriteprob` (`WRITEPROB`) if a new filename is not specified.
3. The Optimizer can write compressed matrix files in the following formats, as long as the command-line tool necessary to compress the file can be located in the path: `bzip2`, `xz`, `lzma`, `Z`, `zip`, `tar`, `tgz`. The Optimizer has built-in support for the `gz` format: no external tools are necessary to write gzipped matrix files.

Related topics

`XPRSreadprob` (`READPROB`).

XPRSwriteprtsol WRITEPRTSOL

Purpose
Writes the current solution to a fixed format ASCII file, *problem_name* `.prt`.

Topic areas
File IO, Solution

Synopsis
```
int XPRS_CC XPRSwriteprtsol(XPRSprob prob, const char *filename, const char
      *flags);
WRITEPRTSOL [filename] [-flags]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name to which the solution is to be written. If omitted, the default *problem_name* will be used. The extension `.prt` will be appended.
`flags`	Flags for `XPRSwriteprtsol` (`WRITEPRTSOL`) are: `x` write the LP solution instead of the current MIP solution; `v` use the provided filename verbatim, without appending the `.prt` extension; `z` write a compressed output file; `s` include classical sensitivity analysis.

Related controls

Integer

`MAXPAGELINES`	Number of lines between page breaks.

Double

`OUTPUTTOL`	Tolerance on print values.

Example 1 (Library)
This example shows the standard use of this function, outputting the solution to file immediately following optimization:
```
XPRSreadprob(prob, "myprob", "");
XPRSlpoptimize(prob, "");
XPRSwriteprtsol(prob, "", "");
```

Example 2 (Console)
```
READPROB
LPOPTIMIZE
PRINTSOL
```

are the equivalent set of commands for Console users who wish to view the output directly on screen.

Further information
1. This function currently does not apply to general nonlinear problems.
2. (*Console*) There is an equivalent command `PRINTSOL` which outputs the same information to the screen. The format is the same as that output to file by `XPRSwriteprtsol` (`WRITEPRTSOL`), except that the user is permitted to enter a response after each screen if further output is required.
3. The fixed width ASCII format created by this command is not as readily useful as that produced by `XPRSwritesol` (`WRITESOL`). The main purpose of `XPRSwriteprtsol` (`WRITEPRTSOL`) is to create a file that can be sent directly to a printer. The format of this fixed format ASCII file is described in Appendix A.
4. To create a prt file for a previously saved solution, the solution must first be loaded with the `XPRSreadbinsol` (`READBINSOL`) function.

Related topics
`XPRSgetsolution`, `XPRSreadbinsol` `XPRSwritebinsol`, `XPRSwritesol`, A.4.

XPRSwriteslxsol WRITESLXSOL

Purpose

Creates an ASCII solution file (.slx) using a similar format to MPS files. These files can be read back into the Optimizer using the XPRSreadslxsol function.

Topic areas

File IO, Solution

Synopsis

```
int XPRS_CC XPRSwriteslxsol(XPRSprob prob, const char *filename, const char
       *flags);
WRITESLXSOL -[flags] [filename]
```

Arguments

prob	The current problem.
filename	A string of up to MAXPROBNAMELENGTH characters containing the file name to which the solution is to be written. If omitted, the default *problem_name* is used with a .slx extension.
flags	Flags to pass to XPRSwriteslxsol (WRITESLXSOL): l write the LP solution in case of a MIP problem; m write the MIP solution; p use full precision for numerical values (obsolete as this is now default behavior); s including slack variables; d LP solution only: including dual variables; r LP solution only: including reduced cost; v use the provided filename verbatim, without appending the .slx extension; z compress the output file.

Example 1 (Library)

```
XPRSwriteslxsol(prob,"lpsolution","");
```

This saves the MIP solution if the problem contains MIP entities, or otherwise saves the LP (barrier in case of quadratic problems) solution of the problem.

Example 2 (Console)

```
WRITESLXSOL lpsolution
```

Which is equivalent to the library example above.

Related topics

XPRSreadslxsol (READSLXSOL, XPRSwriteprtsol (WRITEPRTSOL), XPRSwritebinsol (WRITEBINSOL), XPRSreadbinsol (READBINSOL).

XPRSwritesol — WRITESOL

Purpose

Writes the current solution to a CSV format ASCII file, *problem_name*`.asc` (and `.hdr`).

Topic areas

File IO, Solution

Synopsis

```
int XPRS_CC XPRSwritesol(XPRSprob prob, const char *filename, const char
      *flags);
WRITESOL [filename] [-flags]
```

Arguments

`prob`	The current problem.
`filename`	A string of up to `MAXPROBNAMELENGTH` characters containing the file name to which the solution is to be written. If omitted, the default *problem_name* will be used. The extensions `.hdr` and `.asc` will be appended.
`flags`	Flags to control which optional fields are output: `s` sequence number; `n` name; `t` type; `b` basis status; `a` activity; `c` cost (columns), slack (rows); `l` lower bound; `u` upper bound; `d` dj (column; reduced costs), dual value (rows; shadow prices); `r` right hand side (rows). If no flags are specified, all fields are output. Additional flags: `p` outputs in full precision; `q` only outputs vectors with nonzero optimum value; `x` output the current LP solution instead of the MIP solution; `z` compress the output file.

Related controls

Double

`OUTPUTTOL`	Tolerance on print values.

String

`OUTPUTMASK`	Mask to restrict the row and column names output to file.

Example 1 (Library)

In this example the basis status is output (along with the sequence number) following optimization:

```
XPRSreadprob(prob, "richard", "");
XPRSlpoptimize(prob, "");
XPRSwritesol(prob, "", "sb");
```

See also example `trimloss.c`.

Example 2 (Console)

Suppose we wish to produce files containing

- the names and values of variables starting with the letter `X` which are nonzero and
- the names, values and right hand sides of constraints starting with `CO2`.

The Optimizer commands necessary to do this are:

```
OUTPUTMASK = "X???????"
WRITESOL XVALS -naq
OUTPUTMASK = "CO2?????"
WRITESOL CO2 -nar
```

Further information

1. This function currently does not apply to general nonlinear problems.
2. The command produces two readable files: `filename.hdr` (the solution header file) and `filename.asc` (the CSV foramt solution file). The header file contains summary information, all in one line. The ASCII file contains one line of information for each row and column in the problem. Any fields appearing in the `.asc` file will be in the order the flags are described above. The order that the flags are specified by the user is irrelevant.
3. Additionally, the mask control `OUTPUTMASK` may be used to control which names are reported to the ASCII file. Only vectors whose names match `OUTPUTMASK` are output. `OUTPUTMASK` is set by default to "`????????`", so that all vectors are output.

Related topics

`XPRSgetsolution`, `XPRSwriteprtsol` (`WRITEPRTSOL`).

Chapter 9

Control Parameters

Various controls exist within the Optimizer to govern the solution procedure and the form of output. The majority of these take integer values and act as switches between various types of behavior. The tolerances on values are double precision, and there are a few controls which are character strings, setting names to structures. Any of these may be altered by the user to enhance performance of the Optimizer. However, it should be noted that the default values provided have been found to work well in practice over a range of problems and caution should be exercised if they are changed.

9.1 Retrieving and Changing Control Values

Console Xpress users may obtain control values by issuing the control name at the Optimizer prompt, >, and hitting the RETURN key. Controls may be set using the assignment syntax:

control_name = *new_value*

where *new_value* is an integer value, double or string as appropriate. For character strings, the name must be enclosed in single quotes and all eight characters must be given.

Users of the FICO Xpress Libraries are provided with the following set of functions for setting and obtaining control values:

XPRSgetintcontrol	XPRSgetdblcontrol	XPRSgetstrcontrol
XPRSsetintcontrol	XPRSsetdblcontrol	XPRSsetstrcontrol

It is an important point that the controls as listed in this chapter *must* be prefixed with XPRS_ to be used with the FICO Xpress Libraries and failure to do so will result in an error. An example of their usage is as follows:

```
XPRSgetintcontrol(prob, XPRS_PRESOLVE, &presolve);
printf("The value of PRESOLVE is %d\n", presolve);
XPRSsetintcontrol(prob, XPRS_PRESOLVE, 1-presolve);
printf("The value of PRESOLVE is now %d\n", 1-presolve);
```

9.2 Bit-vector controls

Some Optimizer controls are bit-vector controls, a special type of integer control where each of the 32 bits represents an independent feature toggle–either on or off. The bits are numbered from 0 to 30 (the last bit holds the sign), with each bit k corresponding to the integer value 2^k. For example, enabling bits 1, 4, and 8 in a control means setting it to $2^1 + 2^4 + 2^8 = 2 + 16 + 256 = 274$. Consequently, setting a bit-vector control to zero turns all bits off. Due to the way negative numbers are represented in computer systems (using the so-called "Two's Complement"), the value -1 corresponds to all bits being set to one, effectively enabling all features. For this reason, many (but not all) bit-vector controls in the

Optimizer default to -1. Examples of bit-vector controls include `PRESOLVEOPS`, `SCALING`, and `HEURSEARCHROOTSELECT`. For a list of all bit-vector controls, see 7.2.

Similarly, some functions, such as `XPRSdelcuts`, accept bit-vectors as input, allowing multiple yes/no characteristics to be encoded into a single integer.

ALGAFTERCROSSOVER

Description The algorithm to be used for the final clean up step after the crossover.

Type Integer

Category Control

Topic areas LP, Barrier

Values

1	Automatically determined.
2	Dual simplex.
3	Primal simplex.
4	Concurrent.

Default value 1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`) (when the barrier is used), `XPRScrossoverlpsol`.

ALGAFTERNETWORK

Description The algorithm to be used for the clean up step after the network simplex solver.

Type Integer

Category Control

Topic areas LP, Simplex

Values

-1	Automatically determined.
2	Dual simplex.
3	Primal simplex.

Default value -1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`) (when network simplex is used).

ALTERNATIVEREDCOSTS

Description Controls aggressiveness of searching for alternative reduced cost

Type Integer

Category Control

Topic area Propagation

Values	-1	The solver decides if searching for alternative reduced cost is beneficial or not. This is the default setting.
	0	Searching for alternative reduced cost is disabled.
	1	Searching for alternative reduced cost is enabled.

Default value -1

Affects routines `XPRSoptimize` (`OPTIMIZE`).

AUTOCUTTING

Description Should the Optimizer automatically decide whether to generate cutting planes at local nodes in the tree or not? If the `CUTFREQ` control is set, no automatic selection will be made and local cutting will be enabled.

Type Integer

Category Control

Topic area Cuts

Values	-1	Automatic.
	0	Disabled.
	1	Enabled.

Default value -1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `CUTDEPTH`, `CUTFREQ`.

AUTOSCALING

Description Whether the Optimizer should automatically select between different scaling algorithms. If the `SCALING` control is set, no automatic scaling will be applied.

Type Integer

Category Control

Topic areas Numerics, Problem Transformation

Values	-1	Automatic.
	0	Disabled.
	1	Cautious strategy. Non-standard scaling will only be selected if it appears to be clearly superior.
	2	Moderate strategy.
	3	Aggressive strategy. Standard scaling will only be selected if it appears to be clearly superior.

Default value -1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `SCALING`, `OBJSCALEFACTOR`.

AUTOPERTURB

Description Simplex: This indicates whether automatic perturbation is performed. If this is set to `1`, the problem will be perturbed whenever the simplex method encounters an excessive number of degenerate pivot steps, thus preventing the Optimizer being hindered by degeneracies.

Type Integer

Category Control

Topic area LP

Values

`0`	No perturbation performed.
`1`	Automatic perturbation is performed.

Default value `1`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BACKTRACK

Description Branch and Bound: Specifies how to select the next node to work on when a full backtrack is performed.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

`-1`	Automatically determined.
`1`	Unused.
`2`	Select the node with the best estimated solution.
`3`	Select the node with the best bound on the solution.
`4`	Select the deepest node in the search tree (equivalent to depth-first search).
`5`	Select the highest node in the search tree (equivalent to breadth-first search).
`6`	Select the earliest node created.
`7`	Select the latest node created.
`8`	Select a node randomly.
`9`	Select the node whose LP relaxation contains the fewest number of infeasible MIP entities.
`10`	Combination of 2 and 9.
`11`	Combination of 2 and 4.
`12`	Combination of 3 and 4.

Default value `3`

Note Note When two nodes are rated the same according to the `BACKTRACK` selection, a secondary rating is performed using the method set by `BACKTRACKTIE`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `BACKTRACKTIE`.

BACKTRACKTIE

Description Branch and Bound: Specifies how to break ties when selecting the next node to work on when a full backtrack is performed. The options are the same as for the BACKTRACK control.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

-1	Default selection.
1	Unused.
2	Select the node with the best estimated solution.
3	Select the node with the best bound on the solution.
4	Select the deepest node in the search tree (equivalent to depth-first search).
5	Select the highest node in the search tree (equivalent to breadth-first search).
6	Select the earliest node created.
7	Select the latest node created.
8	Select a node randomly.
9	Select the node whose LP relaxation contains the fewest number of infeasible MIP entities.
10	Combination of 2 and 9.
11	Combination of 2 and 4.
12	Combination of 3 and 4.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also BACKTRACK.

BACKGROUNDMAXTHREADS

Description Limit the number of threads to use in background jobs (for example in parallel to the root cut loop).

Type Integer

Category Control

Topic areas Root, Parallel

Default value -1, let Xpress decide.

Notes If this is set to -1 (the default) then Xpress decides automatically how many threads to run. Otherwise this control limits the number of threads that Xpress will use for jobs in the background (for example in parallel to the root cut loop). Set this to 0 (zero) in order to disable jobs in the background.

The MAXMIPTASKS does not apply to this control. If you want to have the same solution paths between hardware with different CPU counts or between runs with different thread counts then you should set this control to the same non-negative value in all runs.

Affects routines XPRSmipoptimize (MIPOPTIMIZE), XPRSoptimize (OPTIMIZE).

See also BACKGROUNDSELECT, HEURSEARCHBACKGROUNDSELECT.

BACKGROUNDSELECT

Description Bit-vector control (see Section 9.2) to select which tasks to run in background jobs (for example in parallel to the root cut loop).

Type Integer

Category Control

Topic areas Root, Heuristics, Bit-vector

Values

Bit	Meaning
0	Feasibility jump heuristic.
1	Fast branch-and-bound heuristic.
2	Same as bit 1 but with some additional heuristics enabled.
3	Fix-propagate-repair heuristic.

Default value -1

Notes If this is set to -1 (the default) then Xpress decides automatically which tasks to run. Otherwise the value is interpreted as a bitset that indicates which heuristics to run. 0 (zero) means not to run any tasks in the background.

Background jobs that run large neighborhood search heuristics are controlled by `HEURSEARCHBACKGROUNDSELECT`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSoptimize` (`OPTIMIZE`).

See also `HEURSEARCHBACKGROUNDSELECT`, `BACKGROUNDMAXTHREADS`.

BARALG

Description This control determines which barrier algorithm is used to solve the problem. Notably, this is also the control to enable the primal-dual hybrid gradient algorithm.

Type Integer

Category Control

Topic areas LP, Barrier, Primal Dual Hybrid Gradient

Values

-1	Determined automatically.
0	Unused.
1	Use the infeasible-start barrier algorithm.
2	Use the homogeneous self-dual barrier algorithm.
3	Start with 2 and optionally switch to 1 during the execution.
4	Use the hybrid gradient algorithm.

Default value -1

Note The automatic setting uses 1 for LP and QP problems and 3 for QCQP problems. Usually the detection of primal or dual infeasibility is more robust with settings 2 or 3, therefore, it is advantageous to use one of these values if the model is presumably infeasible. Setting 4, the hybrid gradient algorithm, is suitable for very large or very degenerate problems.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSminim` (`MINIM`), `XPRSmaxim` (`MAXIM`).

BARCORES

Description	If set to a positive integer it determines the number of physical CPU cores assumed to be present in the system by the barrier and hybrid gradient algorithms. If the value is set to the default value (`-1`), Xpress will automatically detect the number of cores.
Type	Integer
Category	Control
Topic areas	Barrier, Parallel, Primal Dual Hybrid Gradient
Default value	`-1`(automatically detected)
Note	The control is provided for cross-hardware reproducibility purposes. The count does not include logical cores created by Hyper-Threading.
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`BARTHREADS`.

BARCRASH

Description	Newton barrier and hybrid gradient: This determines the type of crash used for the crossover. During the crash procedure, an initial basis is determined which attempts to speed up the crossover. A good choice at this stage will significantly reduce the number of iterations required to crossover to an optimal solution. The possible values increase proportionally to their time-consumption.	
Type	Integer	
Category	Control	
Topic areas	Barrier, Primal Dual Hybrid Gradient	
Values	`0`	Turns off all crash procedures.
	`1-6`	Available strategies with 1 being conservative and 6 being aggressive.
Default value	`4`	
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).	

BARDUALSTOP

Description	Newton barrier and hybrid gradient: This is a convergence parameter, representing the tolerance for dual infeasibilities. If the difference between the constraints and their bounds in the dual problem falls below this tolerance in absolute value, optimization will stop and the current solution will be returned.	
Type	Double	
Category	Control	
Topic areas	Barrier, Primal Dual Hybrid Gradient, Tolerances	
Values	`0`	The default value is determined automatically based on the problem size, structure and algorithm choice.
	`>=0`	The tolerance for dual infeasibilities.
Default value	`0`	
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).	

BARFAILITERLIMIT

Description Newton barrier: The maximum number of consecutive iterations that fail to improve the solution in the barrier algorithm.

Type Integer

Category Control

Topic areas Barrier, Limits

Values

0	Determined automatically
>0	Maximum number of consecutive barrier iterations allowed without progress.

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARFREESCALE

Description Defines how the barrier algorithm scales free variables.

Type Double

Category Control

Topic area Barrier

Default value 1e-6

Note When using smaller values the barrier algorithm scales free variables more aggressively which can improve performance but may impact numerical stability.

Affects routines XPRSlpoptimize (LPOPTIMIZE).

See also SCALING.

BARGAPSTOP

Description Newton barrier and hybrid gradient: This is a convergence parameter, representing the tolerance for the relative duality gap. When the difference between the primal and dual objective function values falls below this tolerance, the Optimizer determines that the optimal solution has been found.

Type Double

Category Control

Topic areas Barrier, Primal Dual Hybrid Gradient

Values

0	The default value is determined automatically based on the problem size, structure and algorithm choice.
>=0	The tolerance for the relative duality gap.

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARGAPTARGET

Description Newton barrier: The target tolerance for the relative duality gap. The barrier algorithm will keep iterating until either BARGAPTARGET is satisfied or until no further improvements are possible. In the latter case, if BARGAPSTOP is satisfied, it will declare the problem optimal.

Type Double

Category Control

Topic area Barrier

Values

0	The default value is determined automatically based on the problem size, structure and algorithm choice.
>=0	The target tolerance for the relative duality gap.

Default value 0

Note When a solution returned by the barrier algorithm has not converged tightly enough for an application, for example if the dual solution is not accurate enough or crossover is taking too long, setter BARGAPTARGET to a small value often resolves the problem, without the risk of the solve failing due to a complementarity level not being numerically achievable. Typical suggested values can be between 1^{-10} and 1^{-18}.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARHGEXTRAPOLATE

Description Extrapolation parameter for the hybrid gradient algorithm. Although theory suggests that a value of 1 is best, slightly smaller values perform better in general.

Type Double

Category Control

Topic area Primal Dual Hybrid Gradient

Default value 0.15

Affects routines XPRSlpoptimize (LPOPTIMIZE).

BARHGMAXRESTARTS

Description The maximum number of restarts in the hybrid gradient algorithm. Restarts play the role of iterations in the hybrid gradient algorithm. A log line is printed at every restart, unless BAROUTPUT is set to 0.

Type Integer

Category Control

Topic area Primal Dual Hybrid Gradient

Default value 1250

Affects routines XPRSlpoptimize (LPOPTIMIZE).

BARHGOPS

Description Bit-vector control (see Section 9.2) options for the hybrid gradient algorithm. Bits 1, 2 and 3 control which norms of the coefficient matrix are used for solution normalization. The normalization factor is the maximum of the selected norms. By default, or if all three bits are set to 0, the infinity norm is used. The omega parameter referenced in bits 4, 5 and 6 is a measure of the relative magnitudes of the objective and the right-hand side.

Type Integer

Category Control

Topic areas Primal Dual Hybrid Gradient, Bit-vector

Values

Bit	Meaning
0	Use an asymmetric average for the primal averaging.
1	Use the 1-norm of the coefficient matrix in normalizing the initial solution.
2	Use the 2-norm of the coefficient matrix in normalizing the initial solution.
3	Use the infinity norm of the coefficient matrix in normalizing the initial solution.
4	Unused.
5	Contract omega towards 1 if the infeasibility is small enough.
6	Omega is based on the infeasibility.

Default value `8`, only the infinity norm is used for normalization, the other options are all off.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`).

BARINDEFLIMIT

Description Newton Barrier. This limits the number of consecutive indefinite barrier iterations that will be performed. The optimizer will try to minimize (resp. maximize) a QP problem even if the Q matrix is not positive (resp. negative) semi-definite. However, the optimizer may detect that the Q matrix is indefinite and this can result in the optimizer not converging. This control specifies how many indefinite iterations may occur before the optimizer stops and reports that the problem is indefinite. It is usual to specify a value greater than one, and only stop after a series of indefinite matrices, as the problem may be found to be indefinite incorrectly on a few iterations for numerical reasons.

Type Integer

Category Control

Topic areas Barrier, Limits

Default value `15`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARITERATIVE

Description The maximum number of barrier iterations in which an iterative solver is used instead of the Cholesky decomposition.

Type Integer

Category Control

Topic area Barrier

Values

-2	Automatically determined.
-1	Turn iterative solver off.
0	Use iterative solver for the starting point computation.
n>0	Try to apply iterative solver for the first n barrier iterations.

Default value -2

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARITERLIMIT

Description Newton barrier: The maximum number of iterations. While the simplex method usually performs a number of iterations which is proportional to the number of constraints (rows) in a problem, the barrier method standardly finds the optimal solution to a given accuracy after a number of iterations which is independent of the problem size. The penalty is rather that the time for each iteration increases with the size of the problem. BARITERLIMIT specifies the maximum number of iterations which will be carried out by the barrier.

Type Integer

Category Control

Topic areas Barrier, Limits

Default value 500

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARKERNEL

Description Newton barrier: Defines how centrality is weighted in the barrier algorithm.

Type Double

Category Control

Topic area Barrier

Values

>=+1.0	Increases the emphasis on centrality when larger value is set.
<=-1.0	Selects a value adaptively in every iteration from [+1, -BARKERNEL].

Default value 0.0

Note Increasing this parameter may increase the number of iterations, therefore the recommended range is [1,2] and [-2,-1].

Affects routines XPRSlpoptimize (LPOPTIMIZE).

See also BARALG.

BARLARGEBOUND

Description Threshold for the barrier to handle large bounds.

Type Double

Category Control

Topic area Barrier

Default value `0`

BAROBJPERTURB

Description Defines how the barrier perturbs the objective.

Type Double

Category Control

Topic area Barrier

Values

`>0`	Let the optimizer decide if the objective is perturbed or not and use the parameter value as the scale of the perturbation.
`0`	Turn off objective perturbation.
`<0`	Always perturb the objective by the absolute value of the parameter.

Default value `1e-6`

Note The perturbation scale should be set carefully with consideration to the optimality tolerance. The parameter affects only the barrier solve.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `BAROBJSCALE`.

BAROBJSCALE

Description Defines how the barrier scales the objective.

Type Double

Category Control

Topic areas Barrier, Numerics

Values

`-1`	Let the optimizer decide.
`0`	Scale by geometric mean.
`>=0`	Scale such that the largest objective coefficient's largest element does not exceed this number. In quadratic problems, the quadratic diagonal is used as reference valuses instead of the linear objective.

Default value `-1`

Note The scaling perfromed by the barrier is applied on top of any other scaling in the problem and only affects the barrier solve.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `SCALING`.

BARORDER

Description	Newton barrier: This controls the Cholesky factorization in the Newton-Barrier.
Type	Integer
Category	Control
Topic area	Barrier
Values	`0` Choose automatically. `1` Minimum degree method. This selects diagonal elements with the smallest number of nonzeros in their rows or columns. `2` Minimum local fill method. This considers the adjacency graph of nonzeros in the matrix and seeks to eliminate nodes that minimize the creation of new edges. `3` Nested dissection method. This considers the adjacency graph and recursively seeks to separate it into non-adjacent pieces.
Default value	`0`
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARORDERTHREADS

Description	If set to a positive integer it determines the number of concurrent threads for the sparse matrix ordering algorithm in the Newton-barrier method.
Type	Integer
Category	Control
Topic areas	Barrier, Parallel
Values	`0` The default value is determined automatically based on the problem size, structure and algorithm choice. `>=0` The number of concurrent threads for the sparse matrix ordering algorithm in the Newton-barrier method.
Default value	`0`
Note	Larger values than `BARCORES` will be automatically reduced to the value of `BARCORES`.
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`BARORDER`, `BARCORES`.

BAROUTPUT

Description	Newton barrier and hybrid gradient: This specifies the level of solution output provided. Output is provided either after each iteration of the algorithm, or else can be turned off completely by this parameter.
Type	Integer
Category	Control
Topic areas	Barrier, Logging, Primal Dual Hybrid Gradient

Values	0	No output.
	1	At each iteration.

Default value 1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARPERTURB

Description Newton barrier: In numerically challenging cases it is often advantageous to apply perturbations on the KKT system to improve its numerical properties. `BARPERTURB` controlls how much perturbation is allowed during the barrier iterations. By default no perturbation is allowed. Set this parameter with care as larger perturbations may lead to less efficient iterates and the best settings are problem-dependent.

Type Double

Category Control

Topic areas Barrier, Numerics

Default value 0

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARPRESOLVEOPS

Description Newton barrier: This bit-vector (see Section 9.2) controls the Newton-Barrier specific presolve operations.

Type Integer

Category Control

Topic areas Barrier, Presolve, Bit-vector

Values	Bit	Meaning
	0	Use standard presolve.
	1	Extra effort is spent in barrier specific presolve.
	2	Do full matrix eliminations (reduce matrix size).

Default value 0

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARPRIMALSTOP

Description Newton barrier and hybrid gradient: This is a convergence parameter, indicating the tolerance for primal infeasibilities. If the difference between the constraints and their bounds in the primal problem falls below this tolerance in absolute value, the Optimizer will terminate and return the current solution.

Type Double

Category Control

Topic areas Barrier, Tolerances, Primal Dual Hybrid Gradient

Values

0	The default value is determined automatically based on the problem size, structure and algorithm choice.
>=0	The tolerance for primal infeasibilities.

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARREFITER

Description Newton barrier: After terminating the barrier algorithm, further refinement steps can be performed. Such refinement steps are especially helpful if the solution is near to the optimum and can improve primal feasibility and decrease the complementarity gap. It is also often advantageous for the crossover algorithm. BARREFITER specifies the maximum number of such refinement iterations.

Type Integer

Category Control

Topic area Barrier

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARREGULARIZE

Description This bit-vector control (see Section 9.2) determines how the barrier algorithm applies regularization on the KKT system.

Type Integer

Category Control

Topic areas Barrier, Bit-vector

Values

Bit	Meaning
0	Standard regularization is turned on/off.
1	Reduced regularization is turned on/off. This option reduces the perturbation effect of the standard regularization.
2	Forces to keep dependent rows in the KKT system.
3	Forces to preserve degenerate rows in the KKT system.
4	Restrict regularization to infeasible iterates.
5	Disable iterative regularizations.
6	Apply iterative regularization more often.

Default value -1

Note The parameter is a bit set but value -1 (the default value) is treated in a special way: if the parameter is set to -1 then the solver will automatically select the bits it deems most useful.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRSminim (MINIM), XPRSmaxim (MAXIM).

BARRHSSCALE

Description Defines how the barrier scales the right hand side.

Type Double

Category Control

Topic areas Barrier, Numerics

Values

`-1`	Let the optimizer decide.
`0`	Scale by geometric mean.
`>=0`	Scale such that the largest right hand side coefficient's largest element does not exceed this number.

Default value `-1`

Note The scaling perfromed by the barrier is applied on top of any other scaling in the problem and only affects the barrier solve.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `SCALING`.

BARSOLUTION

Description This determines whether the barrier has to decide which is the best solution found or return the solution computed by the last barrier iteration.

Type Integer

Category Control

Topic area Barrier

Values

`-1`	(callback only: do not save current soulution as the best one).
`0`	return the best solution found (in callback: let the barrier decide the current solution is the best or not).
`1`	return the last barrier iteration (in callback: save current solution as the best solution so far).

Default value `0`

Affects routines The barrier algorithm.

BARSTART

Description Controls the computation of the starting point and warm-starting for the Newton barrier and the hybrid gradient algorithms.

Type Integer

Category Control

Topic areas Barrier, Primal Dual Hybrid Gradient

Values	-1	Uses the existing solution for warm-start if one is available.
	0	Warm-start is disabled; the starting point is determined automatically from the next three options.
	1	Uses simple heuristics to compute the starting point based on the magnitudes of the matrix entries.
	2	Uses the pseudoinverse of the constraint matrix to determine primal and dual initial solutions. Less sensitive to scaling and numerically more robust, but in several case less efficient than 1.
	3	Uses the unit starting point for the homogeneous self-dual barrier algorithm.

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARSTARTWEIGHT

Description Newton barrier: This sets a weight for the warm-start point when warm-start is set for the barrier algorithm. Using larger weight gives more emphasis for the supplied starting point.

Type Double

Category Control

Topic area Barrier

Default value 0.85

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also BARSTART.

BARSTEPSTOP

Description Newton barrier: A convergence parameter, representing the minimal step size. On each iteration of the barrier algorithm, a step is taken along a computed search direction. If that step size is smaller than BARSTEPSTOP, the Optimizer will terminate and return the current solution.

Type Double

Category Control

Topic area Barrier

Default value 1.0E-16

Note If the barrier method is making small improvements on BARGAPSTOP on later iterations, it may be better to set this value higher, to return a solution after a close approximation to the optimum has been found.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

BARTHREADS

Description If set to a positive integer it determines the number of threads implemented to run the Newton-barrier and hybrid gradient algorithms. If the value is set to the default value (`-1`), the `THREADS` control will determine the number of threads used.

Type Integer

Category Control

Topic areas Barrier, Parallel, Primal Dual Hybrid Gradient

Default value `-1`(determined by the `THREADS` control)

Note There is a practical upper limit of 50 on the number of parallel threads the optimizer will create.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `BARCORES`, `MIPTHREADS`, `CONCURRENTTHREADS`, `THREADS`.

BIGM

Description The infeasibility penalty used if the "Big M" method is implemented.

Type Double

Category Control

Topic areas LP, Numerics

Default value Dependent on the matrix characteristics.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BIGMMETHOD

Description Simplex: This specifies whether to use the "Big M" method, or the standard phase I (achieving feasibility) and phase II (achieving optimality). In the "Big M" method, the objective coefficients of the variables are considered during the feasibility phase, possibly leading to an initial feasible basis which is closer to optimal. The side-effects involve possible round-off errors due to the presence of the "Big M" factor in the problem.

Type Integer

Category Control

Topic areas LP, Numerics

Values

`0`	For phase I / phase II.
`1`	If "Big M" method to be used.

Default value `1`

Note Reset by `XPRSreadprob` (`READPROB`), `XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp` and `XPRSloadqp`.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BRANCHCHOICE

Description Once a MIP entity has been selected for branching, this control determines which of the branches is solved first.

Type Integer

Category Control

Topic area Branching

Values

0	Minimum estimate branch first.
1	Maximum estimate branch first.
2	If an incumbent solution exists, solve the branch satisfied by that solution first. Otherwise solve the minimum estimate branch first (option 0).
3	Solve first the branch that forces the value of the branching variable to move farther away from the value it had at the root node. If the branching entity is not a simple variable, solve the minimum estimate branch first (option 0).

Default value 0

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

BRANCHDISJ

Description Branch and Bound: Determines whether the optimizer should attempt to branch on general split disjunctions during the branch and bound search.

Type Integer

Category Control

Topic area Branching

Values

-1	Automatic selection of the strategy.
0	Disabled.
1	Cautious strategy. Disjunctive branches will be created only for general integers with a wide range.
2	Moderate strategy.
3	Aggressive strategy. Disjunctive branches will be created for both binaries and integers.

Default value -1

Notes Note Split disjunctions are a special form of disjunctions that can be written as $\sum_j m_j x_j \leqslant m_0 \vee \sum_j m_j x_j \geqslant m_0 + 1$

The split disjunctions created by the optimizer will use a combination of binary or integer variables x_j, with integer coefficients m_j.

Split disjunctions for branching will always be created with a default priority value of 400 instead of the default value of 500 for regular entity branches.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

BRANCHSTRUCTURAL

Description Branch and Bound: Determines whether the optimizer should search for special structure in the problem to branch on during the branch and bound search.

Type Integer

Category Control

Topic area Branching

Values

`-1`	Automatically determined.
`0`	Disabled.
`1`	Enabled.

Default value -1

Notes Structural branches will often involve branching on more than a single MIP entity at a time. As a result of a structural branch, a parent node could therefore end up with more than two child nodes, unlike the standard single entity branches.

Structural branches will always be created with a default priority value of 400 instead of the default value of 500 for regular entity branches.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

BREADTHFIRST

Description The number of nodes to include in the best-first search before switching to the local first search (`NODESELECTION` = `4`).

Type Integer

Category Control

Topic area Branch and Bound Search

Default value `11`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

CACHESIZE

Description *This parameter is deprecated and will be removed in a future release.*
Newton Barrier: L2 or L3 (see notes) cache size in kB (kilobytes) of the CPU. On Intel (or compatible) platforms a value of `-1` may be used to determine the cache size automatically. If the CPU model is new then the cache size may not be correctly detected by an older release of the software.

Type Integer

Category Control

Topic area System

Default value -1

Notes

1. Specifying the correct cache size can give a significant performance advantage with the Newton barrier algorithm. If the size is unknown, it is better to specify a smaller size.
2. If the size cannot be determined automatically, a default size of `128kB` is assumed.
3. Where present, the L3 cache size should be chosen rather than the L2 cache size.
4. For multi-core CPUs, the cache is shared between a subset of the cores. The Optimizer will divide the `CACHESIZE` value by the number of cores sharing the cache if >1 Barrier threads are running.
5. Where the CPU is described as having multiple caches ie. `2x6M` then the correct cache size to use is `6M` not `12M`.
6. Examples:

Intel Core 2 Duo E6400 (2M Cache, 2.13GHz)	`CACHESIZE=2048`
Intel Xeon x5570 (8M Cache, 2.93GHz)	`CACHESIZE=8192`
Intel Core 2 QX6700 (2x4M Cache, 2.93 GHz)	`CACHESIZE=4096`

7. If in doubt, please contact Support for advice.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `L1CACHE`.

CALLBACKFROMMASTERTHREAD

Description Branch and Bound: specifies whether the MIP callbacks should only be called on the master thread.

Type Integer

Category Control

Topic area Callback

Values

`0`	Invoke callbacks on worker threads during parallel MIP;
`1`	Only ever invoke a callback on the thread that called `XPRSmipoptimize`.

Default value `0`

Affects routines `XPRSmipoptimize`.

CALLBACKCHECKTIMEDELAY

Description Minimum delay in milliseconds between two consecutive executions of the CHECKTIME callback in the same solution process

Type Integer

Category Control

Topic area Callback

Values

`0`	Callback delay is disabled - the callback is executed every time;
`n>0`	Callback invocation is suppressed if less than n milliseconds have passed since the last invocation.

Default value 0

Note The CHECKTIME callback may be invoked very often and can lead to an overhead. This control enables users to specify a minimum delay between two consecutive calls, such that the overhead becomes controllable. Note that system time is used to measure the time between two consecutive calls. Therefore, a positive delay value may result in a nondeterministic sequence of callback invocations. Note that the counter is reset at every invocation of a solution process via `XPRSlpoptimize`, `XPRSmipoptimize`, etc.

Affects routines `XPRSaddcbchecktime`.

CALLBACKCHECKTIMEWORKDELAY

Description Minimum delay in work units between two consecutive executions of the CHECKTIME callback in the same solution process

Type Double

Category Control

Topic areas Callback, Determinism

Values

`0`	Callback delay is disabled - the callback is executed every time;
`n>0`	Callback invocation if less than n work units, which may be a fraction of a work unit, have passed since the last invocation.

Default value 0

Note The CHECKTIME callback may be invoked very often and can lead to an overhead. This control enables users to specify a minimum delay between two consecutive calls, such that the overhead becomes controllable. Note that `WORK` is accumulated deterministically during the search in contrast to `TIME`. Therefore, the callback invocation sequence with a positive `CALLBACKCHECKTIMEWORKDELAY` is reproducible across multiple invocations of the same search (same user controls, same problem) even on different machines. Note that the internal counter is reset at every invocation of a solution process via `XPRSlpoptimize`, `XPRSmipoptimize`, etc.

Note The disadvantage of setting a delay in work units is that work increases less uniformly than time.

Note While the related control `CALLBACKCHECKTIMEDELAY` specifies a delay in milliseconds, this double control allows delay specification of an arbitrary fraction of a work unit.

Affects routines `XPRSaddcbchecktime`.

See also `WORKLIMIT`, `CALLBACKCHECKTIMEDELAY`

CHECKINPUTDATA

Description Check input arrays for bad data.

Type Integer

Category Control

Topic areas Problem Creation, File IO

Values

`0`	Do not check.
`1`	Check input arrays.

Default value `1 (on)`

Note If this control is set (the default value) then all floating point input arrays are checked for NAN. An error is raised if a NAN is found in any such array. For certain arrays the Optimizer will also check whether they contain INF and raise an error. For example, matrix coefficients cannot be infinities but bounds can be.

CHOLESKYALG

Description Newton barrier: type of Cholesky factorization used; bit-vector-control (see Section 9.2).

Type Integer

Category Control

Topic areas Barrier, Bit-vector

Values

Bit	Meaning
0	matrix blocking: 0: automatic setting; 1: manual setting.
1	if manual selection of matrix blocking: 0: multi-pass; 1: single-pass.
2	nonseparable QP relaxation: 0: off; 1: on.
3	corrector weight: 0: automatic setting; 1: manual setting.
4	if manual selection of corrector weight: 0: off; 1: on.
5	refinement: 0: automatic setting; 1: manual setting.
6	preconditioned conjugate gradient method (PCGM): 0: PCGM off; 1: PCGM on.
7	Preconditioned quasi minimal residual (QMR) to refine solution: 0: QMR off; 1: QMR on.
8	Perform refinement on the augmented system 0: off; 1: on.
9	Force highest accuracy in refinement 0: off; 1: on.

Default value `-1 (automatic)`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

CHOLESKYTOL

Description Newton barrier: The tolerance for pivot elements in the Cholesky decomposition of the normal equations coefficient matrix, computed at each iteration of the barrier algorithm. If the absolute value of the pivot element is less than or equal to `CHOLESKYTOL`, it merits special treatment in the Cholesky decomposition process.

Type Double

Category Control

Topic area Barrier

Default value `1.0E-15`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

CLAMPING

Description This bit-vector control (see Section 9.2) allows for the adjustment of returned solution values such that they are always within bounds.

Type Integer

Category Control

Topic areas Solution, Numerics, Bit-vector

Values

Bit	Meaning
`-1`	Determined automatically.
`0`	Adjust primal solution to always be within primal bounds. Slacks if provided will be adjusted accordingly.
`1`	Adjust primal slack values to always be within constraint bounds.
`2`	Adjust dual solution to always be within the dual bounds implied by the slacks. Reduced costs, if provided, will be adjusted accordingly.
`3`	Adjust reduced costs to always be within dual bounds implied by the primal solution.

Default value `0`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSminim` (`MINIM`), `XPRSmaxim` (`MAXIM`).

COMPUTE

Description Controls whether the next solve is performed directly or on an Insight Compute Interface.

Type Integer

Category Control

Topic area Compute Interface

Values

`0`	Solve locally.
`1`	Solve using an Insight Compute Interface.

Default value Depends on environment

Note	This control can only be used if Xpress is not in compute mode due to the environment variable of the same name `XPRS_COMPUTE` being set before initialising Xpress. The first time the control is set to 1, Xpress reads the `XPRS_COMPUTE_URL` environment variable for the server address. Once enabled, the same restrictions apply to any solves that normally applies to solves (e.g. a restricted set of callbacks are supported). Enabling compute mode this way does not release the license held by Xpress.
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSiisfirst` (`IIS`), `XPRSrepairinfeas` (`REPAIRINFEAS`).

COMPUTEEXECSERVICE

Description	Selects the Insight execution service that will be used for solving remote optimizations.
Type	String
Category	Control
Topic area	Compute Interface
Default value	Empty string
Note	Set to the name of the execution service you want to use.
Note	When an empty string, the value from the configuration file will be used, or Insight server's default execution service if neither is populated.
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

COMPUTEJOBPRIORITY

Description	Selects the priority that will be used for remote optimization jobs.
Type	Integer
Category	Control
Topic area	Compute Interface
Default value	0
Note	Possible values are from -100 to 100.
Note	A job of priority N will always get scheduled before a job of priority N-1, provided there is sufficient capacity.
Note	If the value is 0, the value from the configuration file will be used, or priority 0 if neither is specified.
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

COMPUTELOG

Description	Controls how the run log is fetched when a solve is performed on an Insight Compute Interface.
Type	Integer
Category	Control
Topic area	Compute Interface
Values	`0` Run log will not be fetched `1` Run log will be fetched in real-time `2` Run log will be fetched at the end of the solve `3` Run log will be fetched at the end of the solve if the solve fails with an error
Default value	1
Note	As run log lines are fetched from the remote solve, they will be sent to the message callback or be displayed by the solver in the console. The run log will not be available if the control is assigned any value other than '1' (real-time) and the solve is terminated from outside a callback, such as by calling `XPRSinterrupt` or hitting `CTRL+C` in the solver console.
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSiisfirst` (`IIS`), `XPRSrepairinfeas` (`REPAIRINFEAS`), `XPRSaddcbmessage`.

CONCURRENTTHREADS

Description	Determines the number of threads used by the concurrent solver.
Type	Integer
Category	Control
Topic areas	LP, Parallel
Values	`-1` Determined automatically `>0` Number of threads to use.
Default value	`-1`
Note	Please refer to section 5.10.1 for a detailed description of the concurrent solver.
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`).
See also	`DETERMINISTIC`, `DUALTHREADS`, `BARTHREADS`, `THREADS`.

CONFLICTCUTS

Description	Branch and Bound: Specifies how cautious or aggressive the optimizer should be when searching for and applying conflict cuts. Conflict cuts are in-tree cuts derived from nodes found to be infeasible or cut off, which can be used to cut off other branches of the search tree.
Type	Integer
Category	Control
Topic area	Branch and Bound Search

Values	`-1`	Automatic.
	`0`	Disable conflict cuts.
	`1`	Cautious application of conflict cuts.
	`2`	Medium application of conflict cuts.
	`3`	Aggressive application of conflict cuts.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`)

CORESPERCPU

Description Used to override the detected value of the number of cores on a CPU. The cache size (either detected or specified via the `CACHESIZE` control) used in Barrier methods will be divided by this amount, and this scaled-down value will be the amount of cache allocated to each Barrier thread

Type Integer

Category Control

Topic area System

Default value `-1`

Affects routines `CACHESIZE`

COVERCUTS

Description Branch and Bound: The number of rounds of lifted cover inequalities at the root node. A lifted cover inequality is an additional constraint that can be particularly effective at reducing the size of the feasible region without removing potential integral solutions. The process of generating these can be carried out a number of times, further reducing the feasible region, albeit incurring a time penalty. There is usually a good payoff from generating these at the root node, since these inequalities then apply to every subsequent node in the tree search.

Type Integer

Category Control

Topic areas Cuts, Root

Default value `-1` — determined automatically.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

CPIALPHA

Description decay term for confined primal integral computation.

Type Double

Category Control

Topic area Misc

Note This control represents the exponential decay term for computing the OBSERVEDPRIMALINTEGRAL. The smaller it is, the more emphasis is put on the early part of the search. A value of 0 corresponds to computing a regular primal integral without exponential decay. For details see Berthold and Csizmadia: *The confined primal integral*, Mathematical Programming volume 188(2), pp. 523-537, 2021.

Default value 0

See also CPISCALEFACTOR, OBSERVEDPRIMALINTEGRAL.

CPUPLATFORM

Description Newton Barrier: Selects the AMD, Intel x86 or ARM vectorization instruction set that Barrier should run optimized code for. On AMD / Intel x86 platforms the SSE2, AVX and AVX2 instruction sets are supported while on ARM platforms the NEON architecture extension can be activated.

Type Integer

Category Control

Topic area System

Values

-2	Highest supported [Generic, SSE2, AVX or AVX2].
-1	Highest supported solve path consistent code [Generic, SSE2 or AVX].
0	Use generic code compatible with all CPUs.
1	Use SSE2 / NEON optimized code.
2	Use AVX optimized code.
3	Use AVX2 optimized code.

Default value -2, using AVX2 instructions if supported by the CPU

Note Generic code, SSE2 and AVX optimized code will all result in the same solution path, regardless of the operating system. The solution paths of AVX2 runs are also the same, regardless of the operating system. Using AVX2 or NEON might result in a solution path that is different from the AVX path.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

CPUTIME

Description How time should be measured when timings are reported in the log and when checking against time limits

Type Integer

Category Control

Topic area System

Values

-1	Disable the timer.
0	Use elapsed time.
1	Use process time.

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

CRASH

Description	Simplex: This determines the type of crash used when the algorithm begins. During the crash procedure, an initial basis is determined which is as close to feasibility and triangularity as possible. A good choice at this stage will significantly reduce the number of iterations required to find an optimal solution. The possible values increase proportionally to their time-consumption.
Type	Integer
Category	Control
Topic areas	LP, Simplex, Bit-vector
Note	For the primal simplex algorithm the following choices are available:

Values

`0`	Turns off all crash procedures.
`1`	For singletons only (one pass).
`2`	For singletons only (multi pass).
`3`	Multiple passes through the matrix considering slacks.
`4`	Multiple ($\leqslant$ `10`) passes through the matrix but only doing slacks at the very end.
`n>10`	As for value `4` but performing at most n - `10` passes.

Note For the dual simplex algorithm, the crash control is interpreted as a bit-vector (see Section 9.2) for adjusting the behavior of the procedure:

Values

Bit	Meaning
`0`	Perform standard crash.
`1`	Perform additional numerical checks during crash.
`2`	Extend the set of column candidates for crash.
`3`	Extend the set of row candidates for crash.
`4`	Force crash, i.e., consider all suitable columns/rows as candidates for crash.

Default value	`2`
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

CROSSOVER

Description	Newton barrier and hybrid gradient: This control determines whether the barrier method will cross over to the simplex method when at optimal solution has been found, to provide an end basis (see `XPRSgetbasis`, `XPRSwritebasis`) and advanced sensitivity analysis information (see `XPRSobjsa`, `XPRSrhssa`, `XPRSbndsa`).
Type	Integer
Category	Control
Topic areas	Barrier, Primal Dual Hybrid Gradient

Values

`-1`	Determined automatically.
`0`	No crossover.
`1`	Primal crossover first.
`2`	Dual crossover first.

Default value `-1`

Note The full primal and dual solution is available whether or not crossover is used. The crossover must not be disabled if the barrier is used to reoptimize nodes of a MIP.
By default crossover will not be performed on QP and MIQP problems, and after the hybrid gradient algorithm.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

CROSSOVERACCURACYTOL

Description Newton barrier: This control determines how crossover adjusts the default relative pivot tolerance. When re-inversion is necessary, crossover will compare the recalculated working basic solution with the assumed ones just before re-inversion took place. If the error is above this threshold, crossover will adjust the relative pivot tolerance to address the build-up of numerical inaccuracies.

Type Double

Category Control

Topic areas Barrier, Primal Dual Hybrid Gradient

Default value `1e-6`

Note The full primal and dual solution is available whether or not crossover is used. The crossover must not be disabled if the barrier is used to reoptimize nodes of a MIP.
By default crossover will not be performed on QP and MIQP problems.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

CROSSOVERITERLIMIT

Description Newton barrier and hybrid gradient: The maximum number of iterations that will be performed in the crossover procedure before the optimization process terminates.

Type Integer

Category Control

Topic areas Barrier, Primal Dual Hybrid Gradient

Default value `2147483647`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `CROSSOVER`.

CROSSOVEROPS

Description Newton barrier and hybrid gradient: a bit-vector (see Section 9.2) for adjusting the behavior of the crossover procedure.

Type Integer

Category Control

Topic areas Barrier, Primal Dual Hybrid Gradient, Bit-vector

Values	Bit	Meaning
	0	Returned solution when the crossover terminates prematurely: 0: Return the last basis from the crossover; 1: Return the barrier solution.
	1	Select the crossover stages to be performed: 0: Perform both crossover stages; 1: Skip second crossover stage.
	2	Set crossover behaviour: 0: Force to perform all pivots; 1: Skip pivots that are numerically less reliable.
	3	Set crossover behaviour: 0: Perform standard crossover; 1: Perform a slower, but numerically more careful crossover.

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also CROSSOVER.

CROSSOVERTHREADS

Description Determines the maximum number of threads that parallel crossover is allowed to use. If CROSSOVERTHREADS is set to the default value (−1), the BARTHREADS control will determine the number of threads used.

Type Integer

Category Control

Topic areas Barrier, Primal Dual Hybrid Gradient, Parallel

Default value −1 (determined by the BARTHREADS control)

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also BARTHREADS, CONCURRENTTHREADS, THREADS.

CUTDEPTH

Description Branch and Bound: Sets the maximum depth in the tree search at which cuts will be generated. Generating cuts can take a lot of time, and is often less important at deeper levels of the tree since tighter bounds on the variables have already reduced the feasible region. A value of 0 signifies that no cuts will be generated.

Type Integer

Category Control

Topic areas Cuts, Branch and Bound Search

Default value −1 — determined automatically.

Note Does not affect cutting on the root node.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also AUTOCUTTING, CUTFREQ.

CUTFACTOR

Description Limit on the number of cuts and cut coefficients the optimizer is allowed to add to the matrix during tree search. The cuts and cut coefficients are limited by CUTFACTOR times the number of rows and coefficients in the initial matrix.

Type Double

Category Control

Topic areas Cuts, Branch and Bound Search

Values

-1	Let the optimizer decide on the maximum amount of cuts based on CUTSTRATEGY.
>=0	Multiple of number of rows and coefficients to use.

Default value -1

Note A value of 0.0 prevents cuts from being added, and a value of e.g. 1.0 will allow the problem to grow to twice the initial number of rows and coefficients.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also CUTSTRATEGY.

CUTFREQ

Description Branch and Bound: This specifies the frequency at which cuts are generated in the tree search. If the depth of the node modulo CUTFREQ is zero, then cuts will be generated.

Type Integer

Category Control

Topic areas Cuts, Branch and Bound Search

Default value -1 — determined automatically.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also AUTOCUTTING, CUTDEPTH.

CUTSELECT

Description A bit-vector (see Section 9.2) providing detailed control of the cuts created for the root node of a MIP solve. Use TREECUTSELECT to control cuts during the tree search.

Type Integer

Category Control

Topic areas Cuts, Root, Bit-vector

Values	Bit	Meaning
	5	Clique cuts.
	6	Mixed Integer Rounding (MIR) cuts.
	7	Lifted cover cuts.
	8	Turn on row aggregation for MIR cuts.
	11	Flow path cuts.
	12	Implication cuts.
	13	Turn on automatic Lift-and-Project cutting strategy.
	14	Disable cutting from cut rows.
	15	Lifted GUB cover cuts.
	16	Zero-half cuts.
	17	Indicator constraint cuts.
	18	Strong Chvatal-Gomory cuts.
	20	Farkas cuts.

Default value `-1`

Note The default value is -1 which enables all bits. Any bits not listed in the above table should be left in their default 'on' state, since the interpretation of such bits might change in future versions of the optimizer.

Note The separation of Multi-Commodity Flow (MCF) cuts is controlled by `MCFCUTSTRATEGY`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `COVERCUTS`, `GOMCUTS`, `TREECUTSELECT`.

CUTSTRATEGY

Description Branch and Bound: This specifies the cut strategy. A more aggressive cut strategy, generating a greater number of cuts, will result in fewer nodes to be explored, but with an associated time cost in generating the cuts. The fewer cuts generated, the less time taken, but the greater subsequent number of nodes to be explored.

Type Integer

Category Control

Topic areas Cuts, Branch and Bound Search

Values	-1	Automatic selection of the cut strategy.
	0	No cuts.
	1	Conservative cut strategy.
	2	Moderate cut strategy.
	3	Aggressive cut strategy.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

DEFAULTALG

Description This selects the algorithm that will be used to solve LPs, standalone or during MIP optimization.

Type Integer

Category Control

Topic areas LP, Branch and Bound Search

Values

1	Automatically determined.
2	Dual simplex.
3	Primal simplex.
4	Newton barrier (or hybrid gradient, if BARALG=4 is set).

Default value 1

Note Please note that when solving MIPs or MINLPs, this will affect cut loop LPs, heuristic LPs, and node LPs. To only change how the initial root LP is solved, please use the `LPFLAGS` control or pass the appropriate flags to `XPRSoptimize` or `XPRSmipoptimize`.

Affects routines `XPRSoptimize` (`OPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `CONCURRENTTHREADS`, `LPFLAGS`.

DENSECOLLIMIT

Description Newton barrier: Columns with more than `DENSECOLLIMIT` elements are considered to be dense. Such columns will be handled specially in the Cholesky factorization of this matrix.

Type Integer

Category Control

Topic areas LP, Barrier

Default value 0 — determined automatically.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

DETERMINISTIC

Description Selects whether to use a deterministic or opportunistic mode when solving a problem using multiple threads.

Type Integer

Category Control

Topic area Determinism

Values

0	Use opportunistic mode.
1	Use deterministic mode.
2	Use deterministic mode, except allow the initial concurrent continuous solve of a MIP to be opportunistic.

Default value 1

Notes In deterministic mode thread synchronization is performed in deterministic order, which guarantees that repeated solves of the same problems under the same starting conditions will always produce the same outcome. This assumes that there are no non-deterministic events affecting the solve, such as interruption due to time limits or non-deterministic interaction through callback functions.

In opportunistic mode the solver will always schedule work on any available thread. This can produce higher CPU utilization, but will sacrifice reproducability.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `THREADS`.

DUALGRADIENT

Description Simplex: This specifies the dual simplex pricing method.

Type Integer

Category Control

Topic areas LP, Simplex

Values

-1	Determined automatically.
0	Devex.
1	Steepest edge.
2	Direct steepest edge.
3	Sparse Devex.

Default value -1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRICINGALG`.

DUALIZE

Description For a linear problem or the initial linear relaxation of a MIP, determines whether to form and solve the dual problem.

Type Integer

Category Control

Topic areas LP, Problem Transformation

Values

-1	Determine automatically which version would be faster.
0	Solve the original problem.
1	Solve the dualized problem.

Default value -1

Affects routines `XPRSoptimize` (`OPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `DUALIZEOPS`

DUALIZEOPS

Description Bit-vector control (see Section 9.2) for adjusting the behavior when a problem is dualized.

Type Integer

Category Control

Topic areas LP, Bit-vector

Values

Bit	Meaning
0	Swap the simplex algorithm to run. If dual simplex is selected for the original problem then primal simplex will be run on the dualized problem, and simiarly if primal simplex is selected.

Default value 1 (bit 0 is set)

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also DUALIZE

DUALPERTURB

Description The factor by which the problem will be perturbed prior to optimization by dual simplex. A value of 0.0 results in no perturbation prior to optimization.

Note the interconnection to the AUTOPERTURB control. If AUTOPERTURB is set to 1, the decision whether to perturb or not is left to the Optimizer. When the problem is automatically perturbed in dual simplex, however, the value of DUALPERTURB will be used for perturbation.

Type Double

Category Control

Topic areas LP, Simplex

Default value -1 — determined automatically.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also AUTOPERTURB, PRIMALPERTURB.

DUALSTRATEGY

Description This bit-vector control (see Section 9.2) specifies the dual simplex strategy.

Type Integer

Category Control

Topic areas LP, Simplex, Bit-vector

Values

Bit	Meaning
0	Switch to primal when re-optimization goes dual infeasible and numerically unstable.
1	When dual intend to switch to primal, stop the solve instead of switching to primal.
2	Use more aggressive cut-off in MIP search.
3	Use dual simplex to remove cost perturbations.
4	Enable more aggressive dual pivoting strategy.
5	Keep using dual simplex even when it's numerically unstable.

Default value 1

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

DUALTHREADS

Description Determines the maximum number of threads that dual simplex is allowed to use. If DUALTHREADS is set to the default value (-1), the THREADS control will determine the number of threads used.

Type Integer

Category Control

Topic areas LP, Parallel

Default value -1 (determined by the THREADS control)

Notes When solving a linear MIP, the dual simplex algorithm will use multiple threads only when solving the initial LP relaxation or when reoptimizing between rounds of cuts on the root node.

The parallel dual simplex algorithm differs from the sequential dual simplex algorithm and might follow a different solve path. For DUALTHREADS > 1 the solve path is independent of the number of threads used, although the practical limit for observing performance benefits is around DUALTHREADS = 8.

Affects routines XPRSlpoptimize (LPOPTIMIZE).

See also CONCURRENTTHREADS, THREADS.

EIGENVALUETOL

Description A quadratic matrix is considered not to be positive semi-definite, if its smallest eigenvalue is smaller than the negative of this value.

Type Double

Category Control

Topic areas Quadratic, Tolerances

Default value 1E-6

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), CHECKCONVEXITY.

See also IFCHECKCONVEXITY.

ELIMFILLIN

Description Amount of fill-in allowed when performing an elimination in presolve .

Type Integer

Category Control

Topic area Presolve

Default value 7

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

ELIMTOL

Description The Markowitz tolerance for the elimination phase of the presolve.

Type Double

Category Control

Topic areas Presolve, Tolerances

Default value `0.001`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

ESCAPENAMES

Description If characters illegal to an mps or lp file should be escaped to guarantee readability, and whether escaped characters should be transformed back when reading such a file.

Type Integer

Category Control

Topic area File IO

Values

`0`	Illegal characters are not escaped.
`1`	Illegal characters are escaped.

Default value `1`

Affects routines `XPRSreadprob` (`READPROB`), `XPRSwriteprob` (`WRITEPROB`).

ETATOL

Description Tolerance on eta elements. During each iteration, the basis inverse is premultiplied by an elementary matrix, which is the identity except for one column - the eta vector. Elements of eta vectors whose absolute value is smaller than `ETATOL` are taken to be zero in this step.

Type Double

Category Control

Topic areas Simplex, Tolerances

Default value `1.0E-13`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSbtran`, `XPRSftran`.

EXTRACOLS

Description The initial number of extra columns to allow for in the matrix. If columns are to be added to the matrix, then, for maximum efficiency, space should be reserved for the columns before the matrix is input by setting the `EXTRACOLS` control. If this is not done, resizing will occur automatically, but more space may be allocated than the user actually requires.

Type Integer

Category Control

Topic area Problem Creation

Default value 0

Affects routines XPRSreadprob (READPROB), XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp.

See also EXTRAROWS, EXTRAELEMS, EXTRAMIPENTS.

EXTRAELEMS

Description The initial number of extra matrix elements to allow for in the matrix, including coefficients for cuts. If rows or columns are to be added to the matrix, then, for maximum efficiency, space should be reserved for the extra matrix elements before the matrix is input by setting the EXTRAELEMS control. If this is not done, resizing will occur automatically, but more space may be allocated than the user actually requires.

Type Integer

Category Control

Topic area Problem Creation

Default value Hardware/platform dependent.

Affects routines XPRSreadprob (READPROB), XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp.

See also EXTRACOLS, EXTRAROWS.

EXTRAMIPENTS

Description The initial number of extra MIP entities to allow for.

Type Integer

Category Control

Topic area Problem Creation

Default value 0

Affects routines XPRSreadprob (READPROB), XPRSloadmip, XPRSloadmiqp.

EXTRAROWS

Description The initial number of extra rows to allow for in the matrix, including cuts. If rows are to be added to the matrix, then, for maximum efficiency, space should be reserved for the rows before the matrix is input by setting the EXTRAROWS control. If this is not done, resizing will occur automatically, but more space may be allocated than the user actually requires.

Type Integer

Category Control

Topic area Problem Creation

Default value 0

Affects routines XPRSreadprob (READPROB), XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp.

See also EXTRACOLS.

EXTRASETELEMS

Description	The initial number of extra elements in sets to allow for in the matrix. If sets are to be added to the matrix, then, for maximum efficiency, space should be reserved for the set elements before the matrix is input by setting the `EXTRASETELEMS` control. If this is not done, resizing will occur automatically, but more space may be allocated than the user actually requires.
Type	Integer
Category	Control
Topic area	Problem Creation
Default value	`0`
Affects routines	`XPRSreadprob` (`READPROB`), `XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`.
See also	`EXTRAMIPENTS`, `EXTRASETS`.

EXTRASETS

Description	The initial number of extra sets to allow for in the matrix. If sets are to be added to the matrix, then, for maximum efficiency, space should be reserved for the sets before the matrix is input by setting the `EXTRASETS` control. If this is not done, resizing will occur automatically, but more space may be allocated than the user actually requires.
Type	Integer
Category	Control
Topic area	Problem Creation
Default value	`0`
Affects routines	`XPRSreadprob` (`READPROB`), `XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`.
See also	`EXTRAMIPENTS`, `EXTRASETELEMS`.

FEASIBILITYJUMP

Description	MIP: Decides if the Feasibility Jump heuristic should be run. The value for this control is either -1 (let Xpress decide), 0 (off) or a value that indicates for which type of models the heuristic should be run.
Type	Integer
Category	Control
Topic area	Heuristics
Values	`-1` Use automatic settings. `0` Turned off. `1` Run the heuristic on models with all integer variables. `2` Run the heuristic on models in which all non-integer variables have bounds [0,1]. `3` Run the heuristic on models in which all non-integer variables have integer bounds.
Default value	`-1`
Affects routines	`XPRSmipoptimize`, `XPRSoptimize` (`MIPOPTIMIZE`, `OPTIMIZE`).

FEASIBILITYPUMP

Description Branch and Bound: Decides if the Feasibility Pump heuristic should be run at the root node.

Type Integer

Category Control

Topic area Heuristics

Values

-1	Automatic.
0	Turned off.
1	Always try the Feasibility Pump.
2	Try the Feasibility Pump only if other heuristics have failed to find an integer solution.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

FEASTOL

Description This tolerance determines when a solution is treated as feasible. If the amount by which a constraint's activity violates its right-hand side or ranged bound is less in absolute magnitude than FEASTOL, then the constraint is treated as satisfied. Similarly, if the amount by which a column violates its bounds is less in absolute magnitude than FEASTOL, those bounds are also treated as satisfied.

Type Double

Category Control

Topic area Tolerances

Default value 1.0E-06

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRSgetinfeas.

FEASTOLPERTURB

Description This tolerance determines how much a feasible primal basic solution is allowed to be perturbed when performing basis changes. The tolerance FEASTOL is always considered as an upper limit for the perturbations, but in some cases smaller value can be more desirable.

Type Double

Category Control

Topic areas Simplex, Tolerances

Default value 1.0E-06

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRSgetinfeas.

FEASTOLTARGET

Description This specifies the target feasibility tolerance for the solution refiner.

Type Double

Category Control

Topic areas Solution Refinement, Tolerances

Default value `0` — use the value specified by `FEASTOL`.

Note Zero and negative values are ignored, and the value of `FEASTOL` is used.

Note Use very small values like 1e-100 to state the refinement should continue as long as an improvement is made. Use very large values like 1e+100 to disable only this aspect of the refiner.

Note Refining solutions to match the feastoltarget can influence and worsen their objective value in case the previous objective could only be achieved through slight infeasibilities.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `REFINEOPS`, `LPREFINEITERLIMIT`, `OPTIMALITYTOLTARGET`.

FORCEOUTPUT

Description Certain names in the problem object may be incompatible with different file formats (such as names containing spaces for LP files). If the optimizer might be unable to read back a problem because of non-standard names, it will first attempt to write it out using an extended naming convention. If the names would not be possible to extend so that they would be reproducible and recognizable, it will give an error message and won't create the file. If the optimizer might be unable to read back a problem because of non-standard names, it will give an error message and won't create the file. This option may be used to force output anyway.

Type Integer

Category Control

Topic area File IO

Values

`0`	Check format compatibility, and in case of failure try to extend names so that they are reproducible and recognizable.
`1`	Force output using problem names as is.
`2`	Always use 'x(' original name ')' in LP files to create a representation that can be read by Xpress. Default for problem having spaces in names
`3`	Substitute spaces by the '_' character in LP files

Default value `0`

Affects routines `XPRSwriteprob` (`WRITEPROB`).

FORCEPARALLELDUAL

Description Dual simplex: specifies whether the dual simplex solver should always use the parallel simplex algorithm. By default, when using a single thread, the dual simplex solver will execute a dedicated sequential simplex algorithm.

Type Integer

Category Control

Topic areas LP, Simplex

Values

0	Disabled.
1	Enabled. Force the dual simplex solver to use the parallel algorithm.

Default value 0

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also THREADS, DUALTHREADS.

GENCONSABSTRANSFORMATION

Description This control specifies the reformulation method for absolute value general constraints at the beginning of the search.

Type Integer

Category Control

Topic areas Presolve, Piecewise Linear and General Constraints

Values

-1	Automatic.
0	Use a formulation based on indicator constraints.
1	Use a formulation based on SOS1-constraints.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

GENCONSDUALREDUCTIONS

Description This parameter specifies whether dual reductions should be applied to reduce the number of columns and rows added when transforming general constraints to MIP structs.

Type Integer

Category Control

Topic areas Presolve, Piecewise Linear and General Constraints

Values

0	Disabled. No dual reductions, add columns and rows.
1	Enabled. Only add neccessary columns and rows, drop those implied by the objective sense.

Default value 1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPDUALREDUCTIONS.

TREEFILELOGINTERVAL

Description This control sets the interval between progress messages output while writing tree data to the tree file, in seconds. The solve is slowed greatly while data is being written to the tree file and this output allows the user to see how much progress is being made.

Type Integer

Category Control

Topic areas Branch and Bound Search, Logging

Default value `60`

See also `TREEDIAGNOSTICS`.

GLOBALBOUNDINGBOX

Description If a nonlinear problem cannot be solved due to appearing unbounded, it can automatically be regularized by the application of a bounding box on the variables. If this control is set to a negative value, in a second solving attempt all original variables will be bounded by the absolute value of this control. If set to a positive value, there will be a third solving attempt afterwards, if necessary, in which also all auxiliary variables are bounded by this value.

Type Double

Category Control

Topic area Global

Values

`0`	Disabled. Problem will return unbounded.
`n<0`	Enabled. Apply lower and upper bounds of this magnitude to all original variables if initial LP is unbounded.
`n>0`	Enabled. Apply lower and upper bounds of this magnitude to all original and auxiliary variables if initial LP and first regularization are unbounded.

Default value `1.0E+06`

See also `GLOBALBOUNDINGBOXAPPLIED`.

Affects routines `XPRSoptimize` (`OPTIMIZE`).

GLOBALLSHEURSTRATEGY

Description When integer-feasible (for MINLP, any solution for NLP) but nonlinear-infeasible solutions are encountered within a global solve, the integer variables can be fixed and a local solver (as defined by the `LOCALSOLVER` control) can be called on the remaining continuous problem. This control defines the frequency and effort of such local solves.

Type Double

Category Control

Topic areas Global, Heuristics

Values	`-1`	Automatic selection of the strategy.
	`0`	Never run a local solver on a partially fixed solution.
	`1`	Conservative strategy.
	`2`	Moderate strategy.
	`3`	Aggressive strategy.

Default value `-1`

See also `LOCALSOLVER`.

Affects routines `XPRSoptimize` (`OPTIMIZE`).

GLOBALNLPCUTS

Description Limit on the number of rounds of outer approximation and convexification cuts generated for the root node, when solving an (MI)NLP to global optimality.

Type Integer

Category Control

Topic areas Global, Cuts

Default value `-1` — determined automatically.

Affects routines `XPRSoptimize` (`OPTIMIZE`).

GLOBALNUMINITNLPCUTS

Description Specifies the maximum number of tangent cuts when setting up the initial relaxation during a global solve. By default, the algorithm chooses the number of cuts automatically. Adding more cuts tightens the problem, resulting in a smaller branch-and-bound tree, at the cost of slowing down each LP solve.

Type Integer

Category Control

Topic areas Global, Cuts

Default value `-1` — determined automatically.

Affects routines `XPRSoptimize` (`OPTIMIZE`).

GLOBALSPATIALBRANCHCUTTINGEFFORT

Description Limits the effort that is spent on creating cuts during spatial branching.

Type Double

Category Control

Topic areas Global, Branching

Values	`-1`	The algorithm chooses the effort limit automatically (default).
	`0`	Disables cuts on branching entities.
	`0<n<1`	Relative effort to spend on cutting on branching entities. Higher values lead to more cuts.

Default value -1.0

See also `GLOBALSPATIALBRANCHPROPAGATIONEFFORT`

GLOBALSPATIALBRANCHIFPREFERORIG

Description Whether spatial branchings on original variables should be preferred over branching on auxiliary variables that were introduced by the reformulation of the global solver.

Type Integer

Category Control

Topic areas Global, Branching

Values

`-1`	Automatically determined.
`0`	Always consider both original and auxiliary variables.
`1`	Always prefer branching on original variables over auxiliaries.
`2`	Always prefer branching on auxiliary variables over originals.

Default value `-1`

GLOBALSPATIALBRANCHPROPAGATIONEFFORT

Description Limits the effort that is spent on propagation during spatial branching.

Type Double

Category Control

Topic areas Global, Propagation

Values

`-1`	The algorithm chooses the effort limit automatically (default).
`0`	Disables propagation on branching entities.
`n>0`	Relative effort to spend on propagating on branching entities. Higher values lead to more propagation.

Default value -1.0

GLOBALTREENLPCUTS

Description Limit on the number of rounds of outer approximation and convexification cuts generated for each node in the tree, when solving an (MI)NLP to global optimality.

Type Integer

Category Control

Topic areas Global, Cuts

Default value `-1` — determined automatically.

Affects routines `XPRSoptimize` (`OPTIMIZE`).

GOMCUTS

Description Branch and Bound: The number of rounds of Gomory or lift-and-project cuts at the root node.

Type Integer

Category Control

Topic areas Cuts, Root

Default value `-1` — determined automatically.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `TREEGOMCUTS`, `LNPBEST`, `LNPITERLIMIT`.

HEURBEFORELP

Description Branch and Bound: Determines whether primal heuristics should be run before the initial LP relaxation has been solved.

Type Integer

Category Control

Topic area Heuristics

Values

`-1`	Automatic - let the optimizer decide if heuristics should be run.
`0`	Disabled.
`1`	Enabled.

Default value `-1`

Note It is possible that a heuristic will find an optimal integer solution that will result in the LP relaxation solution being cut off. If the problem is solved with the "`l`" flag to `XPRSmipoptimize` (i.e., stop after solving the LP relaxation), then `LPSTATUS` might be returned as `XPRS_LP_CUTOFF` or `XPRS_LP_CUTOFF_IN_DUAL`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `HEUREMPHASIS`

HEURDEPTH

Description Branch and Bound: Sets the maximum depth in the tree search at which heuristics will be used to find MIP solutions. It may be worth stopping the heuristic search for solutions after a certain depth in the tree search. A value of 0 signifies that heuristics will not be used. This control no longer has any effect and will be removed from future releases.

Type Integer

Category Control

Topic areas Heuristics, Branch and Bound Search

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

HEURDIVEITERLIMIT

Description Branch and Bound: Simplex iteration limit for reoptimizing during the diving heuristic.

Type Double

Category Control

Topic areas Heuristics, Limits

Values

>=1	Fixed iteration limit.
0	No iteration limit.
<0	Automatic selection of the iteration limit based on the problem size. The absolute value is used as a multiplier on the automatic selection.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also HEUREMPHASIS.

HEURDIVERANDOMIZE

Description The level of randomization to apply in the diving heuristic. The diving heuristic uses priority weights on rows and columns to determine the order in which to e.g. round fractional columns, or the direction in which to round them. This control determines by how large a random factor these weights should be changed.

Type Double

Category Control

Topic area Heuristics

Values 0.0-1.0 Amount of randomization (0.0=none, 1.0=full)

Default value 0.0

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also HEURDIVESTRATEGY, HEURDIVESPEEDUP.

HEURDIVESOFTROUNDING

Description Branch and Bound: Enables a more cautious strategy for the diving heuristic, where it tries to push binaries and integer variables to their bounds using the objective, instead of directly fixing them. This can be useful when the default diving heuristics fail to find any feasible solutions.

Type Integer

Category Control

Topic area Heuristics

Values

-1	Automatic selection.
0	Do not use soft rounding.
1	Cautious use of the soft rounding strategy.
2	More aggressive use of the soft rounding strategy.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also HEURDIVESTRATEGY.

HEURDIVESPEEDUP

Description Branch and Bound: Changes the emphasis of the diving heuristic from solution quality to diving speed.

Type Integer

Category Control

Topic area Heuristics

Values

-2	Automatic selection biased towards quality
-1	Automatic selection biased towards speed.
0-4	manual emphasis bias from emphasis on quality (0) to emphasis on speed (4).

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also HEURDIVESTRATEGY.

HEURDIVESTRATEGY

Description Branch and Bound: Chooses the strategy for the diving heuristic.

Type Integer

Category Control

Topic area Heuristics

Values

-1	Automatic selection of strategy.
0	Disables the diving heuristic.
1-18	Available pre-set strategies for rounding infeasible MIP entities and reoptimizing during the heuristic dive.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also HEUREMPHASIS.

HEUREMPHASIS

Description Branch and Bound: This control specifies an emphasis for the search w.r.t. primal heuristics and other procedures that affect the speed of convergence of the primal-dual gap. For problems where the goal is to achieve a small gap but not neccessarily solving them to optimality, it is recommended to set HEUREMPHASIS to 1. This setting triggers many additional heuristic calls, aiming for reducing the gap at the beginning of the search, typically at the expense of an increased time for proving optimality.

Type Integer

Category Control

Topic area Heuristics

Values

-1	Optimizer default strategy.
0	Disables all heuristics.
1	Focus on reducing the primal-dual gap in the early part of the search.
2	Extremely aggressive search heuristics.

Default value -1

Note A setting of this control to 2 will enable the Pre-root parallel heuristic phase for a limited, problem-dependent amount of work. This phase can be explicitly disabled by setting PREROOTEFFORT to 0.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

HEURFORCESPECIALOBJ

Description Branch and Bound: This specifies whether local search heuristics without objective or with an auxiliary objective should always be used, despite the automatic selection of the Optimiezr. Deactivated by default.

Type Integer

Category Control

Topic area Heuristics

Values

0	Disabled.
1	Enabled. Run special objective heuristics on large problems and even if incumbent exists.

Default value 0

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also HEURSEARCHROOTSELECT, HEURSEARCHTREESELECT.

HEURFREQ

Description Branch and Bound: This specifies the frequency at which heuristics are used in the tree search. Heuristics will only be used at a node if the depth of the node is a multiple of HEURFREQ.

Type Integer

Category Control

Topic areas Heuristics, Branch and Bound Search

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

HEURMAXSOL

Description Branch and Bound: This specifies the maximum number of heuristic solutions that will be found in the tree search. This control no longer has any effect and will be removed from future releases.

Type Integer

Category Control

Topic area Heuristics

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

HEURNODES

Description Branch and Bound: This specifies the maximum number of nodes at which heuristics are used in the tree search. This control no longer has any effect and will be removed from future releases.

Type Integer

Category Control

Topic area Heuristics

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

HEURSEARCHEFFORT

Description Adjusts the overall level of the local search heuristics.

Type Double

Category Control

Topic area Heuristics

Default value `1.0`

Note `HEURSEARCHEFFORT` is used as a multiplier on the default amount of work the local search heuristics should do. A higher value means the local search heuristics will be run more often and that they are allowed to search larger neighborhoods.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `HEUREMPHASIS`, `HEURSEARCHROOTSELECT`, `HEURSEARCHTREESELECT`.

HEURSEARCHFREQ

Description	Branch and Bound: This specifies how often the local search heuristic should be run in the tree.
Type	Integer
Category	Control
Topic areas	Heuristics, Branch and Bound Search
Values	`-1` Automatic. `0` Disabled in the tree. `n>0` Number of nodes between each run.
Default value	`-1`
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`HEURSEARCHROOTCUTFREQ`.

HEURSEARCHROOTCUTFREQ

Description	How frequently to run the local search heuristic during root cutting. This is given as how many cut rounds to perform between runs of the heuristic. Set to zero to avoid applying the heuristic during root cutting. Branch and Bound: This specifies how often the local search heuristic should be run in the tree.
Type	Integer
Category	Control
Topic areas	Heuristics, Root
Values	`-1` Automatic. `0` Disabled heuristic during cutting. `n>0` Number of cutting rounds between each run.
Default value	`-1`
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`HEURSEARCHFREQ`.

HEURSEARCHBACKGROUNDSELECT

Description	Bit-vector control (see Section 9.2) to select which large neighborhood searches to run in the background (for example in parallel to the root cut loop).
Type	Integer
Category	Control
Topic areas	Heuristics, Root, Bit-vector
Values	Bit Meaning `1` Enable `L` heuristic.
Default value	-1

Notes If this is set to -1 (the default) then Xpress decides automatically which heuristics to run. Otherwise the value is interpreted as a bitset that indices which heuristics to run. 0 (zero) means not to run any large neighborhood search heuristics in the background.

Other types of background jobs are controlled by `BACKGROUNDSELECT`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSoptimize` (`OPTIMIZE`).

See also `BACKGROUNDSELECT`, `BACKGROUNDMAXTHREADS`.

HEURSEARCHCOPYCONTROLS

Description Select how user-set controls should affect local search heuristics.

Type Integer

Category Control

Topic areas Heuristics, Controls and Attributes

Default value 1

Values

0	Do not copy any user-set controls
1	Automatic - Let the Optimizer decide which user-set controls to copy
2	Copy all user-set controls

Notes If you set any control to prevent an unfavorable behavior, consider also setting `HEURSEARCHCOPYCONTROLS` to 2.

Termination criteria are not affected by this control.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSoptimize` (`OPTIMIZE`).

See also `BACKGROUNDSELECT`, `HEURSEARCHROOTSELECT`, `HEURSEARCHTREESELECT`.

HEURSEARCHROOTSELECT

Description A bit-vector control (see Section 9.2) for selecting which local search heuristics to apply on the root node of a MIP solve. Use `HEURSEARCHTREESELECT` to control local search heuristics during the tree search.

Type Integer

Category Control

Topic areas Heuristics, Root, Bit-vector

Values

Bit	Meaning
0	Local search with a large neighborhood. Potentially slow but is good for finding solutions that differs significantly from the incumbent.
1	Local search with a small neighborhood centered around a node LP solution.
2	Local search with a small neighborhood centered around an integer solution. This heuristic will often provide smaller, incremental improvements to an incumbent solution.
3	Local search with a neighborhood set up through the combination of multiple integer solutions.
4	*Unused*
5	Local search without an objective function. Called seldom and only when no feasible solution is available.
6	Local search with an auxiliary objective function. Called seldom and only when no feasible solution is available.

Default value `117`

Note Some of the local search heuristics will benefit from having an existing incumbent solution, but it is not required. An initial solution can also be provided by the user through either `XPRSloadmipsol` or `XPRSreadbinsol`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `HEUREMPHASIS`, `HEURSEARCHTREESELECT`, `HEURSEARCHEFFORT`.

HEURSEARCHTREESELECT

Description A bit-vector control (see Section 9.2) for selecting which local search heuristics to apply during the tree search of a MIP solve. Use `HEURSEARCHROOTSELECT` to control local search heuristics on the root node.

Type Integer

Category Control

Topic areas Heuristics, Branch and Bound Search, Bit-vector

Values

Bit	Meaning
0	Local search with a large neighborhood. Potentially slow but is good for finding solutions that differs significantly from the incumbent.
1	Local search with a small neighborhood centered around a node LP solution.
2	Local search with a small neighborhood centered around an integer solution. This heuristic will often provide smaller, incremental improvements to an incumbent solution.
3	Local search with a neighborhood set up through the combination of multiple integer solutions.
4	*Unused*
5	Local search without an objective function. Called seldom and only when no feasible solution is available.
6	Local search with an auxiliary objective function. Called seldom and only when no feasible solution is available.

Default value `17`

Note Some of the local search heuristics will benefit from having an existing incumbent solution, but it is not required. An initial solution can also be provided by the user through either `XPRSloadmipsol` or `XPRSaddmipsol`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `HEUREMPHASIS`, `HEURSEARCHROOTSELECT`, `HEURSEARCHEFFORT`.

HEURSHIFTPROP

Description Determines whether the Shift-and-propagate primal heuristic should be executed. If enabled, Shift-and-propagate is an LP-free primal heuristic that is executed immediately after presolve.

Type Integer

Category Control

Topic area Heuristics

Values

`-1`	The solver decides if Shift-and-propagate should be run. This is the default setting.
`0`	Shift-and-propagate is disabled.
`1`	Shift-and-propagate is enabled.

Default value `-1`

Affects routines `XPRSoptimize` (`OPTIMIZE`).

See also `HEUREMPHASIS`

HEURTHREADS

Description Branch and Bound: The number of threads to dedicate to running heuristics during the root solve.

Type Integer

Category Control

Topic areas Heuristics, Parallel

Values

`-1`	Automatically determined from the `THREADS` control.
`0`	Disabled.
`>=1`	Number of additional threads to dedicate to parallel heuristics.

Default value `0`

Note When heuristic threads are enabled, the heuristics will be run in parallel with the root cutting.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `THREADS`.

HISTORYCOSTS

Description Branch and Bound: How to update the pseudo cost for a MIP entity when a strong branch or a regular branch is applied.

Type Integer

Category Control

Topic area Branching

Values

-1	Automatically determined.
0	No update.
1	Update using only regular branches from the root to the current node.
2	Same as 1, but update with strong branching results as well.
3	Update using any regular branching or strong branching information from all nodes solves before the current node.

Default value -1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `SBBEST`, `SBESTIMATE`, `SBSELECT`

IFCHECKCONVEXITY

Description Determines if the convexity of the problem is checked before optimization. Applies to quadratic, mixed integer quadratic and quadratically constrained problems. Checking convexity takes some time, thus for problems that are known to be convex it might be reasonable to switch the checking off.

Type Integer

Category Control

Topic areas Presolve, Quadratic

Values

0	Turn off convexity checking.
1	Turn on convexity checking.

Default value 1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `EIGENVALUETOL`

IISLOG

Description Selects how much information should be printed during the IIS procedure. Please refer to Appendix A.13 for a more detailed description of the IIS logging format.

Type Integer

Category Control

Topic areas Infeasibility, Logging

Values	0	The IIS procedure does not produce any output.
	1	Prints general information and a progress log of the deletion filter, including bounds on the size of the IIS and timing information.
	2	Complete logging, including the full progress log of all the subproblem solves in the deletion filter. This setting is recommended only for debugging as it may produce a lot of output.

Default value `1`, a progress log is printed

Affects routines `XPRSiisfirst (IIS)`.

IISOPS

Description Selects which part of the restrictions (bounds, constraints, entities) should always be kept in an IIS. This is useful if certain types of restrictions cannot be violated, thus they are known not to be the cause of infeasibility. The IIS obtained this way is irreducible only for the non-protected restrictions.

This bit-vector control (see Section 9.2) has an effect only on the deletion filter of the IIS procedure.

Type Integer

Category Control

Topic areas Infeasibility, Bit-vector

Values

Bit	Meaning
0	Keep binary integralities.
1	Keep the 0 lower bounds of variables.
2	Keep fixed variables.
3	Keep all variable bounds.
4	Keep all general integer entities, except binaries.
5	Keep all equality constraints.
6	Keep all general constraints.
7	Keep all piecewise linear constraints.
8	Keep all specially ordered set (SOS) constraints.
9	Keep all indicator constraints.
10	Keep all delayed rows.
11	Keep all constraints.

Default value `0`, all restrictions are valid candidates for removal

Affects routines `XPRSiisfirst (IIS)`.

INDLINBIGM

Description During presolve, indicator constraints will be linearized using a BigM coefficient whenever that BigM coefficient is small enough. This control defines the largest BigM for which such a linearized version will be added to the problem in addition to the original constraint. If the BigM is even smaller than `INDPRELINBIGM`, then the original indicator constraint will additionally be dropped from the problem.

Type Double

Category Control

Topic areas Presolve, Problem Transformation

Default value `1.0E+05`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

Note INDLINBIGM should always be at least as large as `INDPRELINBIGM`. For any value less or equal to `INDPRELINBIGM`, indicator constraints will never be duplicated and only `INDPRELINBIGM` is taken into account for linearization.

See also `INDPRELINBIGM`

INDPRELINBIGM

Description During presolve, indicator constraints will be linearized using a BigM coefficient whenever that BigM coefficient is small enough. This control defines the largest BigM for which the original constraint will be replaced by the linearized version. If the BigM is larger than INDPRELINBIGM but smaller than `INDLINBIGM`, the linearized row will be added but the original indicator constraint is kept as a numerically stable way to check feasibility.

Type Double

Category Control

Topic areas Presolve, Problem Transformation

Default value `100.0`

Note Replacing an indicator constraint with a BigM row has a side effect on tolerances. In the indicator constraint form, the constraint part is satisfied with FEASTOL tolerance; while after changing it to BigM form, the constraint also includes the binary indicator variable (with a coefficient up to INDPRELINBIGM and an integrality tolerance of MIPTOL), therefore the constraint part of the indicator constraint is satisfied with tolerance FEASTOL+MIPTOL*INDPRELINBIGM.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `INDLINBIGM`

INPUTTOL

Description The tolerance on input values elements. If any value is less than or equal to `INPUTTOL` in absolute value, it is treated as zero. For the internal zero tolerance see `MATRIXTOL`.

Type Double

Category Control

Topic areas Problem Creation, Tolerances, File IO

Note This control needs to be set before inputting the problem, as it has no effect afterwards.

Default value `0.0`

Affects routines `XPRSreadprob` (`READPROB`), `XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSalter` (`ALTER`), `XPRSaddcols`, `XPRSaddcuts`, `XPRSaddrows`, `XPRSchgcoef`, `XPRSchgmcoef`, `XPRSstorecuts`.

INVERTFREQ

Description Simplex: The frequency with which the basis will be inverted. The basis is maintained in a factorized form and on most simplex iterations it is incrementally updated to reflect the step just taken. This is considerably faster than computing the full inverted matrix at each iteration, although after a number of iterations the basis becomes less well-conditioned and it becomes necessary to compute the full inverted matrix. The value of `INVERTFREQ` specifies the maximum number of iterations between full inversions.

Type Integer

Category Control

Topic area Simplex

Default value `-1` — the frequency is determined automatically.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

INVERTMIN

Description Simplex: The minimum number of iterations between full inversions of the basis matrix. See the description of `INVERTFREQ` for details.

Type Integer

Category Control

Topic area Simplex

Default value `3`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

IOTIMEOUT

Description The maximum number of seconds to wait for an I/O operation before it is cancelled.

Type Integer

Category Control

Topic areas File IO, Limits

Default value `30`

Affects routines `XPRSreadprob` (`READPROB`), `XPRSreaddirs` (`READDIRS`), `XPRSreadslxsol` (`READSLXSOL`), `XPRSreadbasis` (`READBASIS`), `XPRSreadbinsol` (`READBINSOL`), `XPRSwritedirs` (`WRITEDIRS`), `XPRSwritebasis` (`WRITEBASIS`), `XPRSwritesol` (`WRITESOL`), `XPRSwritebinsol` (`WRITEBINSOL`), `XPRSwriteprtsol` (`WRITEPRTSOL`), `XPRSwriteslxsol` (`WRITESLXSOL`), `XPRSwriteprob` (`WRITEPROB`), `XPRSsave` (`SAVE`), `XPRSsaveas`, `XPRSrestore` (`RESTORE`), `XPRSiiswrite` (`IIS`), `XPRSalter` (`ALTER`).

KEEPBASIS

Description Simplex: This determines whether the basis should be kept when reoptimizing a problem. The choice is between using a crash basis created at the beginning of simplex or using a basis from a previous solve (if such exists). By default, this control gets (re)set automatically in various situations. By default, it will be automatically set to 1 after a solve that produced a valid basis. This will automatically warmstart a subsequent solve. Explicitly loading a starting basis will also set this control to 1. If the control is explicitly set to 0, any existing basis will be ignored for a new solve, and the Optimizer will start from an ad-hoc crash basis.

Type Integer

Category Control

Topic areas LP, Simplex

Values

`0`	Problem optimization starts from scratch, i.e., any previous basis is ignored.
`1`	The previous basis should be used as a starting basis.
`2`	Use the previous basis only if it is valid for the current problem (the number of basic variables must match the number of rows).
`3`	Use the previous basis only if it is valid and numerically stable in the current problem.

Default value `0`

Note Some pre-LP MIP heuristics are affected by this control (since they use a starting basis if one is available).

Note Keeping a basis can lead to less effective presolving since the Optimizer will avoid presolving operations that invalidate the given basis.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

KEEPNROWS

Description How nonbinding rows should be handled by the MPS reader.

Type Integer

Category Control

Topic areas File IO, Problem Creation

Values

`-1`	Delete N type rows from the matrix.
`0`	Delete elements from N type rows leaving empty N type rows in the matrix.
`1`	Keep N type rows.

Default value `-1`

Affects routines `XPRSreadprob` (`READPROB`).

L1CACHE

Description *This parameter is deprecated and will be removed in a future release.*
Newton barrier: L1 cache size in kB (kilo bytes) of the CPU. On Intel (or compatible) platforms a value of -1 may be used to determine the cache size automatically.

Type Integer

Category Control

Topic areas System, Barrier

Default value Hardware/platform dependent.

Notes Specifying the correct L1 cache size can give a significant performance advantage with the Newton barrier algorithm.

If the size is unknown, it is better to specify a smaller size.

If the size cannot be determined automatically on Intel (or compatible) platforms, a default size of 8 kB is assumed.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

LNPBEST

Description Number of infeasible MIP entities to create lift-and-project cuts for during each round of Gomory cuts at the root node (see `GOMCUTS`).

Type Integer

Category Control

Topic areas Cuts, Root

Default value `50`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

LNPITERLIMIT

Description Number of iterations to perform in improving each lift-and-project cut.

Type Integer

Category Control

Topic areas Cuts, Limits

Default value `-1` — determined automatically.

Note By setting the number to zero a Gomory cut will be created instead.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

LOCALCHOICE

Description Controls when to perform a local backtrack between the two child nodes during a dive in the branch and bound tree.

Type Integer

Category Control

Topic area Branch and Bound Search

Values	1	Never backtrack from the first child, unless it is dropped (infeasible or cut off).
	2	Always solve both child nodes before deciding which child to continue with.
	3	Automatically determined.

Default value 3

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

LPFLAGS

Description A bit-vector control (see Section 9.2) which defines the algorithm for solving an LP problem or the initial LP relaxation of a MIP problem.

Type Integer

Category Control

Topic areas LP, Root, Bit-vector

Values

Bit	Meaning
0	Use the dual simplex method.
1	Use the primal simplex method.
2	Use the barrier method (or hybrid gradient method if BARALG=4 is set).
3	Use the network simplex method.

Default value 0

Notes Setting bit 0, 1, 2, 3 of this control will have the same effect of passing flags `d`, `p`, `b`, `n` to `XPRSmipoptimize` or `XPRSlpoptimize`. When more than one bit are set, then the LP problem will be solved with the concurrent solver. When this control is set and flags are passed at the same time, the flags will overrule the value of the control.

This control can be tuned.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `DEFAULTALG`

LPFOLDING

Description Simplex and barrier: whether to fold an LP problem before solving it.

Type Integer

Category Control

Topic area LP

Values	-1	Automatic.
	0	Disable LP folding.
	1	Enable LP folding. Attempt to fold all LP problems and MIP initial relaxations.

Default value -1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`).

LPITERLIMIT

Description The maximum number of iterations that will be performed by primal simplex or dual simplex before the optimization process terminates. For MIP problems, this is the maximum total number of iterations over all nodes explored by the Branch and Bound method.

Type Integer

Category Control

Topic areas LP, Limits

Default value `2147483647`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

LPLOG

Description Simplex: The frequency at which the simplex log is printed.

Type Integer

Category Control

Topic areas LP, Logging

Values

`n<0`	Detailed output every -*n* iterations.
`0`	Log displayed at the end of the optimization only.
`n>0`	Summary output every *n* iterations.

Default value `100`

Note This control only has an effect if `LPLOGSTYLE` is set to 0.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`).

See also A.7.

LPLOGDELAY

Description Time interval between two LP log lines.

Type Double

Category Control

Topic areas LP, Logging

Default value `1.0`

Note This control only has an effect if `LPLOGSTYLE` is set to 1.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`).

LPLOGSTYLE

Description Simplex: The style of the simplex log.

Type Integer

Category Control

Topic areas LP, Logging

Values

0	Simplex log is printed based on simplex iteration count, at a fixed frequency as specified by the LPLOG control.
1	Simplex log is printed based on an estimation of elapsed time, determined by an internal deterministic timer.

Default value 1

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRSminim (MINIM), XPRSmaxim (MAXIM).

LPREFINEITERLIMIT

Description This specifies the simplex iteration limit the solution refiner can spend in attempting to increase the accuracy of an LP solution.

Type Integer

Category Control

Topic areas Solution Refinement, LP

Default value -1 — determined automatically.

Note The solution refiner iteratively attempts to increase the accuracy of the solution until either both FEASTOLTARGET and OPTIMALITYTOLTARGET is satisfied, or accuracy cannot further be increased, or the effort limit determined by LPREFINEITERLIMIT is exhausted.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also REFINEOPS, FEASTOLTARGET, OPTIMALITYTOLTARGET.

MARKOWITZTOL

Description The Markowitz tolerance used for the factorization of the basis matrix.

Type Double

Category Control

Topic areas LP, Simplex

Default value 0.01

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

MATRIXTOL

Description The zero tolerance on matrix elements. If the value of a matrix element is less than or equal to `MATRIXTOL` in absolute value, it is treated as zero. The control applies when solving a problem, for an input tolerance see `INPUTTOL`.

Type Double

Category Control

Topic area Tolerances

Default value `1.0E-09`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

MAXCHECKSONMAXCUTTIME

Description This control is intended for use where optimization runs that are terminated using the `MAXCUTTIME` control are required to be reproduced exactly. This control is necessary because of the inherent difficulty in terminating algorithmic software in a consistent way using temporal criteria. The control value relates to the number of times the optimizer checks the `MAXCUTTIME` criterion up to and including the check when the termination of cutting was activated. To use the control the user first must obtain the value of the `CHECKSONMAXCUTTIME` attribute after the run returns. This attribute value is the number of times the optimizer checked the `MAXCUTTIME` criterion during the last call to the optimization routine `XPRSmipoptimize`. Note that this attribute value will be negative if the optimizer terminated cutting on the `MAXCUTTIME` criterion. To ensure accurate reproduction of a run the user should first ensure that `MAXCUTTIME` is set to its default value or to a large value so the run does not terminate again on `MAXCUTTIME` and then simply set the control `MAXCHECKSONMAXCUTTIME` to the absolute value of the `CHECKSONMAXCUTTIME` value.

Type Integer

Category Control

Topic area Limits

Values

`0`	Not active.
`n>0`	The number of times the optimizer should check the `MAXCUTTIME` criterion before triggering a termination.

Default value `0`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

MAXCHECKSONMAXTIME

Description This control is intended for use where optimization runs that are terminated using the `TIMELIMIT` (or the deprecated `MAXTIME`) control are required to be reproduced exactly. This control is necessary because of the inherent difficulty in terminating algorithmic software in a consistent way using temporal criteria. The control value relates to the number of times the optimizer checks the `TIMELIMIT` criterion up to and including the check when the termination was activated. To use the control the user first must obtain the value of the `CHECKSONMAXTIME` attribute after the run returns. This attribute value is the number of times the optimizer checked the `TIMELIMIT` criterion during the last call to the optimization routine

XPRSmipoptimize. Note that this attribute value will be negative if the optimizer terminated on the TIMELIMIT criterion. To ensure that a reproduction of a run terminates in the same way the user should first ensure that TIMELIMIT is set to its default value or to a large value so the run does not terminate again on TIMELIMIT and then simply set the control MAXCHECKSONMAXTIME to the absolute value of the CHECKSONMAXTIME value.

Type Integer

Category Control

Topic area Limits

Values
- 0 — Not active.
- n>0 — The number of times the optimizer should check the TIMELIMIT (or MAXTIME) criterion before triggering a termination.

Default value 0

Affects routines XPRSmipoptimize (MIPOPTIMIZE)

MAXCUTTIME

Description The maximum amount of time allowed for generation of cutting planes and reoptimization. The limit is checked during generation and no further cuts are added once this limit has been exceeded.

Type Double

Category Control

Topic areas Limits, Cuts

Values
- 0 — No time limit.
- >0 — Stop cut generation after the given number of seconds.

Default value 0

Note This control changed type with Xpress Optimizer version 9.0 (from integer to double). It is strongly recommended to adapt all access calls to this control accordingly.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

MAXTREEFILESIZE

Description The maximum size, in megabytes, to which the tree file may grow, or 0 for no limit. When the tree file reaches this limit, a second tree file will be created. Useful if you are using a filesystem that puts a maximum limit on the size of a file.

Type Integer

Category Control

Topic areas Limits, File IO

Default value 0

See also TREEFILESIZE.

MAXIIS

Description This function controls the number of Irreducible Infeasible Sets to be found using the `XPRSiisall` (`IIS -a`).

Type Integer

Category Control

Topic areas Infeasibility, Limits

Values

`-1`	Search for all IIS.
`0`	Do not search for IIS.
`n>0`	Search for the first *n* IIS.

Default value `-1`

Note The function `XPRSiisnext` is not affected by this control.

Affects routines `XPRSiisall` (`IIS-a`).

MAXIMPLIEDBOUND

Description Presolve: When tighter bounds are calculated during MIP preprocessing, only bounds whose absolute value are smaller than `MAXIMPLIEDBOUND` will be applied to the problem.

Type Double

Category Control

Topic areas Presolve, Numerics

Default value `1.0E+08`

Note For numerically challenging MIP problems, it can sometimes help make the solve more stable by reducing the value of `MAXIMPLIEDBOUND` to something smaller - e.g. 1.0E+06. It is not recommended to increase this parameter beyond the default of 1.0E+08.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

MAXLOCALBACKTRACK

Description Branch-and-Bound: How far back up the current dive path the optimizer is allowed to look for a local backtrack candidate node.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

`-1`	Automatic.
`n>0`	Local backtrack limit.

Default value `-1`

Note If this control is set to *k*, then the candidate set of nodes for a local backtrack will consist of all active nodes in the subtree rooted at height *k* above the current node. For example, a setting of `1` will result in only sibling nodes of the current node being considered.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also LOCALCHOICE.

MAXMCOEFFBUFFERELEMS

Description The maximum number of matrix coefficients to buffer before flushing into the internal representation of the problem. Buffering coefficients can offer a significant performance gain when you are building a matrix using XPRSchgcoef or XPRSchgmcoef, but can lead to a significant memory overhead for large matrices, which this control allows you to influence.

Type Integer

Category Control

Topic area Problem Creation

Default value 2147483647

Affects routines XPRSchgcoef, XPRSchgmcoef.

MAXMEMORYHARD

Description This control sets the maximum amount of memory in megabytes the optimizer should allocate. If this limit is exceeded, the solve will terminate. This control is designed to make the optimizer stop in a controlled manner, so that the problem object is valid once termination occurs. The solve state will be set to incomplete. This is different to an out of memory condition in which case the optimizer returns an error. The optimizer may still allocate memory once the limit is exceeded to be able to finsish the operations and stop in a controlled manner. When RESOURCESTRATEGY is enabled, the control also has the same effect as MAXMEMORYSOFT and will cause the optimizer to try preserving memory when possible.

Type Integer

Category Control

Topic areas Limits, Memory

Default value 0 (no limit)

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

MAXMEMORYSOFT

Description When RESOURCESTRATEGY is enabled, this control sets the maximum amount of memory in megabytes the optimizer targets to allocate. This may change the solving path, but will not cause the solve to terminate early. To set a hard version of the same, please set MAXMEMORYHARD.

Type Integer

Category Control

Topic areas Limits, Memory

Default value 0 (no limit)

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

MAXMIPSOL

Description Branch and Bound: This specifies a limit on the number of integer solutions to be found by the Optimizer. It is possible that during optimization the Optimizer will find the same objective solution from different nodes. However, MAXMIPSOL refers to the total number of integer solutions found, and not necessarily the number of distinct solutions.

Type Integer

Category Control

Topic areas Limits, Solution

Default value 0

Note Setting MAXMIPSOL=1 can alter the solution path as this will put the emphasis on finding any feasible solution by triggering additional heuristics.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

MAXMIPTASKS

Description Branch-and-Bound: The maximum number of tasks to run in parallel during a MIP solve.

Type Integer

Category Control

Topic areas Limits, Determinism, Parallel

Values

-1	Task limit determined automatically from MIPTHREADS.
>0	Fixed task limit.

Default value -1

Note The MIP solver will create smaller tasks from individual active nodes or based on local search heuristics. These are tasks that will be executed in parallel by the number of threads set by MIPTHREADS.

Note If MAXMIPTASKS is set to a fixed, positive value, the branch-and-bound tree nodes will always be solved in the same deterministic way, independent of the actual number of executing threads implied by MIPTHREADS.

How a MIP is solved will still depend on the number of threads used for solving the continuous relaxation and therefore on the settings for the controls BARTHREADS, BARCORES, DUALTHREADS and CONCURRENTTHREADS).

To obtain a MIP solve that is completely independent of the number of threads, it is sufficient to set MAXMIPTASKS, FORCEPARALLELDUAL, BARTHREADS and BARCORES. The concurrent LP solver should be avoided in this case.

Note While you can set this control to large value, the implementation will limit the number of tasks for performance reasons. This limit is currently 32 on 32-bit platforms and 256 on 64-bit platforms.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPTHREADS, THREADS, DUALTHREADS, BARTHREADS, BARCORES, CONCURRENTTHREADS, FORCEPARALLELDUAL.

MAXNODE

Description Branch and Bound: The maximum number of nodes that will be explored.

Type Integer

Category Control

Topic areas Limits, Branch and Bound Search

Default value `2147483647`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

MAXPAGELINES

Description Number of lines between page breaks in printable output.

Type Integer

Category Control

Topic area Logging

Default value `23`

Affects routines `XPRSwriteprtsol` (`WRITEPRTSOL`).

MAXSCALEFACTOR

Description This determines the maximum scaling factor that can be applied during scaling. The maximum is provided as an exponent of a power of 2.

Type Integer

Category Control

Topic areas Numerics, Presolve

Values `0-64` The maximum is provided an exponent of a power of 2.

Default value `64`

Affects routines `XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSreadprob` (`READPROB`), `XPRSscale` (`SCALE`).

See also `SCALING`.

MAXSTALLTIME

Description The maximum time in seconds that the Optimizer will continue to search for improving solution after finding a new incumbent.

Type Double

Category Control

Topic area Limits

Values	0	No stall time limit.
	>0	If an integer solution has been found, stop MIP search after the given number of seconds without a new incumbent. No effect as long as no solution was found.

Default value 0

Note This control changed type with Xpress Optimizer version 9.0 (from integer to double). It is strongly recommended to adapt all access calls to this control accordingly.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also TIMELIMIT, SOLTIMELIMIT, MAXMIPSOL.

MAXTIME

Description *This parameter is deprecated and will be removed in a future release.*
The maximum time in seconds that the Optimizer will run before it terminates, including the problem setup time and solution time. For MIP problems, this is the total time taken to solve all nodes.

Type Integer

Category Control

Topic area Limits

Values	0	No time limit.
	n>0	If an integer solution has been found, stop MIP search after *n* seconds, otherwise continue until an integer solution is finally found.
	n<0	Stop in LP or MIP search after *n* seconds.

Default value 0

Note This control has been deprecated with Xpress Optimizer version 9.0; the double control TIMELIMIT should be used instead. Note that the meaning of positive values differs between TIMELIMIT and MAXTIME. The functionality of positive MAXTIME values is covered by SOLTIMELIMIT. When both controls are set, the new TIMELIMIT control will take precedence.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also TIMELIMIT, SOLTIMELIMIT.

MCFCUTSTRATEGY

Description Level of Multi-Commodity Flow (MCF) cutting planes separation: This specifies how much aggresively MCF cuts should be separated. If the separation of MCF cuts is enabled, Xpress will try to detect a MCF network structure in the problem and, if such a structure is identified, it will separate specific cutting planes exploiting the identified network.

Type Integer

Category Control

Topic area Cuts

Values	-1	Automatic - let the Optimizer decide.
	0	Separation of MCF cuts disabled.
	1	Moderate separation of MCF cuts.
	2	Aggressive separation of MCF cuts.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also CUTSTRATEGY,COVERCUTS.

MIPABSCUTOFF

Description Branch and Bound: If the user knows that they are interested only in values of the objective function which are better than some value, this can be assigned to MIPABSCUTOFF. This allows the Optimizer to ignore solving any nodes which may yield worse objective values, saving solution time. When a MIP solution is found a new cut off value is calculated and the value can be obtained from the CURRMIPCUTOFF attribute. The value of CURRMIPCUTOFF is calculated using the MIPRELCUTOFF and MIPADDCUTOFF controls.

Type Double

Category Control

Topic area Limits

Default value 1.0E+40 (for minimization problems); -1.0E+40 (for maximization problems).

Note MIPABSCUTOFF can also be used to stop the dual algorithm.

Note Any value outside the interval [-1e40,1e40] will be clipped to that interval.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also MIPRELCUTOFF, MIPADDCUTOFF.

MIPABSGAPNOTIFY

Description Branch and bound: if the gapnotify callback has been set using XPRSaddcbgapnotify, then this callback will be triggered during the tree search when the absolute gap reaches or passes the value you set of the MIPRELGAPNOTIFY control.

Type Double

Category Control

Topic area Callback

Default value -1.0

Affects routines XPRSaddcbgapnotify, XPRSmipoptimize (MIPOPTIMIZE).

See also MIPRELGAPNOTIFY, MIPABSGAPNOTIFYOBJ, MIPABSGAPNOTIFYBOUND

MIPABSGAPNOTIFYBOUND

Description Branch and bound: if the `gapnotify` callback has been set using `XPRSaddcbgapnotify`, then this callback will be triggered during the tree search when the best bound reaches or passes the value you set of the `MIPRELGAPNOTIFYBOUND` control.

Type Double

Category Control

Topic area Callback

Default value `1.0E+20` (for minimization problems); `-1.0E+20` (for maximization problems)

Affects routines `XPRSaddcbgapnotify`, `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPRELGAPNOTIFY`, `MIPABSGAPNOTIFYOBJ`, `MIPABSGAPNOTIFY`

MIPABSGAPNOTIFYOBJ

Description Branch and bound: if the `gapnotify` callback has been set using `XPRSaddcbgapnotify`, then this callback will be triggered during the tree search when the best solution value reaches or passes the value you set of the `MIPRELGAPNOTIFYOBJ` control.

Type Double

Category Control

Topic area Callback

Default value `-1.0E+20` (for minimization problems); `1.0E+20` (for maximization problems)

Affects routines `XPRSaddcbgapnotify`, `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPRELGAPNOTIFY`, `MIPABSGAPNOTIFY`, `MIPABSGAPNOTIFYBOUND`

MIPABSSTOP

Description Branch and Bound: The absolute tolerance determining whether the tree search will continue or not. It will terminate if

$$|\texttt{MIPOBJVAL} - \texttt{BESTBOUND}| \leqslant \texttt{MIPABSSTOP}$$

where `MIPOBJVAL` is the value of the best solution's objective function, and `BESTBOUND` is the current best solution bound. For example, to stop the tree search when a MIP solution has been found and the Optimizer can guarantee it is within 100 of the optimal solution, set `MIPABSSTOP` to 100.

Type Double

Category Control

Topic area Limits

Default value `0.0`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPRELSTOP`, `MIPADDCUTOFF`.

MIPADDCUTOFF

Description Branch and Bound: The amount to add to the objective function of the best integer solution found to give the new CURRMIPCUTOFF. Once an integer solution has been found whose objective function is equal to or better than CURRMIPCUTOFF, improvements on this value may not be interesting unless they are better by at least a certain amount. If MIPADDCUTOFF is nonzero, it will be added to CURRMIPCUTOFF each time an integer solution is found which is better than this new value. This cuts off sections of the tree whose solutions would not represent substantial improvements in the objective function, saving processor time. The control MIPABSSTOP provides a similar function but works in a different way.

Type Double

Category Control

Topic areas Limits, Branch and Bound Search

Default value -1.0E-05

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPRELCUTOFF, MIPABSSTOP, MIPABSCUTOFF.

MIPCOMPONENTS

Description Determines whether disconnected components in a MIP should be solved as separate MIPs. There can be significant performence benefits from solving disconnected components individual instead of being part of the main branch-and-bound search.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

-1	Automatic - let the solver decide.
0	Disable solving disconnected components separately.
1	Solve disconnected components separately.

Default value -1

Note If there are no constraints linking two variables, either directly or indirectly through other variables, they are said to belong to two separate disconnected components. When a problem contains disconnected components of signficant size, it can be advantageous to solve each component as a separate MIP. When significant disconnected components are detected in the problem, the solver will switch to a different solve mode where each component is solved separately. This switch will happen after the root node processing has completed and when the solve is about to enter the branch-and-bound search.

Note Solving disconnected components separately is not compatible with many callbacks that can be used for modifying the branch-and-bound search. Setting any of the following callbacks will automatically disable the separate solving of disconnected components: XPRSaddcboptnode, XPRSaddcbprenode, XPRSaddcbcutmgr, XPRSaddcbchgbranch, XPRSaddcbchgbranchobject

Note Solving disconnected components separately is not compatible with concurrent MIP solves. If concurrent MIP solves has been turned off, disconnected components will be solved as part of the standard branch-and-bound search in each concurrent solve.

Note Disabling MIP dual reductions through `MIPDUALREDUCTIONS` will also disable the separate solve of disconnected components.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRECOMPONENTS`, `MIPCONCURRENTSOLVES`, `XPRSaddcboptnode`, `XPRSaddcbprenode`, `XPRSaddcbcutmgr`, `XPRSaddcbchgbranch`, `XPRSaddcbchgbranchobject`

MIPCONCURRENTNODES

Description Sets the node limit for when a winning solve is selected when concurrent MIP solves are enabled. When multiple MIP solves are started, they each run up to the `MIPCONCURRENTNODES` node limit and only one winning solve is selected for contuinuing the search with.

Type Integer

Category Control

Topic areas Branch and Bound Search, Limits

Values

`-1`	Automatic - let the solver decide on a node limit.
`>0`	Number of nodes each concurrent solve should complete before a winner is selected.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPCONCURRENTSOLVES`

MIPCONCURRENTSOLVES

Description Selects the number of concurrent solves to start for a MIP. Each solve will use a unique random seed for its random number generator, but will otherwise apply the same user controls. The first concurrent solve to complete will have solved the MIP and all the concurrent solves will be terminated at this point. Using concurrent solves can be advantageous when a MIP displays a high level of *performance variability*.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

`-1`	Enabled. The number of concurrent solves depends on `MIPTHREADS`.
`0, 1`	Disabled
`n>1`	Enabled. The number of concurrent solves to start is given by `n`.

Default value `0`

Note A node limit is imposed on each concurrent solve, through `MIPCONCURRENTNODES`. When a concurrent solve reaches this node limit, it will be suspended until all concurrent solves have reached the limit. At this point a winner will be declared, based on which solve made the most progress towards optimality and only the winning solve will continue, using all threading resources. If a concurrent solve completes its MIP search before reaching the node limit, all solves will be stopped.

Note Concurrent solves are not compatible with many callbacks that can be used for modifying the branch-and-bound search. Setting any of the following callbacks will automatically disable concurrent solves: XPRSaddcboptnode, XPRSaddcbprenode, XPRSaddcbcutmgr, XPRSaddcbchgbranch, XPRSaddcbchgbranchobject

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPCONCURRENTNODES, XPRSaddcboptnode, XPRSaddcbprenode, XPRSaddcbcutmgr, XPRSaddcbchgbranch, XPRSaddcbchgbranchobject

MIPDUALREDUCTIONS

Description Branch and Bound: Limits operations that can reduce the MIP solution space.

Type Integer

Category Control

Topic area Presolve

Values

2	Allow dual reductions on continuous variables only.
1	Allow all dual reductions.
0	Prevent all dual reductions.

Default value 1

Note The MIPDUALREDUCTIONS control, when set to a value different from 1 will adjust the values of other controls in order to prevent MIP solver operations that can result in the removal of dominated solutions. For example, dual reductions during preprocessing attempts to remove dominated solutions based on objective arguments, assuming that all constraints are known to the Optimizer. If a problem is detected to have symmetries, the solver might also remove some symmetrical solutions from the search space. In both cases, the set of feasible MIP solutions might be reduced. With default settings, it is only guaranteed that at least one optimal solution remains.

Note When attempting to collect the *n*-best solutions, it is recommended to set MIPDUALREDUCTIONS=2. This will ensure that the only solutions missed by the enumeration are those that only differ from an existing solution in the values of the continuous variables.

Note Advanced users that maintain external constraints, which are applied dynamically to the problem using callbacks during a branch-and-bound solve, it is recommended to set MIPDUALREDUCTIONS=0. This ensures that any solution to the original problem that satisfies all of the user's external constraints maps to a feasible solution in the presolved space.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPPRESOLVE, PRECOMPONENTS, PRESOLVEOPS, SYMMETRY.

MIPFRACREDUCE

Description Branch and Bound: Specifies how often the optimizer should run a heuristic to reduce the number of fractional integer variables in the node LP solutions.

Type Integer

Category Control

Topic areas Branch and Bound Search, Root

Values

-1	Automatic.
0	Disabled.
1	Run before and after cutting on the root node.
2	Run also during root cutting.
3	Run also during the tree search.

Default value -1

Note This heuristic is only applicable to problems that are dual degenerate. These are problems that contain multiple solutions with identical objective function value. The more dual degenerate a problem is, the more likely it will be for this heuristic to have an improving effect.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

MIPKAPPAFREQ

Description Branch and Bound: Specifies how frequently the basis condition number (also known as kappa) should be calculated during the branch-and-bound search.

Type Integer

Category Control

Topic areas Branch and Bound Search, Numerics

Values

0	Do not calculate condition numbers.
1	Calculate conditions numbers on every node, including after each round of root cutting.
n>1	Calculate a condition number once per node of every n'th level of the branch-and-bound tree.

Default value 0

Note The condition number is calculated as the norm of the basis matrix multiplied by the norm of its inverse. This uses the Froebenius norm.

Note A summary will be printed at the end of the solve, summarizing the collected condition numbers collected:

Statistic	Meaning
Nodes kappa stable	No. of stable sampled nodes (kappa < 10^7)
Nodes kappa suspicius	No. of suspicious sampled nodes ($10^7 \leqslant$ kappa < 10^{10})
Nodes kappa unstable	No. of unstable sampled nodes ($10^{10} \leqslant$ kappa < 10^{13})
Nodes kappa ill-posed	No. of ill-posed sampled nodes ($10^{13} \leqslant$ kappa)
Largest kappa seen	The largest condition number calculated through all sampled nodes.
Kappa attention level	A measure of how ill-posed the problem is (between 0 and 1).

Affects routines ATTENTIONLEVEL, MAXKAPPA, XPRSmipoptimize (MIPOPTIMIZE).

See also .

MIPLOG

Description	MIP log print control.
Type	Integer
Category	Control
Topic area	Logging

Values

-n	Print out summary log at each n^{th} node.
0	No printout during MIP tree search.
1	Only print out summary statement at the end.
2	Print out detailed log at all solutions found.
3	Print out detailed log at each node.

Default value -100

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also A.9.

MIPPRESOLVE

Description Branch and Bound: Type of integer processing to be performed. If set to 0, no processing will be performed. This is a bit-vector control (see Section 9.2).

Type Integer

Category Control

Topic areas Presolve, Branch and Bound Search, Bit-vector

Values

Bit	Meaning
0	Reduced cost fixing will be performed at each node. This can simplify the node before it is solved, by deducing that certain variables' values can be fixed based on additional bounds imposed on other variables at this node.
1	Primal reductions will be performed at each node. Uses constraints of the node to tighten the range of variables, often resulting in fixing their values. This greatly simplifies the problem and may even determine optimality or infeasibility of the node before the simplex method commences.
2	[Unused] This bit is no longer used to control probing. Refer to the integer control PREPROBING for setting probing level during presolve.
3	If node preprocessing is allowed to change bounds on continuous columns.
4	Dual reductions will be performed at each node.
5	Allow global (non-bound) tightening of the problem during the tree search.
6	The objective function will be used to find reductions at each node.
7	[Unused] This bit is no longer used to control restarts. Refer to the integer control MIPRESTART for disabling tree restarts.
8	Allow that symmetry is used to presolve the node problem.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also 5.3, PRESOLVE, PRESOLVEOPS, PREPROBING.

MIPRAMPUP

Description Controls the strategy used by the parallel MIP solver during the ramp-up phase of a branch-and-bound tree search.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

`-1`	Automatically determined.
`0`	No special treatment during the ramp-up phase. Always run with the maximal number of tasks.
`1`	Limit the number of tasks until the initial dives have completed.

Default value -1

Note The branch-and-bound tree search starts from the single root node, and only through branching on this root node and the resulting child nodes, are enough active nodes created to produce sufficient tasks to keep all MIP workers busy. This is referred to as the ramp-up phase of a parallel MIP.

In a typical MIP solve, the solutions found during the initial dives will typically provide a significant improvement over the root heuristic solutions. It can therefore be advantageous to let these initial dives run as fast as possible, by limiting resource contention. This can be accomplished by restricting the number of parallel tasks and thereby reducing the memory bus contention. The `MIPRAMPUP` control can be used to turn this initial task restriction of a parallel MIP solve on or off.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPTHREADS`, `MAXMIPTASKS`.

MIPREFINEITERLIMIT

Description This defines an effort limit expressed as simplex iterations for the MIP solution refiner. The limit is per reoptimizations in the MIP refiner.

Type Integer

Category Control

Topic areas Solution Refinement, Limits

Default value `-1` — determined automatically.

Affects routines `XPRSrefinemipsol` (`REFINEMIPSOL`).

MIPRELCUTOFF

Description Branch and Bound: Percentage of the incumbent value to be added to the value of the objective function when an integer solution is found, to give the new value of `CURRMIPCUTOFF`. The effect is to cut off the search in parts of the tree whose best possible objective function would not be substantially better than the current solution. The control `MIPRELSTOP` provides a similar functionality but works in a different way.

Type Double

Category Control

Topic area Limits

Default value `1.0E-04`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPABSCUTOFF`, `MIPADDCUTOFF`, `MIPRELSTOP`.

MIPRELGAPNOTIFY

Description Branch and bound: if the `gapnotify` callback has been set using `XPRSaddcbgapnotify`, then this callback will be triggered during the branch and bound tree search when the relative gap reaches or passes the value you set of the `MIPRELGAPNOTIFY` control.

Type Double

Category Control

Topic area Callback

Default value `-1.0`

Affects routines `XPRSaddcbgapnotify`, `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPABSGAPNOTIFY`, `MIPABSGAPNOTIFYOBJ`, `MIPABSGAPNOTIFYBOUND`

MIPRELSTOP

Description Branch and Bound: This determines when the branch and bound tree search will terminate. Branch and bound tree search will stop if:

$|\texttt{MIPOBJVAL} - \texttt{BESTBOUND}| \leqslant \texttt{MIPRELSTOP} \times \max(|\texttt{BESTBOUND}|, |\texttt{MIPOBJVAL}|)$

where `MIPOBJVAL` is the value of the best solution's objective function and `BESTBOUND` is the current best solution bound. For example, to stop the tree search when a MIP solution has been found and the Optimizer can guarantee it is within 5% of the optimal solution, set `MIPRELSTOP` to 0.05.

Type Double

Category Control

Topic area Limits

Default value `0.0001`

Note This control is a stopping criteria only and different values of the control will not affect the solution path before termination. Unlike other stopping criteria, like time and node count, termination on `MIPRELSTOP` will cause the final solution to be declared optimal and the problem to be returned to its original state.

Note Tolerances, such as `MIPRELCUTOFF` and `MIPABSCUTOFF`, determine how much the objective value of a new MIP solution has to differ from the incumbent for it to be accepted. These controls therefore also influence the final gap at the end of a MIP solve.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPABSSTOP`, `MIPRELCUTOFF`.

MIPRESTART

Description Branch and Bound: controls strategy for in-tree restarts.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

-1	Determined automatically (XPRS_MIPRESTART_DEFAULT).
0	Disable in-tree restarts (XPRS_MIPRESTART_OFF).
1	Allow in-tree restarts at normal aggressiveness (XPRS_MIPRESTART_MODERATE).
2	Allow in-tree restarts at higher aggressiveness (more likely to trigger a restart) (XPRS_MIPRESTART_AGGRESSIVE).

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPRESTARTGAPTHRESHOLD, MIPRESTARTFACTOR

MIPRESTARTGAPTHRESHOLD

Description Branch and Bound: Initial gap threshold to delay in-tree restart. The restart is delayed initially if the gap, given as a fraction between 0 and 1, is below this threshold. The optimizer adjusts the threshold every time a restart is delayed. Note that there are other criteria that can delay or prevent a restart.

Type Double

Category Control

Topic area Branch and Bound Search

Default value 0.02

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPRESTART, MIPRESTARTFACTOR

MIPRESTARTFACTOR

Description Branch and Bound: Fine tune initial conditions to trigger an in-tree restart. Use a value > 1 to increase the aggressiveness with which the Optimizer restarts. Use a value < 1 to relax the aggressiveness with which the Optimizer restarts. Note that this control does not affect the initial condition on the gap, which must be set separately.

Type Double

Category Control

Topic area Branch and Bound Search

Default value 1.0

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPRESTART, MIPRESTARTGAPTHRESHOLD

MIPTERMINATIONMETHOD

Description Branch and Bound: How a MIP solve should be stopped on early termination when there are still active tasks in the system. This can happen when, for example, a time or node limit is reached.

Type Integer

Category Control

Topic areas Branch and Bound Search, Limits

Values

`0`	Terminate tasks at the earliest opportunity. This can result in some unfinished node solves being discarded, although never integer solutions.
`1`	Allow tasks to complete their current work but prevent new tasks from being started.

Default value `0`

Note With `MIPTERMINATIONMETHOD=0`, termination will be quick but the returned state of the MIP solve will not include any work done by interrupted tasks. In particular, it is possible that some user callbacks (not *intsol* or *preintsol*) will have been fired for nodes that are discarded at termination. A user program that relies on the firing of callbacks being completely deterministic should therefore set `MIPTERMINATIONMETHOD=1`, which will produce a slower termination, but guaranteed deterministic firing of all user callbacks.

Note Irrespective of the choice of `MIPTERMINATIONMETHOD`, a MIP solve will always be returned in a deterministic state when `DETERMINISTIC=1`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `DETERMINISTIC`, `MAXMIPTASKS`, `MIPTHREADS`, `THREADS`.

MIPTHREADS

Description If set to a positive integer it determines the number of threads implemented to run the parallel MIP code. If `MIPTHREADS` is set to the default value (`-1`), the `THREADS` control will determine the number of threads used.

Type Integer

Category Control

Topic areas Branch and Bound Search, Parallel

Default value `-1` (determined by the `THREADS` control)

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `DETERMINISTIC`, `MAXMIPTASKS`, `HEURTHREADS`, `THREADS`.

MIPTOL

Description Branch and Bound: This is the tolerance within which a decision variable's value is considered to be integral.

Type Double

Category Control

Topic area Tolerances

Default value `5.0E-06`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

MIPTOLTARGET

Description Target MIPTOL value used by the automatic MIP solution refiner as defined by `REFINEOPS`. Negative and zero values are ignored.

Type Double

Category Control

Topic areas Solution Refinement, Tolerances

Default value `0.0`

Note Refining solutions to match the miptoltarget can influence and worsen their objective value in case the previous objective could only be achieved through slight infeasibilities.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

MIQCPALG

Description This control determines which algorithm is to be used to solve mixed integer quadratic constrained and mixed integer second order cone problems.

Type Integer

Category Control

Topic areas Branch and Bound Search, Quadratic

Values

`-1`	Determined automatically.
`0`	Use the barrier algorithm in the branch and bound algorithm.
`1`	Use outer approximations in the branch and bound algorithm.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

MPS18COMPATIBLE

Description Provides compatibility of MPS file output for older MPS readers.

Type Integer

Category Control

Topic area File IO

Values

`Bit 0`	Do not write objective sense (`OBJSENSE` section).
`Bit 1`	Fixed binaries are written as fixed only (unless used as a base variable for an indicator constraint).

Default value `0`

Affects routines `XPRSwriteprob` (`WRITEPROB`)

MPSBOUNDNAME

Description When reading an MPS file, this control determines which entries from the `BOUNDS` section will be read. As with all string controls, this is of length 64 characters plus a null terminator, `\0`.

Type String

Category Control

Topic area File IO

Default value 64 blanks

Affects routines `XPRSreadprob` (`READPROB`).

MPSECHO

Description Determines whether comments in MPS matrix files are to be printed out during matrix input.

Type Integer

Category Control

Topic area File IO

Values

`0`	MPS comments are *not* to be echoed.
`1`	MPS comments *are* to be echoed.

Default value `0`

Affects routines `XPRSreadprob` (`READPROB`).

MPSFORMAT

Description Specifies the format of MPS files.

Type Integer

Category Control

Topic area File IO

Values

`-1`	To determine the file type automatically.
`0`	For fixed format.
`1`	If MPS files are assumed to be in free format by input.

Default value `1`

Note Setting MPSFORMAT to 0 or -1 disables XSLPreadprob in case Xpress NonLinear is used.

Affects routines `XPRSalter` (`ALTER`), `XPRSreadbasis` (`READBASIS`), `XPRSreadprob` (`READPROB`).

MPSOBJNAME

Description When reading an MPS file, this control determines which neutral row will be read as the objective function. If this control is set when reading a multi-objective MPS file, only the named objective will be read; all other objectives will be ignored. As with all string controls, this is of length 64 characters plus a null terminator, `\0`.

Type String

Category Control

Topic area File IO

Default value 64 blanks

Note To change the name of the objective row when written to an MPS file, use `XPRSaddnames`, passing `XPRS_NAMES_OBJECTIVE` in the `type` parameter.

Affects routines `XPRSreadprob` (`READPROB`).

MPSRANGENAME

Description When reading an MPS file, this control determines which entries from the `RANGES` section will be read. As with all string controls, this is of length 64 characters plus a null terminator, `\0`.

Type String

Category Control

Topic area File IO

Default value 64 blanks

Affects routines `XPRSreadprob` (`READPROB`).

MPSRHSNAME

Description When reading an MPS file, this control determines which entries from the `RHS` section will be read. As with all string controls, this is of length 64 characters plus a null terminator, `\0`.

Type String

Category Control

Topic area File IO

Default value 64 blanks

Affects routines `XPRSreadprob` (`READPROB`).

MULTIOBJOPS

Description Modifies the behaviour of the optimizer when solving multi-objective problems.

Type Integer

Category Control

Topic areas Multiobjective, Bit-vector

Values

Bit	Meaning
0	`XPRS_MULTIOBJOPS_ENABLED` Multi-objective enabled. If this bit is not set, multi-objective problems will treated as single-objective problems, and only objective `0` will be optimized.
1	`XPRS_MULTIOBJOPS_PRESOLVE` Apply multi-objective modifications during presolve. If this bit is not set, the original problem will be modified when solving each subsequent objective, and these modifications will remain in the problem after the solve has completed.
2	`XPRS_MULTIOBJOPS_RCFIXING` Reduced cost fixing. If this bit is set, optimality of earlier objectives will be preserved by fixing all non-basic variables with non-zero reduced costs to their bounds. If not set, optimality of earlier objectives will be preserved by adding constraints to the problem.

Default value 7 (all bits are set, see Section 9.2 for bit-vector controls)

Affects routines `XPRSoptimize`.

MULTIOBJLOG

Description Log level for multi-objective optimization.

Type Integer

Category Control

Topic areas Multiobjective, Logging

Values

0	No logging.
1	Print a summary of each problem that is solved as part of the multi-objective optimization.
2	In addition to summaries, print messages produced by each solve at the level determined by `OUTPUTLOG`.

Default value 2

Affects routines `XPRSoptimize`.

MUTEXCALLBACKS

Description Branch and Bound: This determines whether the callback routines are mutexed from within the optimizer.

Type Integer

Category Control

Topic areas Branch and Bound Search, Callback

Values

0	Callbacks are not mutexed.
1	Callbacks are mutexed.

Default value 1

Note If the users' callbacks take a significant amount of time it may be preferable not to mutex the callbacks. In this case the user must ensure that their callbacks are threadsafe.

Affects routines `XPRSaddcboptnode`, `XPRSaddcbinfnode`, `XPRSaddcbintsol`, `XPRSaddcbnodecutoff`, `XPRSaddcbprenode`.

NETSTALLLIMIT

Description Limit the number of degenerate pivots of the network simplex algorithm, before switching to either primal or dual simplex, depending on ALGAFTERNETWORK.

Type Integer

Category Control

Topic areas Simplex, Limits

Values

-1	Automatically determined limit
0	No limit.
n>0	Limit to *n* network simplex iterations.

Default value -1

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE) (when network simplex is used).

See also ALGAFTERNETWORK

NODEPROBINGEFFORT

Description Adjusts the overall level of node probing.

Type Double

Category Control

Topic areas Branch and Bound Search, Presolve

Default value 1.0

Note NODEPROBINGEFFORT is used as a multiplier on the default amount of work node probing should do. Setting the control to zero disables node probing.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

NODESELECTION

Description Branch and Bound: This determines which nodes will be considered for solution once the current node has been solved.

Type Integer

Category Control

Topic area Branch and Bound Search

Values

1	*Local first*: Choose between descendant and sibling nodes if available; choose from all outstanding nodes otherwise.
2	*Best first*: Choose from all outstanding nodes.
3	*Local depth first*: Choose between descendant and sibling nodes if available; choose from the deepest nodes otherwise.
4	*Best first, then local first*: Best first is used for the first BREADTHFIRST nodes, after which local first is used.
5	*Pure depth first*: Choose from the deepest outstanding nodes.

Default value Dependent on the matrix characteristics.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

NUMERICALEMPHASIS

Description How much emphasis to place on numerical stability instead of solve speed.

Type Integer

Category Control

Topic area Numerics

Values

`-1`	Automatic. The emphasis might be influenced by the setting of other controls.
`0`	Emphasize speed.
`1`	Mild emphasis on numerical stability.
`2`	Medium emphasis on numerical stability.
`3`	Strong emphasis on numerical stability.

Default value -1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

OBJSCALEFACTOR

Description Custom objective scaling factor, expressed as a power of 2. When set, it overwrites the automatic objective scaling factor. A value of 0 means no objective scaling. This control is applied for the full solve, and is independent of any extra scaling that may occur specifically for the barrier or simplex solvers. As it is a power of 2, to scale by 16, set the value of the control to 4.

Type Double

Category Control

Topic area Numerics

Default value `0`

Note This control does not apply to problems solved with the global solver or to problems with multiple objectives.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

OPTIMALITYTOL

Description Simplex: This is the zero tolerance for reduced costs. On each iteration, the simplex method searches for a variable to enter the basis which has a negative reduced cost. The candidates are only those variables which have reduced costs less than the negative value of `OPTIMALITYTOL`.

Type Double

Category Control

Topic area Tolerances

Default value `1.0E-06`

Affects routines `XPRSgetinfeas`, `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

OPTIMALITYTOLTARGET

Description This specifies the target optimality tolerance for the solution refiner.

Type Double

Category Control

Topic areas Solution Refinement, Tolerances

Default value `0` — use the value specified by `OPTIMALITYTOL`.

Note Zero and negative values are ignored, and the value of `OPTIMALITYTOL` is used.

Note Use very small values like 1e-100 to state the refinement should continue as long as an improvement is made. Use very large values like 1e+100 to disable only this aspect of the refiner.

Note Refining solutions to match the optimalitytoltarget can influence and increase their infeasibility in case the previous feasibility could only be achieved through slight dual violations.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `REFINEOPS`, `LPREFINEITERLIMIT`, `FEASTOLTARGET`.

OUTPUTCONTROLS

Description This control toggles the printing of all control settings at the beginning of the search. This includes the printing of controls that have been explicitly assigned to their default value. All unset controls are omitted as they keep their default value.

Type Integer

Category Control

Topic area Logging

Values

`0`	Turn off printing of user-specified control settings.
`1`	Print controls.

Default value `1`

Note Setting `OUTPUTCONTROLS` to 0 has no effect on the function `XPRSdumpcontrols`

OUTPUTLOG

Description This controls the level of output produced by the Optimizer during optimization. In the Console Optimizer, `OUTPUTLOG` controls which messages are sent to the screen (`stdout`). When using the Optimizer library, no output is sent to the screen. If the user wishes output to be displayed, they must define a callback function and print messages to the screen themselves. In this case, `OUTPUTLOG` controls which messages are sent to the user output callback.

Type Integer

Category Control

Topic area	Logging	
Values	0	Turn all output off. Use XPRS_OUTPUTLOG_NO_OUTPUT from xprs.h.
	1	Print all messages. Use XPRS_OUTPUTLOG_FULL_OUTPUT from xprs.h.
	3	Print error and warning messages. Use XPRS_OUTPUTLOG_ERRORS_AND_WARNINGS from xprs.h.
	4	Print error messages only. Use XPRS_OUTPUTLOG_ERRORS from xprs.h.
Default value	1	
Affects routines	XPRSaddcbmessage, XPRSsetlogfile.	

OUTPUTMASK

Description Mask to restrict the row and column names written to file. As with all string controls, this is of length 64 characters plus a null terminator, \0.

Type String

Category Control

Topic area File IO

Default value 64 '?'s

Affects routines XPRSwritesol (WRITESOL).

OUTPUTTOL

Description Zero tolerance on print values.

Type Double

Category Control

Topic areas File IO, Tolerances

Default value 1.0E-05

Affects routines XPRSwriteprtsol (WRITEPRTSOL), XPRSwritesol (WRITESOL).

PENALTY

Description Minimum absolute penalty variable coefficient. BIGM and PENALTY are set by the input routine (XPRSreadprob (READPROB)) but may be reset by the user prior to XPRSlpoptimize (LPOPTIMIZE).

Type Double

Category Control

Topic area Presolve

Default value Dependent on the matrix characteristics.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

PIVOTTOL

Description Simplex: The zero tolerance for matrix elements. On each iteration, the simplex method seeks a nonzero matrix element to pivot on. Any element with absolute value less than `PIVOTTOL` is treated as zero for this purpose.

Type Double

Category Control

Topic areas Simplex, Tolerances

Default value `1.0E-09`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSpivot`.

PPFACTOR

Description The partial pricing candidate list sizing parameter.

Type Double

Category Control

Topic area Simplex

Default value `1.0`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

PREANALYTICCENTER

Description Determines if analytic centers should be computed and used for variable fixing and the generation of alternative reduced costs (-1: Auto 0: Off, 1: Fixing, 2: Redcost, 3: Both)

Type Integer

Category Control

Topic areas Presolve, Barrier

Values

`-1`	Automatic.
`0`	Disable analytic center presolving.
`1`	Use analytic center for variable fixing only.
`2`	Use analytic center for reduced cost computation only.
`3`	Use analytic centers for both, variable fixing and reduced cost computation.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

PREBASISRED

Description	Determines if a lattice basis reduction algorithm should be attempted as part of presolve
Type	Integer
Category	Control
Topic area	Presolve
Values	`-1` Automatic. `0` Disable basis reduction. `1` Enable basis reduction.
Default value	`0`
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`).

PREBNDREDCONE

Description	Determines if second order cone constraints should be used for inferring bound reductions on variables when solving a MIP.
Type	Integer
Category	Control
Topic area	Presolve
Values	`-1` Automatic. `0` Disable bound reductions from second order cone constraints. `1` Enable bound reductions from second order cone constraints.
Default value	`-1`
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`PREBNDREDQUAD`, `MIQCPALG`.

PREBNDREDQUAD

Description	Determines if convex quadratic constraints should be used for inferring bound reductions on variables when solving a MIP.
Type	Integer
Category	Control
Topic areas	Presolve, Quadratic
Values	`-1` Automatic. `0` Disable bound reductions from quadratic constraints. `1` Enable bound reductions from quadratic constraints.
Default value	`-1`
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`PREBNDREDCONE`, `MIQCPALG`.

PRECLIQUESTRATEGY

Description Determines how much effort to spend on clique covers in presolve.

Type Integer

Category Control

Topic area Presolve

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

PRECOEFELIM

Description Presolve: Specifies whether the optimizer should attempt to recombine constraints in order to reduce the number of non zero coefficients when presolving a mixed integer problem.

Type Integer

Category Control

Topic area Presolve

Values

`0`	Disabled.
`1`	Remove as many coefficients as possible.
`2`	Cautious eliminations. Will not perform a reduction if it might destroy problem structure useful to e.g. heuristics or cutting.

Default value `2`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PRECOMPONENTS

Description Presolve: determines whether small independent components should be detected and solved as individual subproblems during root node processing.

Type Integer

Category Control

Topic area Presolve

Values

`-1`	Automatically determined.
`0`	Disable detection of independent components.
`1`	Enable detection of independent components.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PRECOMPONENTSEFFORT

Description Presolve: adjusts the overall effort for the independent component presolver. This control affects working limits for the subproblem solving as well as thresholds when it is called. Increase to put more emphasis on component presolving.

Type Double

Category Control

Topic area Presolve

Default value 1.0

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRECOMPONENTS`.

PRECONEDECOMP

Description Presolve: decompose regular and rotated cones with more than two elements and apply Outer Approximation on the resulting components.

Type Integer

Category Control

Topic areas Presolve, Quadratic

Values

`-1`	Automatically determined.
`0`	Disable cone decomposition.
`1`	Enable cone decomposition by replacing large cones with small ones in the presolved problem.
`2`	Similar to 1, plus decomposition is enabled even if the cone variable is fixed.
`3`	Cones are decomposed within the Outer Approximation domain only, i.e., the problem maintains the original cones.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PRECONFIGURATION

Description MIP Presolve: determines whether binary rows with only few repeating coefficients should be reformulated. The reformulation enumerates the extremal feasible configurations of a row and introduces new columns and rows to model the choice between these extremal configurations. This presolve operation can be disabled as part of the (advanced) IP reductions `PRESOLVEOPS`.

Type Integer

Category Control

Topic area Presolve

Values	-1	Automatically determined.
	0	Disable configuration presolving.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also PRESOLVE, PRESOLVEOPS.

PRECONVERTOBJTOCONS

Description Presolve: convert a linear or quadratic objective function into an objective transfer constraint

Type Integer

Category Control

Topic areas Presolve, Problem Transformation

Values	-1	Automatically determined.
	0	Disable reformulation.
	1	Move only the quadratic part of the objective into a constraint.
	2	Move both the linear and quadratic parts of the objective into a constraint.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also PRESOLVE.

Note This control is only used in MIQPs and MIQCQPs, and has no effect when used on continuous quadratic problems.

PRECONVERTSEPARABLE

Description Presolve: reformulate problems with a non-diagonal quadratic objective and/or constraints as diagonal quadratic or second-order conic constraints.

Type Integer

Category Control

Topic areas Presolve, Problem Transformation

Values	-1	Automatically determined.
	0	Disable reformulation.
	1	Enable reformulation to diagonal quadratic constraints.
	2	Enable reformulation to diagonal quadratic constraints and reduction to second-order cones.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also PRESOLVE.

Note This control is only used in MIQPs and MIQCQPs, and has no effect when used on continuous quadratic problems.

PREDOMCOL

Description Presolve: Determines the level of dominated column removal reductions to perform when presolving a mixed integer problem. Only binary columns will be checked.

Type Integer

Category Control

Topic area Presolve

Values

-1	Automatically determined.
0	Disabled.
1	Cautious strategy, limited effort looking for special structure.
2	Same as 2 but checking all candidates.
3	Includes 1 and 2 but also looks for more generic column domination.

Default value -1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PREDOMROW

Description Presolve: Determines the level of dominated row removal reductions to perform when presolving a problem.

Type Integer

Category Control

Topic area Presolve

Values

-1	Automatically determined.
0	Disabled.
1	Cautious strategy.
2	Medium strategy.
3	Aggressive strategy. All candidate row combinations will be considered.

Default value -1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PREDUPROW

Description Presolve: Determines the type of duplicate rows to look for and eliminate when presolving a problem.

Type Integer

Category Control

Topic area Presolve

Values
- `-1` Automatically determined.
- `0` Do not eliminate duplicate rows.
- `1` Eliminate only rows that are identical in all variables.
- `2` Same as option 1 plus eliminate duplicate rows with simple penalty variable expressions. (MIP only).
- `3` Same as option 2 plus eliminate duplicate rows with more complex penalty variable expressions. (MIP only).

Default value `-1`

Note Duplicate rows can also be disabled by clearing the corresponding bit of the `PRESOLVEOPS` integer control.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PREELIMQUAD

Description Presolve: Allows for elimination of quadratic variables via doubleton rows.

Type Integer

Category Control

Topic areas Presolve, Quadratic

Values
- `-1` Automatically determined.
- `0` Do not eliminate duplicate rows.
- `1` Eliminate at least one quadratic variable for each doubleton row.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PREFOLDING

Description Presolve: Determines if a folding procedure should be used to aggregate continuous columns in an equitable partition.

Type Integer

Category Control

Topic area Presolve

Values
- `-1` Automatically determined.
- `0` Disabled.
- `1` Enabled.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRESOLVE`.

PREIMPLICATIONS

Description Presolve: Determines whether to use implication structures to remove redundant rows. If implication sequences are detected, they might also be used in probing.

Type Integer

Category Control

Topic area Presolve

Values

`-1`	Automatically determined.
`0`	Do not use implications for sparsification.
`1`	Use implications to remove reduandant rows.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`, `PREPROBING`.

PRELINDEP

Description Presolve: Determines whether to check for and remove linearly dependent equality constraints when presolving a problem.

Type Integer

Category Control

Topic area Presolve

Values

`-1`	Automatically determined.
`0`	Do not check for linearly dependent equality constraints.
`1`	Check for and remove linearly dependent equality constraints.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PREOBJCUTDETECT

Description Presolve: Determines whether to check for constraints that are parallel or near parallel to a linear objective function, and which can safely be removed. This reduction applies to MIPs only.

Type Integer

Category Control

Topic area Presolve

Values

`0`	Disable check and reductions.
`1`	Enable check and reductions.

Default value `1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRESOLVE`, `PRESOLVEOPS`.

PREPERMUTE

Description This bit-vector control (see Section 9.2) specifies whether to randomly permute rows, columns and MIP entities when starting the presolve. With the default value 0, no permutation will take place.

Type Integer

Category Control

Topic areas Presolve, Problem Transformation, Bit-vector

Values

Bit	Meaning
0	Permute rows.
1	Permute columns.
2	Permute MIP entities. This bit only affects MIP problems.

Default value 0

Note Random permutations enable trying out different solution paths when solving a problem. The random seed for the permutations can be set using PREPERMUTESEED. When both PRESORT and PREPERMUTE are enabled, it will sort and then permute the problem.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also PREPERMUTESEED, PRESORT, PRESOLVE, MIPPRESOLVE.

PREPERMUTESEED

Description This control sets the seed for the pseudo-random number generator for permuting the problem when starting the presolve. This control only has effects when PREPERMUTE is enabled.

Type Integer

Category Control

Topic areas Presolve, Problem Transformation

Default value 1

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also PREPERMUTE, PRESOLVE, MIPPRESOLVE.

PREPROBING

Description Presolve: Amount of probing to perform on binary variables during presolve. This is done by fixing a binary to each of its values in turn and analyzing the implications.

Type Integer

Category Control

Topic area Presolve

Values

-1	Let the optimizer decide on the amount of probing.
0	Disabled.
+1	Light probing — only few implications will be examined.
+2	Full probing — all implications for all binaries will be examined.
+3	Full probing and repeat as long as the problem is significantly reduced.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also PRESOLVE.

PREPROTECTDUAL

Description Presolve: specifies whether the presolver should protect a given dual solution by maintaining the same level of dual feasibility. Enabling this control often results in a worse presolved model. This control only expected to be optionally enabled before calling XPRScrossoverlpsol.

Type Integer

Category Control

Topic areas Presolve, LP

Values

0	Disabled.
1	Enabled. Protect the dual solution during presolve.

Default value 0

Affects routines XPRScrossoverlpsol

PREROOTEFFORT

Description Dial for the work spent during the Pre-root parallel heuristic phase. A positive value sets a suitable work limit that is dependent on problem-characteristics. Changing the value up/or down dials the work spent in this phase up or down. This control also enables/disables Pre-root parallel heuristics.

Type Double

Category Control

Topic areas Parallel, Heuristics, Limits

Values

-2	Enable Pre-root parallel heuristics without a specific work limit for this phase. The phase will terminate if a different limit is hit, or if it runs out of heuristic work to do.
-1	Enablement of Pre-root parallel heuristics is subject to HEUREMPHASIS.
0	Disable Pre-root parallel heuristics.
x>0	Enable Pre-root parallel heuristics with a work limit dependent on problem characteristics, using x as a factor to dial this work limit up or down.

Default value -1

Note The special value of -2 does not use a phase-specific work limit at all. Make sure to use other limits to stop the Pre-root parallel phase eventually.

Note If a WORKLIMIT on the overall solution process has been set, it also applies during the pre-root parallel phase.

Note This control always sets a work limit that depends on the characteristics of the presolved problem. If you prefer to set an absolute work limit instead, consider the control PREROOTWORKLIMIT.

Affects routines XPRSoptimize (OPTIMIZE)

See also PREROOTWORKLIMIT, HEUREMPHASIS

PREROOTWORKLIMIT

Description Set an explicit work limit in work units for the Pre-root parallel heuristic phase. Any positive value also enables this phase and runs it until the PREROOTWORKLIMIT units of work are hit.

Type Double

Category Control

Topic areas Parallel, Heuristics, Limits

Values

- -1 No explicit work limit for the Pre-root parallel heuristic phase. If enabled, the work limit for this phase is controlled via PREROOTEFFORT.
- 0 Disable Pre-root parallel heuristics.
- x>0 Enable Pre-root parallel heuristics with an explicit work limit of x work units. If set, this work limit has precedence over any work limit set by PREROOTEFFORT.

Default value -1

Note A WORKLIMIT on the overall solution process also applies during pre-root parallel phase.

Note It may happen that even if this control is set, Pre-root parallel heuristics terminate before this work limit is reached. This may happen due to stalling or if the heuristics employed are not applicable to the current problem

Affects routines XPRSoptimize (OPTIMIZE)

See also PREROOTWORKLIMIT, HEUREMPHASIS

PREROOTTHREADS

Description Specifies an explicit number of threads that should be used for the Pre-root parallel heuristic phase. By default, this phase will use all threads available to the solver (as governed by the control THREADS).

Type Integer

Category Control

Topic areas Parallel, Heuristics, Memory

Values

- -1 Use all available threads.
- 0 Disable pre-root parallel heuristics.
- n>0 Use the specified number of threads, superseding the value of THREADS

Default value -1

Note Setting an explicit number of threads does not enable Pre-root parallel heuristics if they are otherwise disabled. To enable this phase, consider the controls HEUREMPHASIS, PREROOTEFFORT, or PREROOTWORKLIMIT.

Note	This control is useful to limit the peak memory usage during Pre-root parallel heuristics, by limiting the amount of problem copies necessary to conduct this phase.
Affects routines	`XPRSoptimize` (`OPTIMIZE`)
See also	`PREROOTEFFORT`, `THREADS`

PRESOLVE

Description This control determines whether presolving should be performed prior to starting the main algorithm. Presolve attempts to simplify the problem by detecting and removing redundant constraints, tightening variable bounds, etc. In some cases, infeasibility may even be determined at this stage, or the optimal solution found.

Type Integer

Category Control

Topic area Presolve

Values

`-1`	Presolve applied, but a problem will not be declared infeasible if primal infeasibilities are detected. The problem will be solved by the LP optimization algorithm, returning an infeasible solution, which can sometimes be helpful.
`0`	Presolve not applied.
`1`	Presolve applied.
`2`	Presolve applied, but redundant bounds are not removed. This can sometimes increase the efficiency of the barrier algorithm.
`3`	Presolve is applied, and bounds detected to be redundant are always removed.

Default value `1`

Note Memory for presolve is dynamically resized. If the Optimizer runs out of memory for presolve, an error message (245) is produced.
Presolve settings 2 and 3 can sometimes make the barrier solves more efficient.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also 5.3, `PRESOLVEOPS`.

PRESOLVEMAXGROW

Description Limit on how much the number of non-zero coefficients is allowed to grow during presolve, specified as a ratio of the number of non-zero coefficients in the original problem.

Type Double

Category Control

Topic areas Presolve, Limits

Default value 0.1

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

PRESOLVEOPS

Description This bit-vector control (see Section 9.2) specifies the operations which are performed during the presolve.

Type Integer

Category Control

Topic areas Presolve, Bit-vector

Values

Bit	Meaning
0	Singleton column removal.
1	Singleton row removal.
2	Forcing row removal.
3	Dual reductions.
4	Redundant row removal.
5	Duplicate column removal.
6	Duplicate row removal.
7	Strong dual reductions.
8	Variable eliminations.
9	No IP reductions.
10	No domain changes for MIP entities (e.g., semi-continuous detection or shifting integers).
11	No advanced IP reductions.
12	No eliminations on integers.
13	No reductions based on solution enumeration.
14	Linearly dependant row removal.
15	No integer variable and SOS detection.
16	No implied bounds.
17	No clique presolve.
18	No mod2 presolve.

Default value `511` (bits `0` — `8` incl. are set)

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSpresolverow`.

See also 5.3, `PRESOLVE`, `MIPPRESOLVE`.

PRESOLVEPASSES

Description Number of reduction rounds to be performed in presolve

Type Integer

Category Control

Topic areas Presolve, Limits

Default value `1`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also 5.3, `PRESOLVE`.

PRESORT

Description	This bit-vector control (see Section 9.2) specifies whether to sort rows, columns and MIP entities by their names when starting the presolve. With the default value `0`, no sorting will take place.
Type	Integer
Category	Control
Topic areas	Presolve, Problem Transformation, Bit-vector
Values	Bit — Meaning `0` — Sort rows. `1` — Sort columns. `2` — Sort MIP entities. This bit only affects MIP problems.
Default value	`0`
Note	Sorting a problem by names can help obtain the same solution path when the rows, columns or MIP entities of the problem is rearranged. It is recommended to enable all three bits when sorting a problem. When both `PRESORT` and `PREPERMUTE` are enabled, it will sort and then permute the problem.
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`PREPERMUTE`, `PRESOLVE`, `MIPPRESOLVE`.

PRICINGALG

Description	Simplex: This determines the primal simplex pricing method. It is used to select which variable enters the basis on each iteration. In general Devex pricing requires more time on each iteration, but may reduce the total number of iterations, whereas partial pricing saves time on each iteration, but may result in more iterations.
Type	Integer
Category	Control
Topic areas	LP, Simplex
Values	`-1` — Partial pricing. `0` — Determined automatically. `1` — Devex pricing. `2` — Steepest edge. `3` — Steepest edge with unit initial weights.
Default value	`0`
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`DUALGRADIENT`.

PRIMALOPS

Description Primal simplex: allows fine tuning the variable selection in the primal simplex solver.

Type Integer

Category Control

Topic areas LP, Simplex, Bit-vector

Values

Bit	Meaning
0	Use aggressive dj scaling.
1	Conventional dj scaling.
2	Use reluctant switching back to partial pricing.
3	Use dynamic switching between cheap and expensive pricing strategies.
4	Keep solving even after potential cycling is detected.

Default value -1

Note If both bits 0 and 1 are both set or unset then the dj scaling strategy is determined automatically. See Section 9.2 for bit-vector controls.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also PRICINGALG.

PRIMALPERTURB

Description The factor by which the problem will be perturbed prior to optimization by primal simplex. A value of 0.0 results in no perturbation prior to optimization.

Note the interconnection to the AUTOPERTURB control. If AUTOPERTURB is set to 1, the decision whether to perturb or not is left to the Optimizer. When the problem is automatically perturbed in primal simplex, however, the value of PRIMALPERTURB will be used for perturbation.

Type Double

Category Control

Topic areas LP, Simplex

Default value -1 — determined automatically.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also AUTOPERTURB, DUALPERTURB.

PRIMALUNSHIFT

Description Determines whether primal is allowed to call dual to unshift.

Type Integer

Category Control

Topic areas LP, Simplex

Values	0	Allow the dual algorithm to be used to unshift.
	1	Don't allow the dual algorithm to be used to unshift.

Default value 0

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRIMALOPS`, `PRICINGALG`, `DUALSTRATEGY`.

PSEUDOCOST

Description Branch and Bound: The default pseudo cost used in estimation of the degradation associated with an unexplored node in the tree search. A pseudo cost is associated with each integer decision variable and is an estimate of the amount by which the objective function will be worse if that variable is forced to an integral value.

Type Double

Category Control

Topic area Branching

Default value `0.01`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSreaddirs` (`READDIRS`).

PWLDUALREDUCTIONS

Description This parameter specifies whether dual reductions should be applied to reduce the number of columns, rows and SOS-constraints added when transforming piecewise linear objectives and constraints to MIP structs.

Type Integer

Category Control

Topic areas Piecewise Linear and General Constraints, Presolve, Problem Transformation

Values	0	Disabled. No dual reductions, add all columns, rows and SOS-constraints.
	1	Enabled. Only add neccessary columns, rows and sets, drop those implied by the objective sense.

Default value 1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPDUALREDUCTIONS`.

PWLNONCONVEXTRANSFORMATION

Description This control specifies the reformulation method for piecewise linear constraints at the beginning of the search.

Note that the chosen formulation will only be used if MIP entities are necessary but not if presolve detected that a convex reformulation is possible. Furthermore, the binary formulation will only be applied to piecewise linear constraints with bounded input variable, otherwise the SOS2-formulation will be used.

Type Integer

Category Control

Topic areas Piecewise Linear and General Constraints, Presolve, Problem Transformation

Values

-1	Automatic.
0	Use a formulation based on SOS2-constraints.
1	Use a formulation based on binary variables.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

QCCUTS

Description Branch and Bound: Limit on the number of rounds of outer approximation cuts generated for the root node, when solving a mixed integer quadratic constrained or mixed integer second order conic problem with outer approximation.

Type Integer

Category Control

Topic areas Cuts, Quadratic

Default value -1 — determined automatically.

Note This control only has an effect for problems with quadratic or second order cone constraints, and only if outer approximation has not been disabled by setting MIQCPALG to 0.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also TREEQCCUTS.

QCROOTALG

Description This control determines which algorithm is to be used to solve the root of a mixed integer quadratic constrained or mixed integer second order cone problem, when outer approximation is used.

Type Integer

Category Control

Topic areas Root, Quadratic

Values

-1	Determined automatically.
0	Use the barrier algorithm.
1	Use the dual simplex on a relaxation of the problem constructed using outer approximation.

Default value -1

Note This control only has an effect if MIQCPALG is set to 1.

Affects routines XPRSmipoptimize (MIPOPTIMIZE), XPRSminim (MINIM), XPRSmaxim (MAXIM).

QSIMPLEXOPS

Description Controls the behavior of the quadratic simplex solvers via a bit-vector (see Section 9.2).

Type Integer

Category Control

Topic areas Quadratic, Simplex, Bit-vector

Values

Bit	Meaning
0	Force traditional primal first phase.
1	Force BigM primal first phase.
2	Force traditional dual first phase.
3	Force BigM dual first phase.
4	Always use artificial bounds in dual.
5	Use original problem basis only when warmstarting the KKT.
6	Skip the primal bound flips for ranged primals (might cause more trouble than good if the bounds are very large).
7	Also do the single pivot crash.
8	Do not apply aggressive perturbation in dual.
9	Applies standard scaling to the KKT system.
10	Do not fall back to using Barrier in case of numerical difficulties with quadratic simplex during a MIP solve.
11	Use primal simplex to solve the phase 1 feasibility problem before applying quadratic primal simplex.
12	Use dual simplex to solve the phase 1 feasibility problem before applying quadratic primal simplex.
13	Use barrier algorithm to solve the phase 1 feasibility problem before applying quadratic primal simplex.
14	Use partial pricing.
15	Use full pricing.
16	Perform cleanup if a superbasic solution is provided for warm-start.

Default value 0

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

QUADRATICUNSHIFT

Description Determines whether an extra solution purification step is called after a solution found by the quadratic simplex (either primal or dual).

Type Integer

Category Control

Topic areas Quadratic, Simplex

Values

-1	Determined automatically.
0	No purification step.
1	Always do the purification step.

Default value `-1`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

RANDOMSEED

Description Sets the initial seed to use for the pseudo-random number generator in the Optimizer. The sequence of random numbers is always reset using the seed when starting a new optimization run.

Type Integer

Category Control

Topic area Determinism

Default value `1`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

REFACTOR

Description Indicates whether the optimization should restart using the current representation of the factorization in memory.

Type Integer

Category Control

Topic areas Branch and Bound Search, Simplex

Values

`-1`	Automatic.
`0`	Do not refactor on reoptimizing.
`1`	Refactor on reoptimizing.

Default value `-1`

Note In the tree search, the optimal bases at the nodes are not refactorized by default, but the optimal basis for an LP problem will be refactorized. If you are repeatedly solving LPs with few changes then it is more efficient to set `REFACTOR` to `0`.

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

REFINEOPS

Description This specifies when the solution refiner should be executed to reduce solution infeasibilities. The refiner will attempt to satisfy the target tolerances for all original linear constraints before presolve or scaling has been applied.

Type Integer

Category Control

Topic areas Solution Refinement, Bit-vector

Values

Bit	Meaning
0	Run the solution refiner on an optimal solution of a continuous problem.
1	Run the solution refiner when a new solution is found during a tree search. The refiner will be applied to the presolved solution before any post-solve operations are applied.
3	Run the solution refiner on each node of the MIP search.
4	Run the solution refiner on an optimal solution before postsolve on a continuous problem.
5	Apply the iterative refiner to refine the solution.
6	Use higher precision in the iterative refinement.
7	If set, the iterative refiner will use the primal simplex algorithm.
8	If set, the iterative refiner will use the dual simplex algorithm.
9	Refine MIP solutions such that rounding them keeps the problem feasible when reoptimized.
10	Attempt to refine MIP solutions such that rounding them keeps the problem feasible when reoptimized, but accept integers solutions even if refinement fails.

Default value 19 (bits 0, 1 and 4 are set)

Notes This is a bit-vector control (see Section 9.2).

If neither the 7th nor 8th bit is set, the refiner will use the primal simplex if the primal violations are larger than the dual violations, otherwise it will use the dual simplex.

If both the 7th and 8th bit are set then the refiner will split the problem into a primal feasible and dual feasible part, and solve the first with primal simplex and the second with dual simplex.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also LPREFINEITERLIMIT, FEASTOLTARGET, OPTIMALITYTOLTARGET, MIPTOLTARGET.

RELAXTREEMEMORYLIMIT

Description When the memory used by the branch and bound search tree exceeds the target specified by the TREEMEMORYLIMIT control, the optimizer will try to reduce this by writing nodes to the tree file. In rare cases, usually where the solve has many millions of very small nodes, the tree structural data (which cannot be written to the tree file) will grow large enough to approach or exceed the tree's memory target. When this happens, optimizer performance can degrade greatly as the solver makes heavy use of the tree file in preference to memory. To prevent this, the solver will automatically relax the tree memory limit when it detects this case; the RELAXTREEMEMORYLIMIT control specifies the proportion of the previous memory limit by which to relax it. Set RELAXTREEMEMORYLIMIT to 0.0 to force the Xpress Optimizer to never relax the tree memory limit in this way.

Type Double

Category Control

Topic areas Limits, Memory

Note While setting higher values of RELAXTREEMEMORYLIMIT can improve performance significantly for a small number of models in low memory situations, the user is advised to use the TREEMEMORYLIMIT control to tune the memory usage of the branch and bound tree, according to the solve characteristics of their problem, rather than increasing RELAXTREEMEMORYLIMIT.

Default value 0.1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also TREEMEMORYLIMIT.

RELPIVOTTOL

Description Simplex: At each iteration a pivot element is chosen within a given column of the matrix. The relative pivot tolerance, RELPIVOTTOL, is the size of the element chosen relative to the largest possible pivot element in the same column.

Type Double

Category Control

Topic areas Simplex, Tolerances

Default value 1.0E-06

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRSpivot.

REPAIRINDEFINITEQ

Description Controls if the optimizer should make indefinite quadratic matrices positive definite when it is possible.

Type Integer

Category Control

Topic area Quadratic

Values

0	Repair if possible.
1	Do not repair.

Default value 1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

REPAIRINFEASMAXTIME

Description *This parameter is deprecated and will be removed in a future release.*
Overall time limit for the repairinfeas tool

Type Integer

Category Control

Topic areas Infeasibility, Limits

Values

0	No time limit.
n>0	If an integer solution has been found, stop MIP search after *n* seconds, otherwise continue until an integer solution is finally found.
n<0	Stop in LP or MIP search after *n* seconds.

Default value 0

Note This control affects the total runtime of repairinfeas, as opposed to MAXTIME which affects the individual solves reapirinfeas carries out.

Affects routines XPRSrepairinfeas (REPAIRINFEAS).

REPAIRINFEASTIMELIMIT

Description Overall time limit for the repairinfeas tool

Type Double

Category Control

Topic areas Infeasibility, Limits

Values `>0` Stop repairinfeas search after the given number of seconds.

Default value `1e+20`

Note This control affects the total runtime of repairinfeas, as opposed to `TIMELIMIT` which affects the individual solves repairinfeas carries out. This control has been newly introduced with Xpress Optimizer version 9.0 and should be used instead of the deprecated `REPAIRINFEASMAXTIME` control.

Affects routines `XPRSrepairinfeas` (`REPAIRINFEAS`).

See also `TIMELIMIT`.

RESOURCESTRATEGY

Description Controls whether the optimizer is allowed to make nondeterministic decisions if memory is running low in an effort to preserve memory and finish the solve. Available memory (or container limits) are automatically detected but can also be changed by `MAXMEMORYSOFT` and `MAXMEMORYHARD`

Type Integer

Category Control

Topic areas Determinism, Memory

Values `1` Allow the optimizer to change the solve path if necessary to preserve memory when getting close to one of the memory limits.

Default value `0`

Affects routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

RLTCUTS

Description Determines whether RLT cuts should be separated in the Xpress Global Solver.

Type Integer

Category Control

Topic areas Quadratic, Cuts

Values

`-1`	The solver decides if RLT cuts are beneficial or not. This is the default setting.
`0`	RLT cuts are disabled.
`1`	RLT cuts are separated.

Default value `-1`

Affects routines `XPRSoptimize` (`OPTIMIZE`).

ROOTPRESOLVE

Description Determines if presolving should be performed on the problem after the tree search has finished with root cutting and heuristics.

Type Integer

Category Control

Topic area Root

Values

`-1`	Let the optimizer decide if the problem should be presolved again.
`0`	Disabled.
`+1`	Always presolve the root problem.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRESOLVE`.

SBBEST

Description Number of infeasible MIP entities to initialize pseudo costs for on each node.

Type Integer

Category Control

Topic area Branching

Values

`-1`	determined automatically.
`0`	disable strong branching.
`n>0`	perform strong branching on up to *n* entities at each node.

Default value `-1`

Note If `SBBEST` is set to zero, the control `HISTORYCOSTS` will also be treated as zero and no past branching or strong branching information will be used in the MIP entity selection.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `SBITERLIMIT`, `SBSELECT`, `SBEFFORT`, `HISTORYCOSTS`.

SBEFFORT

Description Adjusts the overall amount of effort when using strong branching to select an infeasible MIP entity to branch on.

Type Double

Category Control

Topic area Branching

Default value `1.0`

Note `SBEFFORT` is used as a multiplier on other strong branching related controls, and affects the values used for `SBBEST`, `SBSELECT` and `SBITERLIMIT` when those are set to automatic.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also SBBEST, SBITERLIMIT, SBSELECT.

SBESTIMATE

Description Branch and Bound: How to calculate pseudo costs from the local node when selecting an infeasible MIP entity to branch on. These pseudo costs are used in combination with local strong branching and history costs to select the branch candidate.

Type Integer

Category Control

Topic area Branching

Values

-1	Automatically determined.
1-6	Different variants of local pseudo costs.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also SBBEST, SBITERLIMIT, SBSELECT, HISTORYCOSTS.

SBITERLIMIT

Description Number of dual iterations to perform the strong branching for each entity.

Type Integer

Category Control

Topic areas Branching, Limits

Default value -1 — determined automatically.

Note This control can be useful to increase or decrease the amount of effort (and thus time) spent performing strong branching at each node. Setting SBITERLIMIT=0 will disable dual strong branch iterations. Instead, the entity at the head of the candidate list will be selected for branching.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also SBBEST, SBSELECT.

SBSELECT

Description The size of the candidate list of MIP entities for strong branching.

Type Integer

Category Control

Topic area Branching

Values

-2	Automatic (low effort).
-1	Automatic (high effort).
n>=0	Include *n* entities in the candidate list (but always at least SBBEST candidates).

Default value −2

Note Before strong branching is applied on a node of the branch and bound tree, a list of candidates is selected among the infeasible MIP entities. These entities are then evaluated based on the local LP solution and prioritized. Strong branching will then be applied to the SBBEST candidates. The evaluation is potentially expensive and for some problems it might improve performance if the size of the candidate list is reduced.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also SBBEST, SBEFFORT, SBESTIMATE.

SCALING

Description This bit-vector control (see Section 9.2) determines how the Optimizer will rescale a model internally before optimization. If set to 0, no scaling will take place.

Type Integer

Category Control

Topic areas Numerics, Problem Transformation, Bit-vector

Values

Bit	Meaning
0	Row scaling.
1	Column scaling.
2	Row scaling again.
3	Maximum.
4	Curtis-Reid.
5	0: scale by geometric mean. 1: scale by maximum element.
6	Treat big-M rows as normal rows.
7	Scale objective function for the simplex method.
8	Exclude the quadratic part of constraint when calculating scaling factors.
9	Scale before presolve.
10	Do not scale rows up.
11	Do not scale columns down.
12	Do not apply automatic objective scaling.
13	RHS scaling.
14	Disable aggressive quadratic scaling.
15	Enable explicit linear slack scaling.

Default value 163, meaning bits 0, 1, 5 and 7 are set

Note Setting SCALING to 0 will preserve the current scaling of the problem. Note that the Optimizer might automatically select a different scaling strategy, when the control AUTOSCALING is not disabled. However, if SCALING is set to any value by the user, AUTOSCALING will be ignored. The barrier algorithm uses equilibrium scaling.

Affects routines XPRSlpoptimize, XPRSlpoptimize, XPRSmipoptimize, XPRSscale (SCALE).

See also 6.3.1, MAXSCALEFACTOR, OBJSCALEFACTOR, AUTOSCALING.

SERIALIZEPREINTSOL

Description Setting SERIALIZEPREINTSOL to 1 will ensure that the preintsol callback is always fired in a deterministic order during a parallel MIP solve. This applies only when the control DETERMINISTIC is set to 1.

Type Integer

Category Control

Topic areas Callback, Determinism

Values

0	The preintsol callbacks will be fired asynchronously from different threads.
1	The preintsol callbacks will be fired in a deterministic order.

Default value 0

Affects routines XPRSmipoptimize (MIPOPTIMIZE), XPRSaddcbpreintsol.

SIFTING

Description Determines whether to enable sifting algorithm with the dual simplex method.

Type Integer

Category Control

Topic areas LP, Simplex

Values

-1	Automatically determined.
0	Disable sifting.
1	Enable sifting.

Default value -1

Affects routines XPRSmipoptimize (MIPOPTIMIZE), XPRSlpoptimize (LPOPTIMIZE).

SIFTPASSES

Description Determines how quickly we allow to grow the worker problems during the sifting algorithm. Using larger values can increase the number of columns added to the worker problem which often results in increased solve times for the worker problems but the number of necessary sifting iterations may be reduced.

Type Integer

Category Control

Topic areas LP, Simplex

Default value 4

Affects routines XPRSmipoptimize (MIPOPTIMIZE), XPRSlpoptimize (LPOPTIMIZE).

SIFTPRESOLVEOPS

Description Determines the presolve operations for solving the subproblems during the sifting algorithm.

Type Integer

Category Control

Topic areas LP, Simplex

Values

`-1`	Use the PRESOLVEOPS setting specified for the original problem.
`>=0`	Use the value for the PRESOLVEOPS parameter for solving the subproblems during the sifting algorithm.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

See also `PRESOLVEOPS`.

SIFTSWITCH

Description Determines which algorithm to use for solving the subproblems during sifting.

Type Integer

Category Control

Topic areas LP, Simplex

Values

`-1`	Dual simplex.
`0`	Barrier.
`>0`	Use the barrier algorithm while the number of dual infeasibilities is larger than this value, otherwise use dual simplex.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSlpoptimize` (`LPOPTIMIZE`).

SLEEPONTHREADWAIT

Description *This parameter is deprecated and will be removed in a future release.*
In previous versions this was used to determine if the threads should be put into a wait state when waiting for work.

Type Integer

Category Control

Topic area System

Values

`-1`	Automatically determined depending on the CPU the Optimizer is running on.
`0`	Keep the threads busy when waiting for work.
`1`	Put the threads into a wait state when waiting for work.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

SOLTIMELIMIT

Description The maximum time in seconds that the Optimizer will run a MIP solve before it terminates, given that a solution has been found. As long as no solution has been found, this control will have no effect.

Type Double

Category Control

Topic area Limits

Values

`>0`	If an integer solution has been found, stop MIP search after the given number of seconds, otherwise continue until an integer solution is finally found.

Default value `1e+20`

Note This control has been newly introduced with Xpress Optimizer version 9.0 and should be used instead of the deprecated `MAXTIME` control. It can be combined with the `TIMELIMIT` control.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `TIMELIMIT`.

SOSREFTOL

Description The minimum relative gap between the ordering values of elements in a special ordered set. The gap divided by the absolute value of the larger of the two adjacent values must be at least `SOSREFTOL`.

Type Double

Category Control

Topic area Tolerances

Default value `1.0E-06`

Note This tolerance must not be set lower than `1.0E-06`.

Affects routines `XPRSloadmip`, `XPRSloadmiqp`, `XPRSreadprob` (`READPROB`).

SYMMETRY

Description Adjusts the overall amount of effort for symmetry detection.

Type Integer

Category Control

Topic area Presolve

Values

`0`	No symmetry detection.
`1`	Conservative effort.
`2`	Intensive symmetry search.

Default value `1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `SYMSELECT`.

SYMSELECT

Description	Adjusts the overall amount of effort for symmetry detection.
Type	Integer
Category	Control
Topic area	Presolve
Values	`0` Search the whole matrix (otherwise the `0`, `1` and `-1` coefficients only). `1` Search all entities (otherwise binaries only).
Default value	`-1`
Affects routines	`XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`SYMMETRY`.

TIMELIMIT

Description	The maximum time in seconds that the Optimizer will run before it terminates, including the problem setup time and solution time. For MIP problems, this is the total time taken to solve all nodes.
Type	Double
Category	Control
Topic area	Limits
Values	`>0` Stop LP or MIP search after the given number of seconds.
Default value	`1e+20`
Note	This control has been newly introduced with Xpress Optimizer version 9.0 and should be used instead of the deprecated `MAXTIME` control. Note that the meaning of positive values differs between TIMELIMIT and `MAXTIME`. When both controls are set, TIMELIMIT takes precedence. The functionality of positive `MAXTIME` values is covered by `SOLTIMELIMIT`.
Affects routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).
See also	`SOLTIMELIMIT`.

THREADS

Description	The default number of threads used during optimization.
Type	Integer
Category	Control
Topic area	Parallel
Values	`-1` Determined automatically based on hardware configuration. `>0` Number of threads to use.
Default value	`-1`

Note The value may be changed for specific parts of the optimization by the CONCURRENTTHREADS, MIPTHREADS and BARTHREADS controls.

Affects routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also DETERMINISTIC, MIPTHREADS, BARTHREADS, CONCURRENTTHREADS.

TRACE

Description Display the infeasibility diagnosis during presolve. If non-zero, an explanation of the logical deductions made by presolve to deduce infeasibility or unboundedness will be displayed on screen or sent to the message callback function.

Type Integer

Category Control

Topic areas Presolve, Logging

Default value 0

Note Presolve is sometimes able to detect infeasibility and unboundedness in problems. The set of deductions made by presolve can allow the user to diagnose the cause of infeasibility or unboundedness in their problem. However, not all infeasibility or unboundedness can be detected and diagnosed in this way.

Affects routines XPRSlpoptimize (LPOPTIMIZE).

TREECOMPRESSION

Description When writing nodes to the gloal file, the optimizer can try to use data-compression techniques to reduce the size of the tree file on disk. The TREECOMPRESSION control determines the strength of the data-compression algorithm used; higher values give superior data-compression at the affect of decreasing performance, while lower values compress quicker but not as effectively. Where TREECOMPRESSION is set to 0, no data compression will be used on the tree file.

Type Integer

Category Control

Topic areas File IO, Branch and Bound Search, Memory

Default value 2

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also TREEMEMORYLIMIT.

TREECOVERCUTS

Description Branch and Bound: The number of rounds of lifted cover inequalities generated at nodes other than the root node in the tree. Compare with the description for COVERCUTS. A value of -1 indicates the number of rounds is determined automatically.

Type Integer

Category Control

Topic areas Cuts, Branch and Bound Search

Default value -1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

TREECUTSELECT

Description A bit-vector (see Section 9.2) providing detailed control of the cuts created during the tree search of a MIP solve. Use `CUTSELECT` to control cuts on the root node.

Type Integer

Category Control

Topic areas Cuts, Branch and Bound Search, Bit-vector

Values

Bit	Meaning
5	Clique cuts.
6	Mixed Integer Rounding (MIR) cuts.
7	Lifted cover cuts.
8	Turn on row aggregation for MIR cuts.
11	Flow path cuts.
12	Implication cuts.
13	Turn on automatic Lift and Project cutting strategy.
14	Disable cutting from cut rows.
15	Lifted GUB cover cuts.
16	Zero-half cuts.
17	Indicator constraint cuts.
18	Strong Chvatal-Gomory cuts.
20	Farkas cuts.

Default value -1

Note The default value is -1 which enables all bits. Any bits not listed in the above table should be left in their default 'on' state, since the interpretation of such bits might change in future versions of the optimizer.

Note The separation of Multi-Commodity Flow (MCF) cuts is controlled by `MCFCUTSTRATEGY`.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `COVERCUTS`, `GOMCUTS`, `CUTSELECT`.

TREEDIAGNOSTICS

Description A bit-vector (see Section 9.2) providing control over how various tree-management-related messages get printed in the tree log file during the branch-and-bound search.

Type Integer

Category Control

Topic areas Branch and Bound Search, Logging, Bit-vector

Values	Bit	Meaning
	0	Output regular summaries of current tree memory usage.
	1	Output messages whenever tree data is being written to tree file.
	2	Output progress messages while tree data is being written to the tree file, at an interval controlled by the `TREEFILELOGINTERVAL` control.

Default value 7

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPLOG`, `PEAKTOTALTREEMEMORYUSAGE`, `TREEFILELOGINTERVAL`.

TREEGOMCUTS

Description Branch and Bound: The number of rounds of Gomory cuts generated at nodes other than the first node in the tree. Compare with the description for `GOMCUTS`. A value of -1 indicates the number of rounds is determined automatically.

Type Integer

Category Control

Topic areas Cuts, Branch and Bound Search

Default value −1

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

TREEMEMORYLIMIT

Description A soft limit, in megabytes, for the amount of memory to use in storing the branch and bound search tree. This doesn't include memory used for presolve, heuristics, solving the LP relaxation, etc. When set to 0 (the default), the optimizer will calculate a limit automatically based on the amount of free physical memory detected in the machine. When the memory used by the branch and bound tree exceeds this limit, the optimizer will try to reduce the memory usage by writing lower-rated sections of the tree to a file called the "tree file". Though the solve can continue if it cannot bring the tree memory usage below the specified limit, performance will be inhibited and a message will be printed to the log.

Type Integer

Category Control

Topic areas Limits, Memory

Default value 0 (calculate limit automatically)

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `TREEMEMORYSAVINGTARGET`, `TREECOMPRESSION`, `TREEDIAGNOSTICS`.

TREEMEMORYSAVINGTARGET

Description When the memory used by the branch-and-bound search tree exceeds the limit specified by the TREEMEMORYLIMIT control, the optimizer will try to save memory by writing lower-rated sections of the tree to the tree file. The target amount of memory to save will be enough to bring memory usage back below the limit, plus enough extra to give the tree room to grow. The TREEMEMORYSAVINGTARGET control specifies the extra proportion of the tree's size to try to save; for example, if the tree memory limit is 1000Mb and TREEMEMORYSAVINGTARGET is 0.1, when the tree size exceeds 1000Mb the optimizer will try to reduce the tree size to 900Mb. Reducing the value of TREEMEMORYSAVINGTARGET will cause less extra nodes of the tree to be written to the tree file, but will result in the memory saving routine being triggered more often (as the tree will have less room in which to grow), which can reduce performance. Increasing the value of TREEMEMORYSAVINGTARGET will cause additional, more highly-rated nodes, of the tree to be written to the tree file, which can cause performance issues if these nodes are required later in the solve.

Type Double

Category Control

Topic areas Branch and Bound Search, Memory

Default value 0.4

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also TREEMEMORYLIMIT

TREEQCCUTS

Description Branch and Bound: Limit on the number of rounds of outer approximation cuts generated for nodes other than the root node, when solving a mixed integer quadratic constrained or mixed integer second order conic problem with outer approximation.

Type Integer

Category Control

Topic areas Branch and Bound Search, Quadratic

Default value -1 — determined automatically.

Note This control only has an effect for problems with quadratic or second order cone constraints, and only if outer approximation has not been disabled by setting MIQCPALG to 0.

Affects routines XPRSmipoptimize (MIPOPTIMIZE).

See also QCCUTS.

TUNERHISTORY

Description Tuner: Whether to reuse and append to previous tuner results of the same problem.

Type Integer

Category Control

Topic area Tuner

Values	0	Discard any previous tuner results.
	1	Append new results to the previous tuner results, but do not reuse them.
	2	Reuse the previous results and append new results to it.

Default value 2

Notes Please refer to Section 5.13.5 for more information about reusing tuner results.

This control only has an effect on the tuner. This control cannot be tuned.

Affects routines `XPRStune` (`TUNE`).

TUNERMAXTIME

Description Tuner: The maximum time in seconds that the tuner will run before it terminates.

Type Double

Category Control

Topic areas Tuner, Limits

Values	0	No time limit.
	>0	Stop the tuner after the given number of seconds.

Default value 0

Note This control changed type with Xpress Optimizer version 9.0 (from integer to double). It is strongly recommended to adapt all access calls to this control accordingly. This control only has an effect on the tuner.

Affects routines `XPRStune` (`TUNE`).

See also `TIMELIMIT`.

TUNERMETHOD

Description Tuner: Selects a factory tuner method. A tuner method consists of a list of controls with different settings that the tuner will evaluate and try to combine.

Type Integer

Category Control

Topic area Tuner

Values	-1	Automatically determined. The tuner will select the default method based on the problem type.
	0	Select the default LP tuner method.
	1	Select the default MIP tuner method.
	2	Select a more comprehensive MIP tuner method.
	3	Select a root-focus MIP tuner method.
	4	Select a tree-focus MIP tuner method.
	5	Select a simple MIP tuner method.
	6	Select the default SLP tuner method.
	7	Select the default MISLP tuner method.
	8	Select a MIP tuner method focussed on primal heuristics.
	9	Select the default Xpress Global tuner method.

Default value -1

Notes If the tuner has already loaded a user-defined tuner method, then it will not load any factory tuner method.

Please refer to Section 5.13.2 for more information about the tuner method, and Appendix A.8 for the format of the tuner method file.

This control only has an effect on the tuner. This control cannot be tuned.

Affects routines XPRStune (TUNE).

TUNERMETHODFILE

Description Tuner: Defines a file from which the tuner can read user-defined tuner method.

Type String

Category Control

Topic areas Tuner, File IO

Default value (empty)

Notes If the tuner has already loaded a tuner method via XPRStunerreadmethod, then it will not check this control. Otherwise, when this control is defined and a tuner method can be succesfully loaded from this file, then the tuner will not load any factory tuner method.

Please refer to Section 5.13.2 for more information about the tuner method, and Appendix A.8 for the format of the tuner method file.

This control only has an effect on the tuner. This control cannot be tuned.

Affects routines XPRStune (TUNE).

TUNERMODE

Description Tuner: Whether to always enable the tuner or disable it.

Type Integer

Category Control

Topic area Tuner

Values

-1	No effect.
0	Always disable the tuner. XPRStune (TUNE) will have no effect.
1	Always enable the tuner. XPRSmipoptimize (MIPOPTIMIZE), XPRSlpoptimize (LPOPTIMIZE), etc. will call the tuner before solving the problem.

Default value -1

Note This control cannot be tuned.

Affects routines XPRStune (TUNE), XPRSmipoptimize (MIPOPTIMIZE), XPRSlpoptimize (LPOPTIMIZE).

TUNEROUTPUT

Description	Tuner: Whether to output tuner results and logs to the file system.
Type	Integer
Category	Control
Topic areas	Tuner, File IO
Values	`0` Don't output to the file system. `1` Output results and logs to the file system.
Default value	`1`
Notes	Please refer to Section 5.13.3 for more information about the tuner output. This control only has an effect on the tuner. This control cannot be tuned.
Affects routines	`XPRStune` (`TUNE`).

TUNEROUTPUTPATH

Description	Tuner: Defines a root path to which the tuner writes the result file and logs.
Type	String
Category	Control
Topic areas	Tuner, File IO
Default value	tuneroutput
Notes	This control only defines the root path for the tuner output. For each problem, the tuner result will be output to a subfolder underneath this path. For example, by default, the tuner result for a problem called `prob` will be located at `tuneroutput/prob/` Please refer to Section 5.13.3 for more information about the tuner output. This control only has an effect on the tuner. This control cannot be tuned.
Affects routines	`XPRStune` (`TUNE`).

TUNERPERMUTE

Description	Tuner: Defines the number of permutations to solve for each control setting.
Type	Integer
Category	Control
Topic areas	Tuner, Problem Transformation
Values	`0` Solve the original problem only for each setting. `n>0` Solve the original problem and `n` permuted problems for each setting.
Default value	`0`
Notes	Please refer to Section 5.13.7 for more information about tuner problem permutations. This control only has an effect on the tuner. This control cannot be tuned.
Affects routines	`XPRStune` (`TUNE`).

TUNERSESSIONNAME

Description Tuner: Defines a session name for the tuner.

Type String

Category Control

Topic area Tuner

Default value (empty)

Notes When defined, the session name will override the problem name within the tuner. For example, if this control is set to `session`, then the tuner result for a problem will be located at `tuneroutput/session/`

This control can be useful when the problem name is randomly generated.

Please refer to Section 5.13.3 for more information about the tuner output.

This control only has an effect on the tuner. This control cannot be tuned.

Affects routines `XPRStune` (`TUNE`).

TUNERTARGET

Description Tuner: Defines the tuner target – what should be evaluated when comparing two runs with different control settings.

Type Integer

Category Control

Topic area Tuner

Values

`-1`	Automatically determined. The tuner will choose the default target based on problem type.
`0`	Solution time then gap. (MIP/MISLP default)
`1`	Solution time then best bound.
`2`	Solution time then best integer solution.
`3`	The primal dual integral.
`4`	Time only. (LP/SLP default)
`5`	SLP objective only. (SLP/MISLP choice)
`6`	SLP validation number only. (SLP/MISLP choice)
`7`	Gap only.
`8`	Best bound only.
`9`	Best integer solution only.
`10`	Best primal integral. (Only for individual instances, not for problem sets)

Default value `-1`

Notes Please refer to Section 5.13.4 for more information about tuner targets.

This control only has an effect on the tuner. This control cannot be tuned.

Affects routines `XPRStune` (`TUNE`).

TUNERTHREADS

Description Tuner: the number of threads used by the tuner.

Type Integer

Category Control

Topic areas Tuner, Parallel

Values

-1	Choose automaticlly.
1	The tuner will run in sequential.
n>1	The tuner will run in parallel with n threads.

Default value 1

Notes Setting this control will not affect number of threads used by each individual run. It is recommended to have the product of TUNERTHREADS and THREADS less or equal to the number of system threads.

When setting TUNERTHREADS=-1, the tuner will automatically use as many threads as the number of logical processors detected.

Please refer to Section 5.13.6 for more information about tuner with multiple threads.

This control only has an effect on the tuner. This control cannot be tuned.

Affects routines XPRStune (TUNE).

TUNERVERBOSE

Description Tuner: whether the tuner should prints detailed information for each run.

Type Integer

Category Control

Topic areas Tuner, Logging

Values

1	Print extra information.
0	Print less information.

Default value 1

Notes Please refer to Section 5.13.6 for more information about tuner with multiple threads.

This control only has an effect on the tuner. This control cannot be tuned.

Affects routines XPRStune (TUNE).

USERSOLHEURISTIC

Description Determines how much effort to put into running a local search heuristic to find a feasible integer solution from a partial or infeasible user solution.

Type Integer

Category Control

Topic area Heuristics

Values	`-1`	Automatically determined.
	`0`	Search heuristic disabled.
	`1`	Light effort.
	`2`	Moderate effort.
	`3`	High effort.

Default value `-1`

Note When a partial or infeasible user solution is added with `XPRSaddmipsol`, a local search heuristic will be applied to the problem in an attempt to find a feasible, integer solution that either completes the partial solution or is close to the infeasible solution. Whether to run such a heuristic, or how much effort to put into the heuristic can be controlled by this `USERSOLHEURISTIC` parameter.

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `HEURSEARCHROOTSELECT`, `HEURSEARCHTREESELECT`.

VARSELECTION

Description Branch and Bound: This determines the formula used to calculate the estimate of each integer variable, and thus which integer variable is selected to be branched on at a given node. The variable selected to be branched on is the one with the maximum estimate.

Type Integer

Category Control

Topic area Branching

Values	`-1`	Determined automatically.
	`1`	The minimum of the 'up' and 'down' pseudo costs.
	`2`	The 'up' pseudo cost plus the 'down' pseudo cost.
	`3`	The maximum of the 'up' and 'down' pseudo costs, plus twice the minimum of the 'up' and 'down' pseudo costs.
	`4`	The maximum of the 'up' and 'down' pseudo costs.
	`5`	The 'down' pseudo cost.
	`6`	The 'up' pseudo cost.
	`7`	A weighted combination of the 'up' and 'down' pseudo costs, where the weights depend on how fractional the variable is.
	`8`	The product of the 'up' and 'down' pseudo costs.

Default value `-1`

Affects routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

VERSION

Description The Optimizer version number, e.g. `1301` meaning release 13.01.

Type Integer

Category Control

Topic area Misc

Default value Software version dependent

WORKLIMIT

Description	The maximum work (measured in work units) that the Optimizer will run before it terminates. WORK is accumulated during the search and ever increasing. In contrast to TIME, WORK is independent of the hardware and platform on which the search is conducted. The WORKLIMIT serves as a deterministic stopping criterion. When it is reached, it leaves the optimizer in a reproducible state.
Type	Double
Category	Control
Topic areas	Limits, Determinism
Values	>0 Stop LP or MIP search when the given number of work units is reached.
Default value	1e+20
Note	The reproducibility of a search conducted under a WORKLIMIT depends on the chosen determinism of the search, see the control DETERMINISTIC for more information.
Note	The speed at which work is accumulated depends on the hardware, the algorithms involved, and the problem being solved.
Note	
Affects routines	XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).
See also	TIMELIMIT, DETERMINISTIC.

Chapter 10

Problem Attributes

During the optimization process, various properties of the problem being solved are stored and made available to users of the FICO Xpress Libraries in the form of *problem attributes*. These can be accessed in much the same manner as for the controls. Examples of problem attributes include the sizes of arrays, for which library users may need to allocate space before the arrays themselves are retrieved. A full list of the attributes available and their types may be found in this chapter.

10.1 Retrieving Problem Attributes

Library users are provided with the following three functions for obtaining the values of attributes:

XPRSgetintattrib	XPRSgetdblattrib	XPRSgetstrattrib

Much as for the controls previously, it should be noted that the attributes as listed in this chapter *must* be prefixed with XPRS_ to be used with the FICO Xpress Libraries and failure to do so will result in an error. An example of their usage is the following which returns and prints the optimal value of the objective function after the linear problem has been solved:

```
XPRSgetdblattrib(prob, XPRS_LPOBJVAL, &lpobjval);

printf("The objective value is %2.1f\n", lpobjval);
```

ACTIVENODES

Description Number of outstanding nodes.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

ALGORITHM

Description The algorithm the optimizer currently is running / was running just before completition.

Type Integer

Category Attribute

Topic area LP

Values

1	No LP optimization yet.
2	Dual simplex.
3	Primal simplex.
4	Newton barrier or hybrid gradient.
5	Network simplex.

Note If the barrier or hybrid gradient algorithms are used with crossover, then the value of `ALGORITHM` during the crossover and the final clean up will reflect the algorithm used, but will be reset to barrier or hybrid gradient once the optimization is complete.

ATTENTIONLEVEL

Description A measure between 0 and 1 for how numerically unstable the problem is. The attention level is based on a weighted combination of the number of basis condition numbers exceeding certain thresholds. It considers all nodes sampled by `MIPKAPPAFREQ`, with a setting of 1 being the most frequent sampling rate. The higher the attention level, the worse conditioned is the problem.

Type Double

Category Attribute

Topic areas Branch and Bound Search, Numerics

Set by routines `XPRSmipoptimize`.

See also `MAXKAPPA`, `MIPKAPPAFREQ`, `PREDICTEDATTLEVEL`.

AVAILABLEMEMORY

Description The amount of heap memory detected by Xpress as free.

Type Integer

Category Attribute

Topic area Memory

Note On 64-bit systems this is a 64-bit integer, use `XPRSgetintattrib64` to retrieve its value.

See also `PEAKMEMORY`, `CURRENTMEMORY`, `TOTALMEMORY`.

BARAASIZE

Description Number of nonzeros in AA^T.

Type Integer

Category Attribute

Topic area Barrier

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARCGAP

Description Convergence criterion for the Newton barrier algorithm.

Type Double

Category Attribute

Topic area Barrier

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARCONDA

Description Absolute condition measure calculated in the last iteration of the barrier algorithm.

Type Double

Category Attribute

Topic area Barrier

Set by routines The barrier algorithm.

BARCONDD

Description Condition measure calculated in the last iteration of the barrier algorithm.

Type Double

Category Attribute

Topic area Barrier

Set by routines The barrier algorithm.

BARCROSSOVER

Description Indicates whether or not the basis crossover phase has been entered.

Type Integer

Category Attribute

Topic area Barrier

Values

`0`	the crossover phase has not been entered.
`1`	the crossover phase has been entered.

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARDENSECOL

Description	Number of dense columns found in the matrix.
Type	Integer
Category	Attribute
Topic area	Barrier
Set by routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARDUALINF

Description	Sum of the dual infeasibilities for the Newton barrier algorithm.
Type	Double
Category	Attribute
Topic area	Barrier
Set by routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARDUALOBJ

Description	Dual objective value calculated by the Newton barrier algorithm.
Type	Double
Category	Attribute
Topic area	Barrier
Set by routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARITER

Description	Number of Newton barrier iterations.
Type	Integer
Category	Attribute
Topic area	Barrier
Set by routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARLSIZE

Description	Number of nonzeros in L resulting from the Cholesky factorization.
Type	Integer
Category	Attribute
Topic area	Barrier
Set by routines	`XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARPRIMALINF

Description Sum of the primal infeasibilities for the Newton barrier algorithm.

Type Double

Category Attribute

Topic area Barrier

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARPRIMALOBJ

Description Primal objective value calculated by the Newton barrier algorithm.

Type Double

Category Attribute

Topic area Barrier

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

BARSING

Description Number of linearly dependent binding constraints at the optimal barrier solution. These results in singularities in the Cholesky decomposition during the barrier that may cause numerical troubles. Larger dependence means more chance for numerical difficulties.

Type Double

Category Attribute

Topic area Barrier

Set by routines The barrier algorithm.

BARSINGR

Description Regularized number of linearly dependent binding constraints at the optimal barrier solution. These results in singularities in the Cholesky decomposition during the barrier that may cause numerical troubles. Larger dependence means more chance for numerical difficulties.

Type Double

Category Attribute

Topic area Barrier

Set by routines The barrier algorithm.

BESTBOUND

Description Value of the best bound determined so far by the MIP search.

Type Double

Category Attribute

Topic area Branch and Bound Search

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

BOUNDNAME

Description Active bound name.

Type String

Category Attribute

Topic area Misc

Set by routines `XPRSreadprob`.

BRANCHVALUE

Description The value of the branching variable at a node of the Branch and Bound tree.

Type Double

Category Attribute

Topic area Branching

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

BRANCHVAR

Description The branching variable at a node of the Branch and Bound tree.

Type Integer

Category Attribute

Topic area Branching

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

CALLBACKCOUNT_CUTMGR

Description This attribute counts the number of times the cut manager callback set by `XPRSaddcbcutmgr` has been called for the current node, including the current callback call. The value of this attribute should only be used from within the cut manager callback.

Type Integer

Category Attribute

Topic areas Callback, Cuts

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

CALLBACKCOUNT_OPTNODE

Description This attribute counts the number of times the optimal node callback set by XPRSaddcboptnode has been called for the current node, including the current callback call. The value of this attribute should only be used from within the optimal node callback.

Type Integer

Category Attribute

Topic areas Callback, Branch and Bound Search

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

CHECKSONMAXCUTTIME

Description This attribute is used to set the value of the MAXCHECKSONMAXCUTTIME control. Its value is the number of times the optimizer checked the MAXCUTTIME criterion during the last call to the optimization routine XPRSmipoptimize. If a run terminates cutting operations on the MAXCUTTIME criterion then the attribute is the negative of the number of times the optimizer checked the MAXCUTTIME criterion up to and including the check when the termination was activated. Note that the attribute is set to zero at the beginning of each call to an optimization routine.

Type Integer

Category Attribute

Topic area Limits

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

CHECKSONMAXTIME

Description This attribute is used to set the value of the MAXCHECKSONMAXTIME control. Its value is the number of times the optimizer checked the MAXTIME criterion during the last call to the optimization routine XPRSmipoptimize. If a run terminates on the MAXTIME criterion then the attribute is the negative of the number of times the optimizer checked the MAXTIME criterion up to and including the check when the termination was activated. Note that the attribute is set to zero at the beginning of each call to an optimization routine.

Type Integer

Category Attribute

Topic area Limits

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

COLS

Description Number of columns (i.e. variables) in the matrix.

Type Integer

Category Attribute

Topic area Problem Information

Note If the matrix is in a presolved state, this attribute returns the number of columns in the **presolved** matrix. If you require the value for the original matrix then use the INPUTCOLS attribute instead. The PRESOLVESTATE attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines XPRSaddcols, XPRSdelcols, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE) XPRSreadprob.

See also INPUTCOLS, ROWS.

COMPUTEEXECUTIONS

Description The number of solves executed on a compute server.

Type Integer

Category Attribute

Topic area Compute Interface

Set by routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRStune, XPRSrepairinfeas, XPRSiisfirst, XPRSiisnext, XPRSiisall.

CONEELEMS

Description Number of second order cone coefficients in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Note If the matrix is in a presolved state, this attribute returns the number of the second order (including rotated second order) cone coefficients in the **presolved** matrix. Second order conic quadratic constraints are automaticly detected at optimization time, and this attribute is not set before optimizing the problem.

Set by routines Optimizing the problem.

CONES

Description Number of second order and rotated second order cones in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Note If the matrix is in a presolved state, this attribute returns the number of second order (including rotated second order) cones in the **presolved** matrix. Conic quadratic constraints are automaticly detected at optimization time, and this attribute is not set before optimizing the problem.

Set by routines Optimizing the problem.

CORESDETECTED

Description Number of logical cores detected by the optimizer, which is the total number of threads the hardware can execute across all CPUs.

Type Integer

Category Attribute

Topic area System

Values `>=1` Detected number of logical cores.

Notes If Xpress is running in a container and the container limits the number of cores then this limit is applied to the attribute value by default. This clipping can be disabled by setting `IGNORECONTAINERCPULIMIT` to `1`.

The optimizer will automatically use as many solver threads as the number of logical cores detected.

If the detection fails, the optimizer will default to using a single thread only.

Set by routines `XPRSinit`.

See also `THREADS`, `CORESPERCPUDETECTED`, `CPUSDETECTED`, `PHYSICALCORESDETECTED`, `PHYSICALCORESPERCPUDETECTED`.

CORESPERCPUDETECTED

Description Number of logical cores per CPU unit detected by the optimizer, which is the number of threads each CPU can execute.

Type Integer

Category Attribute

Topic area System

Values `>=1` Detected number of logical cores per CPU unit.

Set by routines `XPRSinit`.

See also `THREADS`, `CORESDETECTED`, `CPUSDETECTED`, `PHYSICALCORESDETECTED`, `PHYSICALCORESPERCPUDETECTED`.

CPISCALEFACTOR

Description scale factor from primal integral computation.

Type Double

Category Attribute

Topic area Misc

Note This attribute represents the scaling factor that was used when computing the primal integral. It can be used to compute an updated (or correlated) primal integral with respect to a new reference solution. For details see Berthold and Csizmadia: *The confined primal integral*, Mathematical Programming volume 188(2), pp. 523-537, 2021.

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `OBSERVEDPRIMALINTEGRAL`, `CPIALPHA`.

CPUSDETECTED

Description Number of CPU units detected by the optimizer.

Type Integer

Category Attribute

Topic area System

Values `>=1` Detected number of CPU units.

Set by routines `XPRSinit`.

See also `THREADS`, `CORESDETECTED`, `CORESPERCPUDETECTED`, `PHYSICALCORESDETECTED`, `PHYSICALCORESPERCPUDETECTED`.

CROSSOVERITER

Description Number of simplex iterations performed in crossover.

Type Integer

Category Attribute

Topic areas LP, Barrier

Note In case of a barrier solve with crossover, this will only include simplex iterations performed during crossover. For the number of simplex iterations performed during cleanup, refer to `SIMPLEXITER`.

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `SIMPLEXITER`.

CURRENTMEMORY

Description The amount of dynamically allocated heap memory by the problem being solved.

Type Integer

Category Attribute

Topic area Memory

Note On 64-bit systems this is a 64-bit integer, use `XPRSgetintattrib64` to retrieve its value.

See also `PEAKMEMORY`, `CURRENTMEMORY`, `TOTALMEMORY`.

CURRENTNODE

Description The unique identifier of the current node in the tree search.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Note The root node is always identified as node 1.

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PARENTNODE`.

CURRMIPCUTOFF

Description The current MIP cut off.

Type Double

Category Attribute

Topic area Branch and Bound Search

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPABSCUTOFF`.

CUTROUNDS

Description Number of rounds of cuts applied to the current node of a branch-and-bound search.

Type Integer

Category Attribute

Topic area Cuts

Note Can be used within the `cutround` callback set by `XPRSaddcbcutround` to retrieve the number of rounds of cutting that has been applied since the start of the current node solve. Since the `cutround` callback is fired *before* a round of cuts, the `CUTROUNDS` attribute will start from zero.

Set by routines `XPRSmipoptimize`, `XPRSoptimize`.

CUTS

Description Number of cuts being added to the matrix.

Type Integer

Category Attribute

Topic area Cuts

Set by routines `XPRSaddcuts`, `XPRSdelcpcuts`, `XPRSdelcuts`, `XPRSloadcuts`, `XPRSloadmodelcuts`.

DUALINFEAS

Description Number of dual infeasibilities.

Type Integer

Category Attribute

Topic area LP

Note If the matrix is in a presolved state, this attribute returns the number of dual infeasibilities in the **presolved** matrix. If you require the value for the original matrix, make sure you obtain the value when the matrix is not presolved. The `PRESOLVESTATE` attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRIMALINFEAS`.

ELEMS

Description Number of matrix nonzeros (elements).

Type Integer

Category Attribute

Topic area Problem Information

Note If the matrix is in a presolved state, this attribute returns the number of matrix nonzeros in the **presolved** matrix. If you require the value for the original matrix, make sure you obtain the value when the matrix is not presolved. The `PRESOLVESTATE` attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines `XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`), `XPRSreadprob`.

ERRORCODE

Description The most recent Optimizer error number that occurred. This is useful to determine the precise error that has occurred, after an Optimizer function has signalled an error by returning a non-zero value. The return value itself is **not** the error number. Refer to the section 11.2 for a list of possible error numbers, the errors they indicate, and advice on what they mean and how to resolve them. A short error message may be obtained using `XPRSgetlasterror`, and all messages may be intercepted using the user output callback function; see `XPRSaddcbmessage`.

Type Integer

Category Attribute

Topic area Misc

Set by routines Any.

GENCONCOLS

Description Number of input variables in general constraints (i.e. MIN/MAX/AND/OR/ABS constraints) in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note General constraints are transformed during presolve, so `GENCONCOLS` will always be zero in a presolved problem. If you require the value for the original problem then use the `ORIGINALGENCONCOLS` attribute instead.

Set by routines `XPRSaddgencons`, `XPRSdelgencons`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `ORIGINALGENCONCOLS`, `GENCONS`, `GENCONVALS`.

GENCONS

Description The number of general constraints (i.e. MIN/MAX/AND/OR/ABS constraints) in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note General constraints are transformed during presolve, so `GENCONS` will always be zero in a presolved problem. If you require the value for the original problem then use the `ORIGINALGENCONS` attribute instead.

Set by routines `XPRSaddgencons`, `XPRSdelgencons`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `ORIGINALGENCONSGENCONCOLS`, `GENCONVALS`.

GENCONVALS

Description Number of constant values in general constraints (MIN/MAX constraints) in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note General constraints are transformed during presolve, so `GENCONVALS` will always be zero in a presolved problem. If you require the value for the original problem then use the `ORIGINALGENCONVALS` attribute instead.

Set by routines XPRSaddgencons, XPRSdelgencons, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also ORIGINALGENCONVALS, GENCONS, GENCONCOLS.

GLOBALNLPINFEAS

Description Number of nonlinear infeasibilities at the current node of a global solve, measured as the number of violated atomic formulas.

Type Integer

Category Attribute

Topic area Global

Set by routines XPRSnlpoptimize (NLPOPTIMIZE).

See also MIPINFEAS.

GLOBALBOUNDINGBOXAPPLIED

Description Whether a bounding box equal to the absolute value of the GLOBALBOUNDINGBOX control was applied to the problem after the initial solve came back infeasible and if so, to which variables.

Type Double

Category Attribute

Topic area Global

Values

0	No bounding box applied(XPRS_GLOBALBOUNDINGBOX_NOT_APPLIED).
1	Bounding box applied to all original variables(XPRS_GLOBALBOUNDINGBOX_ORIGINAL).
2	Bounding box applied to all original and auxiliary variables(XPRS_GLOBALBOUNDINGBOX_AUXILIARY).

Set by routines XPRSoptimize (OPTIMIZE).

IISSOLSTATUS

Description IIS solution status.

Type Integer

Category Attribute

Topic area Infeasibility

Values

0	Unstarted (XPRS_IIS_UNSTARTED), typically due to a license or input error.
1	Feasible (XPRS_IIS_FEASIBLE), so no IIS can be produced.
2	Completed (XPRS_IIS_COMPLETED); this is the normal termination. The number of IIS sets can be queried through the NUMIIS attribute.
3	Unfinished (XPRS_IIS_UNFINISHED), either because of hitting the time limit or because of a user interrupt. When the IIS procedure terminates with this status the last IIS set may not be irreducible.

Note The possible return values are defined as constants in the Optimizer C header file.

Set by routines `IIS`, `XPRSiisfirst`, `XPRSiisnext`, `XPRSiisall`.

See also `NUMIIS`.

INDICATORS

Description Number of indicator constrains in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note When the matrix is in a presolved state, the indicator constraints are stored in a special pool and not part of the matrix. Otherwise the indicator constraints are rows of the matrix and their details can be retrieved with the `XPRSgetindicators` function. To find the number of indicators in the original matrix, use the `ORIGINALINDICATORS` attribute. The `PRESOLVESTATE` attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines `XPRSsetindicators`, `XPRSdelindicators`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `ORIGINALINDICATORS`.

INPUTCOLS

Description Number of columns (i.e. variables) in the original matrix before nonlinear reformulations.

Type Integer

Category Attribute

Topic area Problem Information

Note During nonlinear solves, the `ORIGINALCOLS` attribute may include columns introduced by nonlinear reformulations, while `INPUTCOLS` will always refer to the number of columns in the problem that was originally input. Column indices used in calls to `XPRSgetsolution` and `XPRSgetredcosts` must be in the range `0` to `INPUTCOLS-1`.

Set by routines `XPRSaddcols`, `XPRSdelcols`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `COLS`, `ORIGINALCOLS`, `INPUTROWS`.

INPUTROWS

Description Number of rows (i.e. constraints) in the original matrix before nonlinear reformulations.

Type Integer

Category Attribute

Topic area Problem Information

Note During nonlinear solves, the ORIGINALROWS attribute may include rows introduced by nonlinear reformulations, while INPUTROWS will always refer to the number of rows in the problem that was originally input. Row indices used in calls to XPRSgetslacks and XPRSgetduals must be in the range 0 to INPUTROWS-1.

Set by routines XPRSaddrows, XPRSdelrows, XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp, XPRSreadprob.

See also ROWS, ORIGINALROWS, INPUTCOLS.

LPOBJVAL

Description Value of the objective function of the last LP solved.

Type Double

Category Attribute

Topic area LP

Set by routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

See also OBJVAL, MIPOBJVAL, OBJRHS.

LPSTATUS

Description LP solution status.

Type Integer

Category Attribute

Topic area LP

Values

0	Unstarted (XPRS_LP_UNSTARTED).
1	Optimal (XPRS_LP_OPTIMAL).
2	Infeasible (XPRS_LP_INFEAS).
3	Objective worse than cutoff (XPRS_LP_CUTOFF).
4	Unfinished (XPRS_LP_UNFINISHED).
5	Unbounded (XPRS_LP_UNBOUNDED).
6	Cutoff in dual (XPRS_LP_CUTOFF_IN_DUAL).
7	Problem could not be solved due to numerical issues. (XPRS_LP_UNSOLVED).
8	Problem contains quadratic data, which is not convex (XPRS_LP_NONCONVEX). Consider using FICO Xpress Global.

Note The possible return values are defined as constants in the Optimizer C header file.

Set by routines XPRSlpoptimize (LPOPTIMIZE).

See also MIPSTATUS.

MATRIXNAME

Description The matrix name.

Type String

Category Attribute

Topic area Problem Information

Note This is the name read from the `MATRIX` field in an MPS matrix, and is *not* related to the problem name used in the Optimizer. Use `XPRSgetprobname` to get the problem name.

Set by routines `XPRSreadprob`, `XPRSsetprobname`.

MAXABSDUALINFEAS

Description Maximum calculated absolute dual infeasibility in the unscaled original problem.

Type Double

Category Attribute

Topic areas LP, Numerics

Set by routines `XPRSlpoptimize`, `XPRSmipoptimize`.

MAXABSPRIMALINFEAS

Description Maximum calculated absolute primal infeasibility in the unscaled original problem.

Type Double

Category Attribute

Topic areas LP, Numerics

Set by routines `XPRSlpoptimize`, `XPRSmipoptimize`, `XPRSoptimize`, `XPRSrefinemipsol`.

MAXKAPPA

Description Largest basis condition number (also known as kappa) calculated through all nodes sampled by `MIPKAPPAFREQ`.

Type Double

Category Attribute

Topic areas Branch and Bound Search, Numerics

Set by routines `XPRSmipoptimize`.

See also `MIPKAPPAFREQ`.

MAXMIPINFEAS

Description Maximum integer fractionality in the solution.

Type Double

Category Attribute

Topic areas Numerics, MIP Entities

Set by routines `XPRSmipoptimize`, `XPRSoptimize`.

MAXPROBNAMELENGTH

Description Maximum size of the problem name and also the maximum allowed length of the file or path string for any function that accepts such an argument (not including the NULL terminator).

Type Integer

Category Attribute

Topic area Misc

Set by routines `XPRSgetprobname`, `XPRSsetprobname`.

MAXRELDUALINFEAS

Description Maximum calculated relative dual infeasibility in the unscaled original problem.

Type Double

Category Attribute

Topic areas LP, Numerics

Set by routines `XPRSlpoptimize`, `XPRSmipoptimize`.

MAXRELPRIMALINFEAS

Description Maximum calculated relative primal infeasibility in the unscaled original problem.

Type Double

Category Attribute

Topic areas LP, Numerics

Set by routines `XPRSlpoptimize`, `XPRSmipoptimize`, `XPRSoptimize`.

MEMORYLIMITDETECTED

Description The detected amount of memory accessible to the solver process, in megabytes. This is the minimum of physical memory, virtual memory limitations, and detected container limitations (Linux only).

Type Integer

Category Attribute

Topic area Memory

Set by routines `XPRSinit`.

See also `CORESDETECTED`, `CORESPERCPUDETECTED`, `CPUSDETECTED`, `PHYSICALCORESDETECTED`, `PHYSICALCORESPERCPUDETECTED`, `PEAKMEMORY`, `CURRENTMEMORY`, `TOTALMEMORY`.

MIPBESTOBJVAL

Description Objective function value of the best integer solution found.

Type Double

Category Attribute

Topic area Branch and Bound Search

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPOBJVAL`, `OBJVAL`.

MIPENTS

Description Number of MIP entities (i.e. binary, integer, semi-continuous, partial integer, and semi-continuous integer variables) but excluding the number of special ordered sets.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note If the matrix is in a presolved state, this attribute returns the number of MIP entities in the **presolved** matrix. If you require the value for the original matrix then use the `ORIGINALMIPENTS` attribute instead. The `PRESOLVESTATE` attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines `XPRSchgcoltype`, `XPRSdelcols`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `SETS`, `ORIGINALMIPENTS`.

MIPINFEAS

Description Number of integer infeasibilities, including violations of special ordered sets, at the current node.

Type Integer

Category Attribute

Topic areas Branch and Bound Search, MIP Entities

Note In callbacks set by `XPRSaddcbpreintsol` or `XPRSaddcbintsol` this attribute is set to -1 if the solution was found by a heuristic.

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRIMALINFEAS`.

MIPOBJVAL

Description Objective function value of the last integer solution found.

Type Double

Category Attribute

Topic areas Branch and Bound Search, Solution

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `OBJVAL`, `MIPBESTOBJVAL`.

MIPSOLNODE

Description Node at which the last integer feasible solution was found.

Type Integer

Category Attribute

Topic areas Branch and Bound Search, Solution

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `MIPSOLTIME`.

MIPSOLS

Description Number of integer solutions that have been found.

Type Integer

Category Attribute

Topic areas Branch and Bound Search, Solution

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

MIPSOLTIME

Description Time at which the last integer feasible solution was found.

Type Double

Category Attribute

Topic areas Branch and Bound Search, Solution

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

Note This attribute changed type with Xpress Optimizer version 9.0 (from integer to double). It is strongly recommended to adapt all query calls to this attribute accordingly.

See also `MIPSOLNODE`, `MAXSTALLTIME`.

MIPSTATUS

Description (MIP) solution status.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Values

0	Problem has not been loaded (`XPRS_MIP_NOT_LOADED`).
1	MIP search incomplete - the initial continuous relaxation has not been solved and no integer solution has been found (`XPRS_MIP_LP_NOT_OPTIMAL`).
2	MIP search incomplete - the initial continuous relaxation has been solved and no integer solution has been found (`XPRS_MIP_LP_OPTIMAL`).
3	MIP search incomplete - no integer solution found (`XPRS_MIP_NO_SOL_FOUND`).
4	MIP search incomplete - an integer solution has been found (`XPRS_MIP_SOLUTION`).
5	MIP search complete - no integer solution found (`XPRS_MIP_INFEAS`).
6	MIP search complete - integer solution found (`XPRS_MIP_OPTIMAL`).
7	MIP search incomplete - the initial continuous relaxation was found to be unbounded. A solution may have been found (`XPRS_MIP_UNBOUNDED`).

Note The possible return values are defined as constants in the Optimizer C header file.

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `LPSTATUS`.

MIPTHREADID

Description The ID for the MIP thread.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Note The first MIP thread has ID 0 and is the same as the main thread. All other threads are new threads and are destroyed when the MIP search is halted.

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

See also MIPTHREADS.

NAMELENGTH

Description The length (in 8 character units) of row and column names in the matrix. To allocate a character array to store names, you must allow 8*NAMELENGTH+1 characters per name (the +1 allows for the string terminator character).

Type Integer

Category Attribute

Topic area Misc

Set by routines XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp, XPRSreadprob.

NODEDEPTH

Description Depth of the current node.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

OBJECTIVES

Description Number of objectives in the problem.

Type Integer

Category Attribute

Topic area Multiobjective

Set by routines XPRSreadprob, XPRSaddobj, XPRSdelobj, XPRSchgobjn.

See also SOLVEDOBJS, OBJSTOSOLVE.

NODES

Description Number of nodes solved so far in the MIP search. A node is counted as solved when it is either dropped or branched on.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Note The root node has depth 1.

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

NUMIIS

Description Number of IISs found. You should first query the `IISSOLSTATUS` attribute to make sure that the IIS procedure terminated successfully.

Type Integer

Category Attribute

Topic area Infeasibility

Set by routines `IIS`, `XPRSiisfirst`, `XPRSiisnext`, `XPRSiisall`.

OBJNAME

Description *This parameter is deprecated and will be removed in a future release. To find the name of the objective function use `XPRSgetnamelist`, passing `XPRS_NAMES_OBJECTIVE` in the `type` parameter.*

Type String

Category Attribute

Topic area Misc

Set by routines `XPRSreadprob`.

OBJRHS

Description Fixed part of the objective function.

Type Double

Category Attribute

Topic area Problem Information

Note If the matrix is in a presolved state, this attribute returns the fixed part of the objective in the **presolved** matrix. If you require the value for the original matrix, make sure you obtain the value when the matrix is not presolved. The `PRESOLVESTATE` attribute can be used to test if the matrix is presolved or not. See also 5.3. If an MPS file contains an objective function coefficient in the RHS then the negative of this will become `OBJRHS`.

Set by routines `XPRSchgobj`.

See also `LPOBJVAL`.

OBJVAL

Description Value of the objective function of the last problem solved with XPRSoptimize.

Type Double

Category Attribute

Topic areas Solution Process, Problem Information

Note Also set up by XPRSlpoptimize, XPRSmipoptimize and XPRSnlpoptimize, however is it best practice to access the relevant objective attribute LPOBJVAL, MIPOBJVAL or NLPOBJVAL when a specific problem type has been optimized.

Set by routines XPRSoptimize.

See also LPOBJVAL, MIPOBJVAL, NLPOBJVAL.

OBJSENSE

Description Sense of the optimization being performed.

Type Double

Category Attribute

Topic area Problem Information

Values

-1.0	For maximization problems.
1.0	For minimization problems.

Note The objective sense of a problem can be changed using XPRSchgobjsense.

Set by routines XPRSchgobjsense (CHGOBJSENSE).

OBJSTOSOLVE

Description Number of objectives that will be solved during the current multi-objective solve.

Type Integer

Category Attribute

Topic area Multiobjective

Note During a multi-objective solve, OBJSTOSOLVE will be equal to the number of distinct objective priorities, exluding objectives with zero weight (which are considered to be disabled). This can be fewer than OBJECTIVES, since all objectives with the same priority are combined into a single objective and solved at once.

Note Outside of a multi-objective solve, OBJSTOSOLVE is set to zero.

Set by routines XPRSoptimize.

See also OBJECTIVES, SOLVEDOBJS.

OBSERVEDPRIMALINTEGRAL

Description Value of the (observed) primal integral.

Type Double

Category Attribute

Topic area Branch and Bound Search

Note This attribute represents the integral of the primal gap over time. It measures the convergence of the primal bound `MIPBESTOBJVAL` over the whole solving time. The observed primal integral uses the best MIP solution found in a solve as a reference value. Consequently, different solves might use different reference values, and the observed primal integral values might not be readily comparable. If they do finish with the same `MIPBESTOBJVAL` (e.g., because both solved to optimality), the primal integrals are comparable, and the lower value indicates a better convergence of the best solution value. For details on the primal integral see Berthold: *Measuring the impact of primal heuristics*, OR Letters 41(6), pp. 611-614, 2013 as well as Berthold and Csizmadia: *The confined primal integral*, Mathematical Programming volume 188(2), pp. 523-537, 2021.

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `PRIMALDUALINTEGRAL`, `CPIALPHA`, `CPISCALEFACTOR`, `MIPBESTOBJVAL`.

OPTIMIZETYPEUSED

Description The type of solver used in the last call to `XPRSoptimize`, `XPRSmipoptimize`, `XPRSlpoptimize` or `XPRSnlpoptimize`.

Type Integer

Category Attribute

Topic area Solution Process

Values

`-1`	No solver was selected yet (`XPRS_OPTIMIZETYPE_NONE`). This can occur if the solve was interrupted while determining the problem type.
`0`	The LP solver was selected (`XPRS_OPTIMIZETYPE_LP`). The LP algorithm used by default is controlled by `DEFAULTALG`.
`1`	The MIP solver was selected (`XPRS_OPTIMIZETYPE_MIP`).
`2`	A local nonlinear solver was selected (`XPRS_OPTIMIZETYPE_LOCAL`). See XPRS_LOCALSOLVERSELECTED for which local solver was selected.
`3`	The global nonlinear solver was selected (`XPRS_OPTIMIZETYPE_GLOBAL`).

Note This attribute can be used to determine which type of solver was automatically selected when `XPRSoptimize` was called. If the global solver is enabled, this attribute can be used to determine whether a quadratic problem was found to be convex, since the global solver will only be selected if the problem is found to be non-convex.

Set by routines `XPRSoptimize`, `XPRSmipoptimize`, `XPRSlpoptimize`, `XPRSnlpoptimize`.

ORIGINALCOLS

Description Number of columns (i.e. variables) in the original matrix before presolving.

Type Integer

Category Attribute

Topic area Problem Information

Notes

1. If you require the value for the presolved matrix then use the `COLS` attribute.
2. During nonlinear solves, the `ORIGINALCOLS` attribute may include columns introduced by nonlinear reformulations, while `INPUTCOLS` will always refer to the number of columns in the problem that was originally input.

Set by routines `XPRSaddcols`, `XPRSdelcols`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `COLS`, `INPUTCOLS`, `ORIGINALROWS`.

ORIGINALGENCONS

Description Number of general constraints in the original problem before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note If you require the value for the presolved problem then use the `GENCONS` attribute.

Set by routines `XPRSaddgencons`, `XPRSdelgencons`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `GENCONS`, `ORIGINALGENCONCOLS`, `ORIGINALGENCONVALS`.

ORIGINALGENCONCOLS

Description Number of input variables in general constraints in the original problem before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note If you require the value for the presolved problem then use the `GENCONCOLS` attribute.

Set by routines `XPRSaddgencons`, `XPRSdelgencons`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `GENCONCOLS`, `ORIGINALGENCONS`, `ORIGINALGENCONVALS`.

ORIGINALGENCONVALS

Description Number of constant values in general constraints in the original problem before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note If you require the value for the presolved problem then use the `GENCONVALS` attribute.

Set by routines `XPRSaddgencons`, `XPRSdelgencons`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `GENCONVALS`, `ORIGINALGENCONS`, `ORIGINALGENCONCOLS`.

ORIGINALINDICATORS

Description Number of indicator constraints in the original matrix before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note If you require the value for the presolved matrix then use the INDICATORS attribute.

Set by routines XPRSsetindicators, XPRSdelindicators, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also INDICATORS.

ORIGINALMIPENTS

Description Number of MIP entities (i.e. binary, integer, semi-continuous, partial integer, and semi-continuous integer variables) but excluding the number of special ordered sets in the original matrix before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note If you require the value for the presolved matrix then use the MIPENTS attribute.

Set by routines XPRSchgcoltype, XPRSdelcols, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also ORIGINALSETS, MIPENTS.

ORIGINALPWLS

Description Number of piecewise linear constraints in the original problem before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note If you require the value for the presolved problem then use the PWLCONS attribute.

Set by routines XPRSaddpwlcons, XPRSdelpwlcons, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also PWLCONS, ORIGINALPWLPOINTS.

ORIGINALPWLPOINTS

Description Number of breakpoints of piecewise linear constraints in the original problem before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note If you require the value for the presolved problem then use the `PWLPOINTS` attribute.

Set by routines `XPRSaddpwlcons`, `XPRSdelpwlcons`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `PWLPOINTS`, `ORIGINALPWLS`.

ORIGINALQCELEMS

Description Number of quadratic row coefficients in the original matrix before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Note If you require the value for the presolved matrix then use the `QCELEMS` attribute.

Set by routines `XPRSaddqmatrix`, `XPRSdelqmatrix`, `XPRSchgqrowcoeff`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `ORIGINALQCONSTRAINTS`, `QCELEMS`.

ORIGINALQCONSTRAINTS

Description Number of rows with quadratic coefficients in the original matrix before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Note If you require the value for the presolved matrix then use the `QCONSTRAINTS` attribute.

Set by routines `XPRSaddqmatrix`, `XPRSdelqmatrix`, `XPRSchgqrowcoeff`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `ORIGINALQCELEMS`, `QCONSTRAINTS`.

ORIGINALQELEMS

Description Number of quadratic non-zeros in the original objective before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Notes

1. Xpress ensures that the matrix is symmetric, so off-diagonal elements are counted twice as follows: if `i` is different from `j` and you specify a non-zero at `(i,j)`, then the implicit element at `(j,i)` will also be counted, for a total of 2 non-zeros.
2. If you require the value for the presolved matrix then use the `QELEMS` attribute.

Set by routines `XPRSaddqmatrix`, `XPRSdelqmatrix`, `XPRSchgqrowcoeff`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `QELEMS`.

ORIGINALROWS

Description Number of rows (i.e. constraints) in the original matrix before presolving.

Type Integer

Category Attribute

Topic area Problem Information

Notes

1. If you require the value for the presolved matrix then use the `ROWS` attribute.
2. During nonlinear solves, the `ORIGINALROWS` attribute may include rows introduced by nonlinear reformulations, while `INPUTROWS` will always refer to the number of rows in the problem that was originally input.

Set by routines `XPRSaddrows`, `XPRSdelrows`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `ROWS`, `INPUTROWS`, `ORIGINALCOLS`.

ORIGINALSETMEMBERS

Description Number of variables within special ordered sets (set members) in the original matrix before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note If you require the value for the presolved matrix then use the `SETMEMBERS` attribute.

Set by routines `XPRSaddsets`, `XPRSdelsets`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `SETMEMBERS`, `ORIGINALSETS`.

ORIGINALSETS

Description Number of special ordered sets in the original matrix before presolving.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note If you require the value for the presolved matrix then use the `SETS` attribute.

Set by routines `XPRSaddsets`, `XPRSdelsets`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `ORIGINALSETMEMBERS`, `ORIGINALMIPENTS`, `SETS`.

PARENTNODE

Description The parent node of the current node in the tree search.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

PEAKMEMORY

Description An estimate of the peak amount of dynamically allocated heap memory by the problem.

Type Integer

Category Attribute

Topic area Memory

Note On 64-bit systems this is a 64-bit integer, use `XPRSgetintattrib64` to retrieve its value.

See also `CURRENTMEMORY`,`SYSTEMMEMORY`.

PEAKTOTALTREEMEMORYUSAGE

Description The peak size, in megabytes, that the branch-and-bound search tree reached during the solve. Note that this value will include the uncompressed size of any compressed data and the size of any data saved to the tree file.

Type Integer

Category Attribute

Topic areas Branch and Bound Search, Memory

Set by routines `XPRSmipoptimize`.

See also `TREEMEMORYUSAGE`.

PENALTYVALUE

Description The weighted sum of violations in the solution to the relaxed problem identified by the infeasibility repair function.

Type Double

Category Attribute

Topic area Infeasibility

Set by routines XPRSrepairinfeas (REPAIRINFEAS), XPRSrepairweightedinfeas.

PHYSICALCORESDETECTED

Description The total number of physical cores across all CPUs detected by the optimizer.

Type Integer

Category Attribute

Topic area System

Values >=1 Detected number of physical cores.

Set by routines XPRSinit.

See also CORESDETECTED, CORESPERCPUDETECTED, CPUSDETECTED, PHYSICALCORESPERCPUDETECTED.

PHYSICALCORESPERCPUDETECTED

Description The number of physical cores per CPU detected by the optimizer.

Type Integer

Category Attribute

Topic area System

Values >=1 Detected number of physical cores per CPU.

Set by routines XPRSinit.

See also CORESDETECTED, CORESPERCPUDETECTED, CPUSDETECTED, PHYSICALCORESPERCPUDETECTED.

PREDICTEDATTLEVEL

Description A measure between 0 and 1 to predict how numerically unstable the current MIP solve can be expected to be. After the root LP solve, a machine learning model is used to predict the actual ATTENTIONLEVEL which will only be computed if MIPKAPPAFREQ is set to a nonzero value. If the predicted attention level exceeds a value of 0.1, a message will be printed to the log.

Type Double

Category Attribute

Topic area Numerics

Set by routines XPRSmipoptimize.

See also ATTENTIONLEVEL, MAXKAPPA.

PRESOLVEINDEX

Description Presolve: The row or column index on which presolve detected a problem to be infeasible or unbounded.

Type Integer

Category Attribute

Topic area Presolve

Note Row indices are in the range 0 to ROWS−1, and column indices are in the range ROWS+SPAREROWS to ROWS+SPAREROWS+COLS−1.

Set by routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

PRESOLVESTATE

Description Problem status as a bit-vector (compare Section 9.2).

Type Integer

Category Attribute

Topic areas Presolve, Bit-vector

Values

Bit	Meaning
0	Problem has been loaded.
1	Problem has been LP presolved.
2	Problem has been MIP presolved.
7	Solution in memory is valid.

Note Other bits are reserved.

Set by routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

PRIMALDUALINTEGRAL

Description Value of the primal-dual integral.

Type Double

Category Attribute

Topic area Misc

Note This attribute represents the integral of the primal-dual gap over time. It measures the convergence of the best (dual) bound BESTBOUND and the primal bound MIPBESTOBJVAL over the whole solving time. Lower values are better. For details on the primal(-dual) integral see Berthold: *Measuring the impact of primal heuristics*, OR Letters 41(6), pp. 611-614, 2013.

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

See also BESTBOUND, MIPBESTOBJVAL.

PRIMALINFEAS

Description Number of primal infeasibilities.

Type Integer

Category Attribute

Topic area LP

Note If the matrix is in a presolved state, this attribute returns the number of primal infeasibilities in the **presolved** matrix. If you require the value for the original matrix, make sure you obtain the value when the matrix is not presolved. The PRESOLVESTATE attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines XPRSlpoptimize (LPOPTIMIZE).

See also SUMPRIMALINF, DUALINFEAS, MIPINFEAS.

PWLCONS

Description Number of piecewise linear constraints in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note Piecewise linear constraints are transformed during presolve, so PWLCONS will always be zero in a presolved problem. If you require the value for the original problem then use the ORIGINALPWLS attribute instead.

Set by routines XPRSaddpwlcons, XPRSdelpwlcons, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also ORIGINALPWLS, PWLPOINTS.

PWLPOINTS

Description Number of breakpoints of piecewise linear constraints in the problem.

Type Integer

Category Attribute

Topic areas Problem Information, Piecewise Linear and General Constraints

Note Piecewise linear constraints are transformed during presolve, so PWLPOINTS will always be zero in a presolved problem. If you require the value for the original problem then use the ORIGINALPWLPOINTS attribute instead.

Set by routines XPRSaddpwlcons, XPRSdelpwlcons, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also ORIGINALPWLPOINTS, PWLCONS.

QCELEMS

Description Number of quadratic row coefficients in the matrix.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Note If the matrix is in a presolved state, this attribute returns the number of quadratic row coefficients in the **presolved** matrix. If you require the value for the original matrix then use the `ORIGINALQCELEMS` attribute instead. The `PRESOLVESTATE` can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines `XPRSaddqmatrix`, `XPRSdelqmatrix`, `XPRSchgqrowcoeff`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `QCONSTRAINTS`, `ORIGINALQCELEMS`.

QCONSTRAINTS

Description Number of rows with quadratic coefficients in the matrix.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Note If the matrix is in a presolved state, this attribute returns the number of rows with quadratic coefficients in the **presolved** matrix. If you require the value for the original matrix then use the `ORIGINALQCONSTRAINTS` attribute instead. The `PRESOLVESTATE` attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines `XPRSaddqmatrix`, `XPRSdelqmatrix`, `XPRSchgqrowcoeff`, `XPRSloadlp`, `XPRSloadmip`, `XPRSloadqp`, `XPRSloadqcqp`, `XPRSloadmiqp`, `XPRSloadmiqcqp`, `XPRSreadprob`.

See also `QCELEMS`, `ORIGINALQCONSTRAINTS`.

QELEMS

Description Number of quadratic non-zeros in the objective.

Type Integer

Category Attribute

Topic areas Problem Information, Quadratic

Notes

1. Xpress ensures that the matrix is symmetric, so off-diagonal elements are counted twice as follows: if `i` is different from `j` and you specify a non-zero at `(i,j)`, then the implicit element at `(j,i)` will also be counted, for a total of 2 non-zeros.
2. If the matrix is in a presolved state, this attribute returns the number of quadratic elements in the **presolved** matrix. If you require the value for the original matrix then use the `ORIGINALQELEMS` attribute instead. The `PRESOLVESTATE` attribute can be used to test if the matrix is presolved or not.See also 5.3.

Set by routines XPRSchgqobj, XPRSchgmqobj, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also ORIGINALQELEMS.

RANGENAME

Description Active range name.

Type String

Category Attribute

Topic area File IO

Set by routines XPRSreadprob.

RESTARTS

Description Total number of restarts performed.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Note As opposed to TREERESTARTS this not only includes the number of times the branch-and-bound tree is reset to the root node, but also includes the number of times presolve is repeated at the root node. If your application caches information about the presolved model then it has to refresh this information whenever the value of this attribute changed.

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

See also TREERESTARTS.

RHSNAME

Description Active right hand side name.

Type String

Category Attribute

Topic area File IO

Set by routines XPRSreadprob.

ROWS

Description Number of rows (i.e. constraints) in the matrix.

Type Integer

Category Attribute

Topic area Problem Information

Note If the matrix is in a presolved state, this attribute returns the number of rows in the **presolved** matrix. If you require the value for the original matrix then use the INPUTROWS attribute instead. The PRESOLVESTATE attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines XPRSaddrows, XPRSdelrows, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE), XPRSreadprob.

See also COLS, INPUTROWS.

SETMEMBERS

Description Number of variables within special ordered sets (set members) in the matrix.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note If the matrix is in a presolved state, this attribute returns the number of variables within special ordered sets in the **presolved** matrix. If you require the value for the original matrix then use the ORIGINALSETMEMBERS attribute instead. The PRESOLVESTATE attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines XPRSaddsets, XPRSdelsets, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also SETS, ORIGINALSETMEMBERS.

SETS

Description Number of special ordered sets in the matrix.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Note If the matrix is in a presolved state, this attribute returns the number of special ordered sets in the **presolved** matrix. If you require the value for the original matrix then use the ORIGINALSETS attribute instead. The PRESOLVESTATE attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines XPRSaddsets, XPRSdelsets, XPRSloadlp, XPRSloadmip, XPRSloadqp, XPRSloadqcqp, XPRSloadmiqp, XPRSloadmiqcqp, XPRSreadprob.

See also SETMEMBERS, MIPENTS, ORIGINALSETS.

SIMPLEXITER

Description Number of simplex iterations performed.

Type Integer

Category Attribute

Topic areas LP, Simplex

Note In case of a barrier solve with crossover, this will only include simplex iterations performed during cleanup. For the number of crossover iterations, refer to `CROSSOVERITER`.

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `CROSSOVERITER`.

SOLSTATUS

Description Status of the solution of the last problem solved with XPRSoptimize.

Type Integer

Category Attribute

Topic areas Solution Process, Solution

Values

0	`XPRS_SOLSTATUS_NOTFOUND`	No solution available.
1	`XPRS_SOLSTATUS_OPTIMAL`	An optimal solution has been found.
2	`XPRS_SOLSTATUS_FEASIBLE`	A solution that is not proven optimal is found.
3	`XPRS_SOLSTATUS_INFEASIBLE`	No solution exists.
4	`XPRS_SOLSTATUS_UNBOUNDED`	The problem is unbounded, if feasible.

Note Also set up by `XPRSlpoptimize`, `XPRSmipoptimize` and `XPRSnlpoptimize`, however is it best practice to access the relevant status by checking `LPSTATUS`, `MIPSTATUS` or `NLPSTATUS` when a specific problem type has been optimized.

SOLVEDOBJS

Description Number of objectives that have been solved so far during a multi-objective solve.

Type Integer

Category Attribute

Topic area Multiobjective

Note After an optimal multi-objective solve, `SOLVEDOBJS` will be equal to the number of distinct objective priorities, exluding objectives with zero weight (which are considered to be disabled). This can be fewer than `OBJECTIVES`, since all objectives with the same priority are combined into a single objective and solved at once. During multi-objective optimization, `0 <= SOLVEDOBJS <= OBJSTOSOLVE`.

Note After a single-objective solve, `SOLVEDOBJS` is set to `1`.

Set by routines `XPRSoptimize`.

See also `OBJECTIVES`, `OBJSTOSOLVE`.

SOLVESTATUS

Description Status of the solve of the last problem solved with XPRSoptimize.

Type Integer

Category Attribute

Topic area Solution Process

Values

0	XPRS_SOLVESTATUS_UNSTARTED	The solve has not been started.
1	XPRS_SOLVESTATUS_STOPPED	Optimization has been interrupted.
2	XPRS_SOLVESTATUS_FAILED	Optimization has run into a nonrecoverable problem and failed.
3	XPRS_SOLVESTATUS_COMPLETED	Search completed.

Note Also set up by XPRSlpoptimize, XPRSmipoptimize and XPRSnlpoptimize, however is it best practice to access the relevant status by checking LPSTATUS, MIPSTATUS or NLPSTATUS when a specific problem type has been optimized.

SPARECOLS

Description Number of spare columns in the matrix.

Type Integer

Category Attribute

Topic area Problem Information

Set by routines XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp, XPRSreadprob.

SPAREELEMS

Description Number of spare matrix elements in the matrix.

Type Integer

Category Attribute

Topic area Problem Information

Set by routines XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp, XPRSreadprob.

SPAREMIPENTS

Description Number of spare MIP entities in the matrix.

Type Integer

Category Attribute

Topic areas Problem Information, MIP Entities

Set by routines XPRSloadmip, XPRSloadlp, XPRSloadmiqp, XPRSloadqp, XPRSreadprob.

SPAREROWS

Description	Number of spare rows in the matrix.
Type	Integer
Category	Attribute
Topic area	Problem Information
Set by routines	`XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSreadprob`.

SPARESETELEMS

Description	Number of spare set elements in the matrix.
Type	Integer
Category	Attribute
Topic areas	Problem Information, MIP Entities
Set by routines	`XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSreadprob`.

SPARESETS

Description	Number of spare sets in the matrix.
Type	Integer
Category	Attribute
Topic areas	Problem Information, MIP Entities
Set by routines	`XPRSloadmip`, `XPRSloadlp`, `XPRSloadmiqp`, `XPRSloadqp`, `XPRSreadprob`.

STOPSTATUS

Description	Status of the optimization process.
Type	Integer
Category	Attribute
Topic area	Solution Process
Note	Possible values are:

Value	Description
XPRS_STOP_NONE	no interruption - the solve completed normally
XPRS_STOP_TIMELIMIT	time limit hit
XPRS_STOP_WORKLIMIT	work limit hit
XPRS_STOP_CTRLC	user hit control C or sent a corresponding signal
XPRS_STOP_NODELIMIT	node limit hit in branch-and-bound
XPRS_STOP_ITERLIMIT	LP iteration limit hit
XPRS_STOP_MIPGAP	MIP gap is sufficiently small
XPRS_STOP_SOLLIMIT	solution limit hit
XPRS_STOP_MEMORYERROR	insufficient memory led to error in memory allocation
XPRS_STOP_USER	user interrupt
XPRS_STOP_SOLVECOMPLETE	MIP solve completed without fully exploring the tree
XPRS_STOP_LICENSELOST	license has been lost, all flexible licenses are in use
XPRS_STOP_NUMERICALERROR	solver encountered numerical errors

Note XPRS_STOP_SOLVECOMPLETE can occur when a disconnected component is found to be infeasible, proving the whole problem to be infeasible without finishing the branch and bound search.

Note XPRS_STOP_NUMERICALERROR can occur during a multi-objective solve, when the first objective solves to optimality, but when solving a subsequent objective the problem is found to be infeasible due to numerical issues associated caused by fixing the earlier objectives.

Set by routines XPRSlpoptimize (LPOPTIMIZE), XPRSmipoptimize (MIPOPTIMIZE).

SUMPRIMALINF

Description Scaled sum of primal infeasibilities.

Type Double

Category Attribute

Topic areas LP, Numerics

Note If the matrix is in a presolved state, this attribute returns the scaled sum of primal infeasibilities in the **presolved** matrix. If you require the value for the original matrix, make sure you obtain the value when the matrix is not presolved. The PRESOLVESTATE attribute can be used to test if the matrix is presolved or not. See also 5.3.

Set by routines XPRSlpoptimize (LPOPTIMIZE).

See also PRIMALINFEAS.

SYSTEMMEMORY

Description The amount of non problem specific memory used by the solver.

Type Integer

Category Attribute

Topic areas System, Memory

Note On 64-bit systems this is a 64-bit integer, use `XPRSgetintattrib64` to retrieve its value.

See also `CURRENTMEMORY`, `PEAKMEMORY`.

TIME

Description Time spent solving the problem as measured by the optimizer.

Type Double

Category Attribute

Topic area Solution Process

Note This attribute changed type with Xpress Optimizer version 9.0 (from integer to double). It is strongly recommended to adapt all query calls to this attribute accordingly.

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `TIMELIMIT`, `SOLTIMELIMIT`.

TOTALMEMORY

Description The amount of dynamically allocated heap memory by the optimizer, including all problems currently exsisting.

Type Integer

Category Attribute

Topic area Memory

Note On 64-bit systems this is a 64-bit integer, use `XPRSgetintattrib64` to retrieve its value.

See also `PEAKMEMORY`, `CURRENTMEMORY`.

TREECOMPLETION

Description Estimation of the relative completion of the search tree as fraction between 0 and 1. Its accuracy mainly depends on the level of degeneracy of a problem and the balancedness of the search tree.

Type Double

Category Attribute

Topic area Branch and Bound Search

Set by routines `XPRSmipoptimize` (`MIPOPTIMIZE`).

TREEFILESIZE

Description The allocated size of the tree file, in megabytes. Because data can be removed from the tree file during the branch and bound search, the size of the tree file is usually greater than the amount of data currently within it (represented by the `TREEFILEUSAGE` attribute).

Type Integer

Category Attribute

Topic areas Branch and Bound Search, Memory

See also TREEFILEUSAGE.

TREEFILEUSAGE

Description The number of megabytes of data from the branch-and-bound tree that have been saved to the tree file. Note that the actual allocated size of the tree file (represented by the TREEFILESIZE attribute) may be greater than this value.

Type Integer

Category Attribute

Topic areas Branch and Bound Search, Memory

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

See also TREEFILESIZE, TREEMEMORYLIMIT.

TREEMEMORYUSAGE

Description The amount of physical memory, in megabytes, currently being used to store the branch-and-bound search tree.

Type Integer

Category Attribute

Topic areas Branch and Bound Search, Memory

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

See also TREEMEMORYLIMIT, TREEFILEUSAGE.

TREERESTARTS

Description Number of in-tree restarts performed.

Type Integer

Category Attribute

Topic area Branch and Bound Search

Set by routines XPRSmipoptimize (MIPOPTIMIZE).

See also RESTARTS.

UUID

Description Universally Unique Identifier for the problem instance.

Type String

Category Attribute

Topic area Misc

WORK

Description Amount of deterministic algorithmic "work" spent since the invocation of the search process. Work is measured in work units. In contrast with `TIME`, `WORK` is deterministic. It is also independent of the platform and hardware used. At the end of a search, the optimizer reports both the accumulated total work as well as the speed at which WORK has been accumulated (measured in work units per second) in the log file. The accumulation of work depends on the size of the model being solved, the involved algorithms, and the stage of the search.

Type Double

Category Attribute

Topic areas Solution Process, Determinism

Note Work is independent of the platform and machine used. In contrast to wall clock time and the related controls, work is deterministic, i.e., multiple invocations of the optimizer will always yield the same amount of work unless a time limit has been hit.

Set by routines `XPRSlpoptimize` (`LPOPTIMIZE`), `XPRSmipoptimize` (`MIPOPTIMIZE`).

See also `WORKLIMIT`.

XPRESSVERSION

Description The Xpress version number.

Type String

Category Attribute

Topic area Misc

Note The version number of Xpress.

CHAPTER 11

Return Codes and Error Messages

11.1 Optimizer Return Codes

When the Console Optimizer terminates after the STOP command, it may set an exit code that can be tested by the operating system or by the calling program. The exit code is set as follows:

Return Code	Description
0	Program terminated normally (with STOP).
63	LP optimization unfinished.
64	LP feasible and optimal.
65	LP infeasible.
66	LP unbounded.
67	IP optimal solution found.
68	IP search incomplete but an IP solution has been found.
69	IP search incomplete, no IP solution found.
70	IP infeasible.
99	LP optimization not started.
255	Xpress Optimizer has not been initialized.

11.2 Optimizer Error and Warning Messages

Following an error, the Optimizer can be interrogated as necessary to obtain more information about the specific error which occurred. Library functions will return a non-zero error code to indicate that an error occurred. The last error code can also be found by querying the ERRORCODE attribute. A description of the last error can be retrieved using the function XPRSgetlasterror. Errors and warnings are also printed to the screen (for Console users) and the message callback (for library users) when they are encountered.

The table below shows the values that can be returned by a library function or found in ERRORCODE, and a possible resolution of the error or warning.

3 ***Extension not allowed - ignored.***

The specified extension is not allowed. The Optimizer ignores the extension and truncates the filename.

5 ***Error on .<ext> file.***

An error has occurred on the . <ext> file. Please make sure that there is adequate disk space for the file and that it has not become corrupted.

6 ***No match for column <col> in matrix.***

Column <col> has not been defined in the `COLUMNS` section of the matrix and cannot be used in subsequent sections. Please check that the spelling of <col> is correct and that it is not written outside the field reserved for column names.

7 ***Empty matrix. Please increase `EXTRAROWS`.***

There are too few rows or columns. Please increase `EXTRAROWS` before input, or make sure there is at least one row in your matrix and try to read it again.

9 ***Error on read of basis file.***

The basis file `.BSS` is corrupt. Please make sure that there is adequate disk space for the file and that it has not been corrupted.

11 ***Not allowed - solution not optimal.***

The operation you are trying to perform is not allowed unless the solution is optimal. Please call `XPRSmaxim` (MAXIM) or `XPRSminim` (`MINIM`) to optimize the problem and make sure the process is completed. If the control `LPITERLIMIT` has been set, make sure that the optimal solution can be found within the maximum number of iterations allowed.

17 ***Critical numerical error - unrecoverable.***

Numerical breakdown, the problem is poorly scaled and/or badly conditioned.

18 ***Bound conflict for column <col>.***

Specified upper bound for column <col> is smaller that the specified lower bound. Please change one or both bounds to solve the conflict and try again.

19 ***Eta overflow straight after invert - unrecoverable.***

There is not enough memory for eta arrays. Either increase the virtual paging space or the physical memory.

20 ***Insufficient memory for array <array>.***

There is not enough memory for an internal data structure. Either increase the virtual paging space or the physical memory.

21 ***Unidentified section The command is not recognized by the Optimizer.***

Please check the spelling and try again. Please refer to the Reference Manual for a list of valid commands.

29 ***Input aborted.***

Input has encountered too many problems in reading your matrix and it has been aborted. This message will be preceded by other error messages whose error numbers will give information about the nature of each of the problems. Please correct all errors and try again.

36 ***Linear Optimizer only***

You are only authorized to use the Linear Optimizer. Please contact your local sales office to discuss upgrading to the IP Optimizer if you wish to use this command.

38 ***Invalid option.***

One of the options you have specified is incorrect. Please check the input option and retype the command. A list of valid options for each command can be found in 8.

41 ***Unspecified MIP error - contact the Xpress support team.***

Internal error. Please contact your local support office.

45 ***Failure to open tree file - aborting. (Perhaps disk is full).***

The Optimizer cannot open the `.GLB` file. This usually occurs when your disk is full. If this is not the case it means that the `.GLB` file has been corrupted.

50 ***Inconsistent basis.***

Internal basis held in memory has been corrupted. Please contact your local support office.

52 ***Too many non-zero elements.***

The number of matrix elements exceeds the maximum allowed. Please contact your local sales office to discuss upgrading your license to solve larger problems.

56 ***Reference row entries too close for set <set> member <col>.***

The coefficient of column <col> in the constraint being used as reference row for set <set> is too close to the coefficient of some other column in the reference row. Please make sure the coefficients in the reference row differ enough from one another. One way of doing this is to create a non computational constraint (`N` type) that contains all the variables members of the set <set> and then assign coefficients whose distance from each other is of at least 1 unit.

58 ***Duplicate element for column <col> row <row>.***

The coefficient for column <col> appears more than once in row <row>. The elements are added together but please make sure column <col> only has one coefficient in <row> to avoid this warning message.

61 ***Unexpected EOF on workfile.***

An internal workfile has been corrupted. Please make sure that there is adequate disk space and try again. If the problem persists please contact your local support office.

64 ***Error closing file <file>.***

The Optimizer could not close file <file>. Please make sure that the file exists and that it is not being used by another application.

65 ***Fatal error on read from workfile <file> - program aborted.***

An internal workfile has been corrupted. Please make sure that your disk has enough space and try again. If the problem persists please contact your local support office.

66 ***Unable to open file <file>.***

The Optimizer has failed to open the file <file>. Please make sure that the file exists and there is adequate disk space.

67 ***Error on read of file <file>.***

The Optimizer has failed to read the file <file>. Please make sure that the file exists and that it has not been corrupted.

68 ***Error on write of file <file>.***

The Optimizer has failed to write the file <file>. Please make sure that there is adequate disk space, that the volume is not corrupt, and that the directory is not write-protected.

69 ***Path is too long: <path>.***

The path <path> cannot be opened in the filesystem because it is too long.

73 ***Problem has too many rows. The maximum is <num>***

The Optimizer cannot solve your problem since the number of rows exceeds the maximum allowed. Please contact your local sales office to discuss upgrading your license to solve larger problems.

76 ***Illegal priority: entity <ent> value <num>.***

Entity <ent> has been assigned an invalid priority value of <num> in the directives files and this priority will be ignored. Please make sure that the priority value lies between 0 and 1000 and that it is written inside the corresponding field in the `.DIR` file.

77 ***Illegal set card <line>.***

The set definition in line <line> of the `.MAT` or .MPS file creates a conflict. Please make sure that the set has a correct type and has not been already defined. Please refer to the Reference Manual for a list of valid set types.

80 ***File creation error.***

The Optimizer cannot create a file. Please make sure that these is adequate disk space and that the volume is not corrupt.

81 ***Fatal error on write to workfile <file> - program aborted.***

The Optimizer cannot write to the file <file>. Please make sure that there is adequate disk space and that the volume is not corrupt.

83 ***Fatal error on write to file - program aborted.***

The Optimizer cannot write to an internal file. Please make sure that there is adequate disk space and that the volume is not corrupt.

84 ***Input line too long. Maximum line length is <num>***

A line in the `.MAT` or `.MPS` file has been found to be too long. Please reduce the length to be less or equal than <num> and input again.

85 ***File not found: <file>.***

The Optimizer cannot find the file <file>. Please check the spelling and that the file exists. If this file has to be created by the Optimizer, make sure that the process which creates the file has been performed.

86 ***Problem name attribute must not be `NULL`.***

This error occurs if a user attempts to read a problem file or set a problem name with a `NULL` argument. Please make sure to always use a non-NULL string argument for problem and file names in functions `XPRSreadprob`, `XPRSsetprobname`, and `XPRSgetprobname`.

89 ***No optimization has been attempted.***

The operation you are trying to perform is not allowed unless the solution is optimal. Please call `XPRSmaxim` (`MAXIM`) or `XPRSminim` (`MINIM`) to optimize the problem and make sure the process is completed. If you have set the control `LPITERLIMIT` make sure that the optimal solution can be found within the maximum number of iterations allowed.

90 ***Infeasibility repair cannot be called for problems with multiple objectives.***

You can disable the multi-objective solve and call infeasibility repair again.

91 ***No problem has been input.***

An operation has been attempted that requires a problem to have been input. Please make sure that `XPRSreadprob` (`READPROB`) is called and that the problem has been loaded successfully before trying again.

97 ***Split vector <vector>.***

The declaration of column <vector> in the `COLUMN` section of the `.MAT` or `.MPS` file must be done in contiguous line. It is not possible to interrupt the declaration of a column with lines corresponding to a different vector.

98 ***At line <num> no match for row <row>.***

A non existing row <row> is being used at line number <num> of the `.MAT` or `.MPS` file. Please check spelling and make sure that <row> is defined in the `ROWS` section.

102 ***Eta file space exceeded - optimization aborted.***

The Optimizer requires more memory. Please increase your virtual paging space or physical memory and try to optimize again.

107 ***Problem has too many MIP entities. The maximum is <num>***

The Optimizer cannot solve your problem since the number of MIP entities exceeds the maximum allowed. Please contact your local sales office to discuss upgrading your license to solve larger problems.

111 ***Duplicate row <row> - ignored.***

Row <row> is used more than once in the same section. Only the first use is kept and subsequent ones are ignored.

112 ***Postoptimal analysis not permitted on presolved problems.***

Re-optimize with `PRESOLVE = 0`. An operation has been attempted on the presolved problem. Please optimize again calling `XPRSmaxim` (`MAXIM`), `XPRSminim` (`MINIM`) with the `l` flag or turning presolve off by setting `PRESOLVE` to 0.

113 ***Unable to restore version <ver> save files.***

The svf file was created by a different version of the Optimizer and cannot be restored with this version.

114 ***Fatal error - pool hash table full at vector <vector>.***

Internal error. Please contact your local support office.

120 ***Problem has too many rows and columns. The maximum is <num>***

The Optimizer cannot solve your problem since the number of rows plus columns exceeds the maximum allowed. Please contact your local sales office to discuss upgrading your license to solve larger problems.

122 ***Corrupt solution file.***

Solution file `.SOL` could not be accessed. Please make sure that there is adequate disk space and that the file is not being used by another process.

124 ***Invalid parameter value passed to <function>. Parameter value <param_name> is not allowed***

A parameter lookup by name has failed. The provided parameter name does not match any parameters in Xpress.

127 ***Not found: <vector>.***

An attempt has been made to use a row or column <vector> that cannot be found in the problem. Please check spelling and try again.

128 ***Cannot load directives for problem with no MIP entities.***

The problem does not have MIP entities and so directives cannot be loaded.

129 ***Access denied to problem state : '<name>' (<routine>).***

The user is not licensed to have set or get access to problem control (or attribute) <name>. The routine used for access was <routine>.

131 ***No column: <col>.***

Column <col> used in basis file `.BSS` does not exist in the problem. A new basis will be created internally from where column <col> will have been removed and the rest of columns and rows will maintain their basic/non-basic status.

132 ***No row: <row>.***

Row <row> used in basis file `.BSS` does not exist in the problem. A new basis will be created internally from where row <row> will have been removed and the rest of columns and rows will maintain their basic/non-basic status.

136 ***Cannot access control <control_name> via attribute routine <function>***

When accessing controls and attributes, the API function called must be matched appropriately to the type (double, int, string) and access type (control / attribute) of the parameter.

137 ***Bad internal state found in 'Struct' lookup : <parameter_table> (<parameter_name>)***

A parameter provided could not be found in the parameters table. This is an internal error, please contact FICO support.

140 ***Basis lost - recovering.***

The number of rows in the problem is not equal to the number of basic rows + columns in the problem, which means that the existing basis is no longer valid. This will be detected when re-optimizing a problem that has been altered in some way since it was last optimized (see below). A correct basis is generated automatically and no action needs to be taken. The basis can be lost in two ways: (1) if a row is deleted for which the slack is non-basic: the number of rows will decrease by one, but the number of basic rows + columns will be unchanged. (2) if a basic column is deleted: the number of basic rows + columns will decrease by one, but the number of rows will be unchanged. You can avoid losing the basis by only deleting rows for which the slack is basic, and columns which are non-basic. (The `XPRSgetbasis` function can be used to determine the basis status.) To delete a non-basic row without losing the basis, bring it into the basis first, and to delete a basic column without losing the basis, take it out of the basis first - the functions `XPRSgetpivots` and `XPRSpivot` may be useful here. However, remember that the message is only a warning and the Optimizer will generate a new basis automatically if necessary.

143 ***No entity <ent>.***

Entity <ent> used in directives file `.DIR` cannot be found in the problem and its corresponding priority will be ignored. Please check spelling and that the column <ent> is actually declared as an entity in the `BOUNDS` section or is a set member.

151 ***Illegal MARKER.***

The line marking the start of a set of integer columns or a set of columns belonging to a Special Ordered Set in the `.MPS` file is incorrect.

152 ***Unexpected EOF.***

The Optimizer has found an unexpected EOF marker character. Please check that the input file is correct and input again.

153 ***Illegal card at line <line>.***

Line <line> of the `.MPS` file could not be interpreted. Please refer to the Reference Manual for information about the valid MPS format.

155 ***Cannot access control '<id>' via attribute routine <routine>.***

Controls cannot be accessed from attribute access routines.

156 ***Cannot access attribute '<id>' via control routine <routine>.***

Attributes cannot be accessed from control access routines.

157 ***Cannot access attribute <attribute_name> via control routine <routine>.***

Attributes cannot be accessed from control access routines.

158 ***Unrecognized callback name <callback> (<function>)***

The callback name provided to the API function is not recognized.

159 ***Failed to set default controls.***

Attempt failed to set controls to their defaults.

160 ***Cannot access <typename> type '<id>' via routine <routine>.***

Accessing an attribute or control requires using a routine with matching type.

161 ***Cannot access <typename> type '<name>' via routine <routine>.***

Accessing an attribute or control requires using a routine with matching type.

162 ***Recording and playback error : <info>.***

An error occurred in the recording and playback tool.

163 ***Failed to copy controls.***

Attempt failed to copy controls defined for one problem to another.

164 ***Problem is not presolved.***

Action requires problem to be presolved and the problem is not presolved.

167 ***Failed to allocate memory of size <bytes> bytes.***

The Optimizer failed to allocate required memory of size <bytes>.

168 ***Required resource not currently available : '<name>'.***

The resource <name> is required by an action but is unavailable.

169 ***Failed to create resource : '<name>'.***

The resource <name> failed to create.

170 ***Corrupt tree file.***

Tree file `.GLB` cannot be accessed. Please make sure that there is adequate disk space and that the file is not being used by another process.

171 ***Invalid row type for row <row>.***

`XPRSalter` (`ALTER`) cannot change the row type of <row> because the new type is invalid. Please correct and try again.

173 ***Name not recognized : <name>.***

The control name cannot be recognized.

178 ***Not enough spare rows to remove all violations.***

The Optimizer could not add more cuts to the matrix because there is not enough space. Please increase `EXTRAROWS` before input to improve performance.

179 ***Load MIP solution failed : '<status description>'.***

Attempt failed to load MIP solution into the Optimizer. See <status description» for details of the failure.

180 ***No change to this control allowed.***

The Optimizer does not allow changes to this control.

181 ***Cannot alter bound on BV, SC, UI, PI, or set member.***

`XPRSalter` (`ALTER`) cannot be used to change the upper or lower bound of a variable if its variable type is binary, semi-continuous, integer, partial integer, semi-continuous integer, or if it is a set member.

186 ***Inconsistent number of variables in problem.***

A compact format basis is being read into a problem with a different number of variables than the one for which the basis was created.

187 ***Unable to restore alternative system <system> save files.***

The svf file was created on a different operating system and cannot be restored on the current system.

188 ***Unable to restore - save file is invalid.***

The svf file was not in the expected format. It may have been corrupted.

191 ***Solution in file '<file>' (rows:<nrow>, cols:<ncol>) not compatible with problem.***

The size of the loaded problem is not compatible with problem size from the solution file.

192 ***Bad flags <flag string>.***

A flag string passed into a command line call is invalid.

193 ***Possible unexpected results from `XPRSreadbinsol` (`READBINSOL`): <message>.***

A call to the `XPRSreadbinsol` (`READBINSOL`) may produce unexpected results. See <message> for details.

195 ***Cannot read LP solution into presolved problem.***

An LP solution cannot be read into a problem in a presolved state.

197 ***Failed to register callback for event : '<event>'.***

Registering callback for an event failed.

202 ***Control parser: <error>.***

A generic control parser error, for example a memory allocation failure.

243 ***The Optimizer requires a newer version of the XPRL library.***

You are using the XPRS library from one Xpress distribution and the XPRS library from a previous Xpress distribution. You should remove all other Xpress distributions from your system library path environment variable.

245 ***Not enough memory to presolve matrix.***

The Optimizer required more memory to presolve the matrix. Please increase your virtual paging space or physical memory. If this is not possible try setting `PRESOLVE` to `0` before optimizing, so that the presolve procedure is not performed.

247 ***Directive on non-MIP entity not allowed: <col>.***

Column <col> used in directives file `.DIR` is not a MIP entity and its corresponding priority will be ignored. A variable is referred to as a 'MIP entity' if it is not continuous or if it is member of a set. Please refer to Appendix A for details about valid entities and set types.

249 ***Insufficient improvement found.***

Insufficient improvement was found between barrier iterations which has caused the barrier algorithm to terminate.

250 ***Too many numerical errors.***

Too many numerical errors have been encountered by the barrier algorithm and this has caused the barrier algorithm to terminate.

251 ***Out of memory.***

There is not enough memory for the barrier algorithm to continue.

259 ***Warning: The Q matrix may not be semi-definite.***

The Q matrix must be positive (negative) semi-definite for a minimization (maximization) problem in order for the problem to be convex. The barrier algorithm has encountered numerical problems which indicate that the problem is not convex.

261 ***<ent> already declared as a MIP entity - old declaration ignored.***

Entity <ent> has already been declared as MIP entity. The new declaration prevails and the old declaration prevails and the old declaration will be disregarded.

262 ***Unable to remove shift infeasibilities of <sumprimalinf>.***

Perturbations to the right hand side of the constraints which have been applied to enable problem to be solved cannot be removed. It may be due to round off errors in the input data or to the problem being badly scaled.

263 ***The problem has been presolved.***

The problem in memory is the presolved one. An operation has been attempted on the presolved problem. Please optimize again calling `XPRSmaxim` (`MAXIM`), `XPRSminim` (`MINIM`) with the `l` flag or tuning presolve off by setting `PRESOLVE` to `0`. If the operation does not need to be performed on an optimized problem just load the problem again.

264 ***Not enough spare matrix elements to remove all violations.***

The Optimizer could not add more cuts to the matrix because there is not enough space. Please increase `EXTRAELEMS` before input to improve performance.

266 ***Cannot read basis for presolved problem. Re-input matrix.***

The basis cannot be read because the problem in memory is the presolved one. Please reload the problem with `XPRSreadprob` (`READPROB`) and try to read the basis again.

268 ***Cannot perform operation on presolved matrix. Please postsolve or re-input matrix.***

The problem in memory is the presolved one. Please postsolve or reload the problem and try the operation again.

279 ***The Optimizer has not been initialized.***

The Optimizer could not be initialized successfully. Please initialize it before attempting any operation and try again.

287 ***Cannot read in directives after the problem has been presolved.***

Directives cannot be read if the problem in memory is the presolved one. Please reload the problem and read the directives file `.DIR` before optimizing. Alternatively, re-optimize using the `-l` flag or set `PRESOLVE` to 0 and try again.

293 ***This license file does not specify the permitted problem size. Contact your vendor to obtain a valid license.***

The license file is invalid as it doesn't specify the permitted problem size. Please contact your local sales office.

314 ***Invalid number.***

The input is not a number. Please check spelling and try again.

319 ***No Optimizer license found. Please contact your vendor to obtain a license.***

Your license does not authorize the direct use of the Optimizer. You probably have a license that authorizes other Xpress products, for example Mosel or BCL.

320 ***An internal error has occured. Please report to " SUPPORT_CONTACT_NAME " the circumstances under which this happened.***

An internal error has occured. Please report to " SUPPORT_CONTACT_NAME " the circumstances under which this happened.

324 ***Not enough extra matrix elements to complete elimination phase.***

Increase `PRESOLVEMAXGROW` before input to improve performance. The elimination phase performed by the presolve procedure created extra matrix elements. If the number of such elements is larger than allowed by the `PRESOLVEMAXGROW` parameter, the elimination phase will stop. Please increase `PRESOLVEMAXGROW` before loading the problem to improve performance.

352 ***Command not authorized in this version.***

There has been an attempt to use a command for which your Optimizer is not authorized. Please contact your local sales office to upgrade your authorization if you wish to use this command.

361 ***`QMATRIX` or `QUADOBJ` section must be after `COLUMN` section.***

Error in matrix file. Please make sure that the `QMATRIX` or `QUADOBJ` sections are after the `COLUMNS` section and try again.

362 ***Duplicate elements not allowed in `QUADOBJ` section.***

The coefficient of a column appears more than once in the `QUADOBJ` section. Please make sure all columns have only one coefficient in this section.

363 ***Quadratic matrix must be symmetric in `QMATRIX` section.***

Only symmetric matrices can be input in the `QMATRIX` section of the `.MAT` or `.MPS` file. Please correct and try again.

368 ***`QSECTION` second element in line ignored: <line>.***

The second element in line <line> will be ignored.

381 ***Bug in lifting of cover inequalities.***

Internal error. Please contact you local support office.

392 ***This version is not authorized to be called from BCL.***

This version of the Optimizer cannot be called from the subroutine library BCL. Please contact your local sales office to upgrade your authorization if you wish to run the Optimizer from BCL.

394 ***Fatal communications error.***

There has been a communication error between the master and the slave processes. Please check the network and try again.

395 ***This version is not authorized to be called from the Optimizer library.***

This version of the Optimizer cannot be called from the Optimizer library. Please contact your local sales office to upgrade your authorization if you wish to run the Optimizer using the libraries.

401 ***Invalid row type passed to <function>.***

Elements <num> of your array has invalid row type <type>. There has been an error in one of the arguments of function <function>. The row type corresponding to element <num> of the array is invalid. Please refer to the section corresponding to function <function> in 8 for further information about the row types that can be used.

402 ***Invalid row number passed to <function> (at index <idx>). Row number <num> is invalid***

There has been an error in one of the arguments of function <function>. The row number corresponding to element <idx> of the array is invalid. Please make sure that the row numbers are at least `0` and less than the total number of rows in the problem.

403 ***Invalid MIP entity passed to <function>.***

Element <num> of your array has invalid entity type <type>. There has been an error in one of the arguments of function <function>. The column type <type> corresponding to element <num> of the array is invalid for a MIP entity.

404 ***Invalid set type passed to <function>.***

Element <num> of your array has invalid set type <type>. There has been an error in one of the arguments of function <function>. The set type <type> corresponding to element <num> of the array is invalid for a set entity.

405 ***Invalid column number passed to <function>. Column number <num> is invalid***

There has been an error in one of the arguments of function <function>. Please make sure that the column numbers are at least `0` and less than the total number of columns in the problem. If the function being called is `XPRSgetobj` or `XPRSchgobj` a column number of `-1` is valid and refers to the constant in the objective function.

406 ***Invalid row range passed to <function>.***

Limit <lim> is out of range. There has been an error in one of the arguments of function <function>. The row numbers lie between 0 and the total number of rows of the problem. Limit <lim> is outside this range and therefore is not valid.

407 ***Invalid column range passed to <function>.***

Limit <lim> is out of range. There has been an error in one of the arguments of function <function>. The column numbers lie between 0 and the total number of columns of the problem. Limit <lim> is outside this range and therefore is not valid.

409 ***Invalid directive passed to <function>.***

Element <num> of your array has invalid directive <type>. There has been an error in one of the arguments of function <function>. The directive type <type> corresponding to element <num> of the array is invalid. Please refer to the Reference Manual for a list of valid directive types.

410 ***Invalid row basis type passed to <function>.***

Element <num> of your array has invalid row basis type <type>. There has been an error in one of the arguments of function <function>. The row basis type corresponding to element <num> of the array is invalid.

411 ***Invalid column basis type passed to <function>.***

Element <num> of your array has invalid column basis type <type>. There has been an error in one of the arguments of function <function>. The column basis type corresponding to element <num> of the array is invalid.

412 ***Invalid parameter number passed to <function>.***

Parameter number <num> is out of range. LP or MIP parameters and controls can be used in functions by passing the parameter or control name as the first argument or by passing an associated number. In this case number <num> is an invalid argument for function <function> because it does not correspond to an existing parameter or control. If you are passing a number as the first argument, please substitute it with the name of the parameter or control whose value you wish to set or get. If you are already passing the parameter or control name, please check 8 to make sure that is valid for function <function>.

413 ***Not enough spare rows in <function>.***

Increase `EXTRAROWS` before input. There are not enough spare rows to complete function <function> successfully. Please increase `EXTRAROWS` before `XPRSreadprob` (`READPROB`) and try again.

414 ***Not enough spare columns in <function>.***

Increase `EXTRACOLS` before input. There are not enough spare columns to complete function <function> successfully. Please increase `EXTRACOLS` before `XPRSreadprob` (`READPROB`) and try again.

415 ***Not enough spare matrix elements in <function>.***

Increase `EXTRAELEMS` before input. There are not enough spare matrix elements to complete function <function> successfully. Please increase `EXTRAELEMS` before `XPRSreadprob` (`READPROB`) and try again.

416 ***Invalid bound type passed to <function>.***

Element <elem> of your array has invalid bound type <type>. There has been an error in one of the arguments of function <function>. The bound type <type> of element number <num> of the array is invalid.

417 ***Invalid complement flag passed to <function>. Element <elem> of your array has invalid complement flag <flag>.***

Element <elem> of your array has an invalid complement flag <flag>. There has been an error in one of the arguments of function <function>. The complement flag corresponding to indicator constraint <num> of the array is invalid.

418 ***Invalid cut number passed to <function>.***

Element <num1> of your array has invalid cut number <num2>. Element number <num1> of your array contains a cut which is not stored in the cut pool. Please check that <num2> is a valid cut number.

419 ***Not enough space to store cuts in <function>.***

There is not enough space to complete function <function> successfully.

420 ***Too many saved matrices in savmat***

Version 12 compatibility interface only. There is a hard limit of at most 64 matrices that can be saved with `savmat` or `cpymat`

421 ***Matrix no. <mat> has not been saved. Cannot restore in resmat***

Version 12 compatibility interface only. No matrix with number <mat> has been saved with `savmat` or `cpymat`.

422 ***Solution is not available.***

There is no solution available. This could be because the problem in memory has been changed or optimization has not been performed. Please optimize and try again.

423 ***Duplicate rows/columns passed to <function>.***

Element <elem> of your array has duplicate row/col number <num>. There has been an error in one of the arguments of function <function>. The element number <elem> of the argument array is a row or column whose sequence number <num> is repeated.

424 ***Not enough space to store cuts in <function>.***

There is not enough space to complete function <function> successfully.

425 ***Column already basic.***

The column cannot be pivoted into the basis since it is already basic. Please make sure the variable is non-basic before pivoting it into the basis.

426 ***Column not eligible to leave basis.***

The column cannot be chosen to leave the basis since it is already non-basic. Please make sure the variable is basic before forcing it to leave the basis.

427 ***Invalid column type passed to <function>.***

Element <num> of your array has invalid column type <type>. There has been an error in one of the arguments of function <function>. The column type <type> corresponding to element <num> of the array is invalid.

429 ***No basis is available.***

No basis is available.

430 ***Column types cannot be changed during the tree search.***

The Optimizer does not allow changes to the column type while the tree search is in progress. Please call this function before starting the tree search or after the tree search has been completed. You can call `XPRSmaxim` (`MAXIM`) or `XPRSminim` (`MINIM`) with the `l` flag if you do not want to start the tree search automatically after finding the LP solution of a problem with MIP entities.

434 ***Invalid name passed to `XPRSgetindex`.***

A name has been passed to `XPRSgetindex` which is not the name of a row or column in the matrix.

436 ***Cannot trace infeasibilities when integer presolve is turned on.***

Try `XPRSmaxim` (`XPRSmaxim`) / `XPRSminim` (`MINIM`) with the `l` flag. Integer presolve can set upper or lower bounds imposed by the column type as well as those created by the interaction of the problem constraints. The infeasibility tracing facility can only explain infeasibilities due to problem constraints.

459 ***Not enough memory for branch and bound tree***
Not enough resources for branch and bound tree (<type>)
Failure locking branch and bound tree (probably out of memory)
Failure in handling of branch and bound tree (<type>)

Functions to signal that an unexpected error happened during the management of the branch-and-bound tree for storing information from a MIP solve. The string <type> will provide more information about the particular failure. These errors are typical of running out of memory.

473 ***Row classification not available.***

474 ***Column passed to <routine> has inconsistent bounds. See column <index> of <count>.***

The bounds are inconsistent for column <index> of the <count> columns passed into routine <routine>.

475 ***Inconsistent bounds [<lb>,<ub>] for column <column name> in call to <routine>.***

The lower bound <lb> is greater than the upper bound <ub> in the bound pair given for column <column name> passed into routine <routine>.

476 ***Unable to round bounds [<lb>,<ub>] for integral column <column name> in call to <routine>.***

Either the lower bound <lb> is greater than the upper bound <ub> in the bound pair given for the integer column <column name> passed into routine <routine> or the interval defined by <lb> and <ub> does not contain an integer value.

501 ***Error at <line> Empty file.***

Read aborted. The Optimizer cannot read the problem because the file is empty.

502 ***Warning: ’min’ or ’max’ not found at <line.col>. No objective assumed.***

An objective function specifier has not been found at column <col>, line <line> of the LP file. If you wish to specify an objective function please make sure that ’max’, ’maximize’, ’maximum’, ’min’, ’minimize’ or ’minimum’ appear.

503 ***Objective not correctly formed at <line.col>. Aborting.***

The Optimizer has aborted the reading of the problem because the objective specified at line <line> of the LP file is incorrect.

504 ***No keyword or empty problem at <line.col>.***

There is an error in column <col> at line <line> of the LP file. Neither ’Subject to’, ’subject to:’, ’subject to’, ’such that’ ’s.t.’, or ’st’ can be found. Please correct and try again.

505 ***A keyword was expected at <line.col>.***

A keyword was expected in column <col> at line <line> pf the LP file. Please correct and try again.

506 ***The constraint at <line.col> has no term.***

A variable name is expected at line <line> column <col>: either an invalid character (like ’+’ or a digit) was encountered or the identifier provided is unknown (new variable names are declared in constraint section only).

507 ***RHS at <line.col> is not a constant number.***

Line <line> of the LP file will be ignored since the right hand side is not a constant.

508 ***The constraint at <line> has no term.***

The LP file contains a constraint with no terms.

509 ***The type of the constraint at <line.col> has not been specified.***

The constraint defined in column <col> at line <line> of the LP file is not a constant and will be ignored.

510 ***Upper bound at <line.col> is not a numeric constant.***

The upper bound declared in column <col> at line <line> of the LP file is not a constant and will be ignored.

511 ***Bound at <line.col> is not a numeric constant.***

The bound declared in column <col> at line <line> of the LP file is not a constant and will be ignored.

512 ***Unknown word starting with an ’f’ at <line.col>. Treated as ’free’.***

A word starting with an ’f’ and not known to the Optimizer has been found in column <col> at line <line> of the LP file. The word will be read into the Optimizer as ’free’.

513 ***Wrong bound statement at <line.col>.***

The bound statement in column <col> at line <line> is invalid and will be ignored.

514 ***Lower bound at <line.col> is not a numeric constant. Treated as -inf.***

The lower bound declared in column <col> at line <line> of the LP file is not a constant. It will be translated into the Optimizer as the lowest possible bound.

515 ***Sign ’<’ expected at <line.col>.***

A character other than the expected sign ’<’ has been found in column <col> at line <line> of the LP file. This line will be ignored.

516 ***Problem has not been loaded.***

The problem could not be loaded into the Optimizer. Please check the other error messages appearing with this message for more information.

517 ***Error when loading names at <line.col>.***

The name of the rows, columns, objectives, special ordered sets, piecewise linear constraints or general constraints could not be loaded into the Optimizer. Please check the other error messages appearing with this message for more information.

518 ***Unexpected character at <line.col>.***

The LP file contains a character that was not expected. Please check that the syntax of the file is correct.

519 ***Not enough memory at <line.col>.***

The information in column <col> at line <line> of the LP file cannot be read because all the allocated memory has already been used. Please increase your virtual page space or physical memory and try again.

520 ***Unexpected EOF at <line.col>.***

An unexpected EOF marker character has been found at line <line> of the LP file and the loading of the problem into the Optimizer has been aborted. Please correct and try again.

521 ***Number expected for exponent at <line.col>.***

The entry in column <col> at line <line> of the LP file is not a properly expressed real number and will be ignored.

522 ***Line <line> too long (length>255).***

Line <line> of the LP file is too long and the loading of the problem into the Optimizer has been aborted. Please check that the length of the lines is less than 255 and try again.

523 ***The Optimizer cannot reach line <line.col>.***

The reading of the LP file has failed due to an internal problem. Please contact your local support office.

524 ***Constraints could not be read into the Optimizer. Error found at <line.col>.***

The reading of the LP constraints has failed due to an internal problem. Please contact your local support office.

525 ***Bounds could not be set into the Optimizer. Error found at <line.col>.***

The setting of the LP bounds has failed due to an internal problem. Please contact your local support office.

526 ***LP problem could not be loaded into the Optimizer. Error found at <line.col>.***

The reading of the LP file has failed due to an internal problem. Please contact your local support office.

527 ***Copying of rows unsuccessful.***

The copying of the LP rows has failed due to an internal problem. Please contact your local support office.

528 ***Copying of columns unsuccessful.***

The copying of the LP columns has failed due to an internal problem. Please contact your local support office.

529 ***Redefinition of constraint at <line.col>.***

A constraint is redefined in column <col> at line <line> of the LP file. This repeated definition is ignored.

530 ***Name too long. Truncating it.***

The LP file contains an identifier longer than 64 characters: it will be truncated to respect the maximum size.

531 ***Sign '>' expected here <line>.***

A greater than sign was expected in the LP file.

532 ***Quadratic term expected here <pos>***

The LP file reader expected to read a quadratic term at position <pos>: a variable name and "2' or the product of two variables. Please check the quadratic part of the objective in the LP file.

533 ***Wrong exponent value. Treated as 2 <pos>***

The LP file reader encountered an exponent different than 2 at position <pos>. Such exponents are automatically replaced by 2.

535 ***A constraint name was expected here.***

The LP file reader expected to read a row name in the ranges section.

536 ***Wrong range statement at <pos>.***

The LP file reader expected to read a range type in the ranges section.

538 ***Quadratic objective functions are not supported in multi-objective problems.***

The problem being read has more than one objective, and at least one objective contains quadratic terms. This is not currently supported by the Optimizer. This error can be avoided by reformulating the problem with a linear objective function and expressing the quadratic terms via transfer rows.

539 ***Invalid indicator constraint condition at <line.col>***

The condition part in column <col> of the indicator constraint at line <line> is invalid.

545 ***A variable name was found but ignored at <pos> due to not appearing before.***

The LP file reader read a variable in bounds or integer type sections that does not appear in the matrix.

548 ***Repeated bound at <pos>. Previous bound ignored.***

The LP file reader read more than one definition of a lower bound or an upper bound in the *bounds* section.

552 ***'S1|2:' expected here. Skipping <pos>***

Unknown set type read while reading the LP file at position <pos>. Please use set type 'S1' or 'S2'.

553 ***This set has no member. Ignoring it <pos>***

An empty set encountered while reading the LP file at position <pos>. The set has been ignored.

554 ***Weight expected here. Skipping <pos>***

A missing weight encountered while reading sets in the LP file at position <pos>. Please check definitions of the sets in the file.

555 ***Cannot presolve cut with `PRESOLVEOPS` bits 0, 5 or 8 set or bit 11 cleared.***

Can not presolve cut with `PRESOLVEOPS` bits 0, 5 or 8 set or bit 11 cleared.
No cuts can be presolved if the following presolve options are turned on:
bit 0: singleton column removal,
bit 5: duplicate column removal,
bit 8: variable eliminations
or if the option
bit 11: No advanced IP reductions is turned off. Please check the presolve settings.

557 ***Integer solution is not available***

Failed to retrieve an integer solution because no integer solution has been identified yet.

558 ***Column <col> duplicated in basis file - new entry ignored.***

Column <col> is defined in the basis file more than once. Any repeated definitions are ignored.

559 ***The old feature <feature> is no longer supported***

The feature <feature> is no longer supported and has been removed. Please contact Xpress support for help about replacement functionality.

602 ***Values must be specified for all columns when column indices are not provided.***

In a call to `XPRSaddmipsol` the column index array is optional. When this argument is omitted (given as `NULL`), the length of the solution value array must match `ORIGINALCOLS`.

603 ***Duplicate value for <type> <value> at index <index> passed to <function>***

In a call to <function> a duplicate entry <value> is passed at index <index> in the array of type <type>.

604 ***String passed as parameter is too long***

The file name passed to `XPRSsetlogfile` can be at most 200 characters long.

606 ***Failed to parse list of diving heuristic strategies at position <pos>***

Invalid diving heuristic strategy number provided in position <pos> of the string controls `HEURDIVEUSE` or `HEURDIVETEST`. Please check control `HEURDIVESTRATEGY` for valid strategy numbers.

706 ***Not enough memory to add sets.***

Insufficient memory while allocating memory for the new sets. Please free up some memory, and try again.

707 ***Function can not be called during the tree search***

The function being called cannot be used during the tree search. Please call the function before starting the tree search.

708 ***Invalid input passed to <function>***
Must specify mstart or mnel when creating matrix with columns

No column information is available when calling function <function>. If no columns were meant to be passed to the function, then please set the column number to zero. Note, that `mstart` and `mnel` should be set up for empty columns as well.

710 ***`MIPTOL` (`FEASTOL`) <val1> must not be less than `FEASTOL` (`MATRIXTOL`) <val2>***

The integer tolerance `MIPTOL`, feasibility tolerance `FEASTOL`, and zero tolerance `MATRIXTOL` should be ordered as `MIPTOL` >= `FEASTOL` >= `MATRIXTOL`. Please increase `MIPTOL` or decrease `MATRIXTOL` before setting `FEASTOL`.

711 ***MIPTOL(FEASTOL) <val1> must not be less than FEASTOL(MATRIXTOL) <val2>. Adjusting MIPTOL(FEASTOL)***

The integer tolerance `MIPTOL`, feasibility tolerance `FEASTOL`, and zero tolerance `MATRIXTOL` should be ordered as `MIPTOL` >= `FEASTOL` >= `MATRIXTOL`. The value of `MIPTOL` or `FEASTOL` has been increased to (`val2`) for the search.

712 ***Function not permitted when problem is presolved: <func>***

The problem is currently in a presolved state and the function <func> can only be called when the problem is in its original state. `XPRSpostsolve` can be called to return the problem to its original state.

713 ***<row/column> index out of bounds calling <function>. <index1> is '<' or '>' <bound>***

An index is out of its bounds when calling function <function>. Please check the indices.

715 ***Invalid objective sense passed to <function>. Must be XPRS_OBJ_MINIMIZE or XPRS_OBJ_MAXIMIZE.***

Invalid objective sense was passed to function <function>. Please use either `XPRS_OBJ_MINIMIZE` or `XPRS_OBJ_MAXIMIZE`.

716 ***Invalid names type passed to XPRSgetnamelist.***
Type code <num> is unrecognized.

An invalid name type was passed to `XPRSgetnamelist`.

717 ***Generic error.***

Used to promote license manager errors.

721 ***No IIS has been identified yet***

No irreducible infeasible set (IIS) has been found yet. Before running the function, please use `IIS -f`, `IIS -n` or `IIS -a` to identify an IIS.

722 ***IIS number <num> is not yet identified***

Irreducible infeasible set (IIS) with number <num> is not available. The number <num> stands for the ordinal number of the IIS. The value of <num> should not be larger than `NUMIIS`.

723 ***Unable to create an IIS sub-problem***

The irreducible infeasible set (IIS) procedure is unable to create the IIS approximation. Please check that there is enough free memory.

724 ***Error while optimizing the IIS sub-problem***

An error occurred while minimizing an irreducible infeasible set (IIS) sub-problem. Please check the return code set by the Optimizer.

725 ***Problems with variables for which shift infeasibilities cannot be removed are considered infeasible in the IIS***

The irreducible infeasible set (IIS) sub-problem being solved by the IIS procedure is on the boundary of being feasible or infeasible. For problems that are only very slightly infeasible, the Optimizer applies a technique called infeasibility shifting to produce a solution. Such solutions are considered feasible, although if solved as a separate problem, a warning message is given. For consistency reasons however, in the case of the IIS procedure such problems are treated as being infeasible.

726 ***This function is not valid for the IIS approximation. Please specify an IIS with count number > 0***

Irreducible infeasible set (IIS) number 0 (the ordinal number of the IIS) refers to the IIS approximation, but the functionality called is not available for the IIS approximation. Please use an IIS number between 1 and `NUMIIS`.

727 ***Bound conflict on column <col>; IIS will not continue***

There is a bound conflict on column <col>. Please check the bounds on the column, and remove any conflicts before running the irreducible infeasible set (IIS) procedure again (bound conflicts are trivial IISs by themselves).

728 ***Unknown file type specification <type>***

Unknown file type was passed to the irreducible infeasible set (IIS) sub-problem writer. Please refer to `XPRSiiswrite` for the valid file types.

729 ***Writing the IIS failed***

Failed to write the irreducible infeasible set (IIS) sub-problem or the comma separated file (.csv) containing the IIS information to disk. Please check access permissions.

730 ***Failed to retrieve data for IIS <num>***

The irreducible infeasible set (IIS) procedure failed to retrieve the internal description for IIS number <num>. This may be an internal error, please contact your local support office.

731 ***IIS stability error: reduced or modified problem appears feasible***

Some problems are on the boundary of being feasible or infeasible. For such problems, it may happen that the irreducible infeasible set (IIS) working problem becomes feasible unexpectedly. If the problem persists, please contact your local support office.

732 ***Unknown parameter or wrong parameter combination***

The wrong parameter or parameter combination was used when calling the irreducible infeasible set (IIS) console command. Please refer to the IIS command documentation for possible combinations.

733 ***Filename parameter missing***

No filename is provided for the `IIS -w` or `IIS -e` console command. Please provide a file name that should contain the irreducible infeasible set (IIS) information.

734 ***Problem data relevant to IISs is changed***

This failure is due to the problem being changed between iterative calls to IIS functions. Please start the IIS analysis from the beginning.

736 ***Initial infeasible subproblem is not available. Run `IIS -f` to set it up***

The initial infeasible subproblem requested is not available. Please use the `IIS -f` function to generate it.

738 ***The approximation may be inaccurate. Please use `IIS` or `IIS -n` instead.***

The irreducible infeasible set (IIS) procedure was run with the option of generating the approximation of an IIS only. However, ambiguous duals or reduces costs are present in the initial infeasible subproblem. This message is always preceded by warning 737. Please continue with generating IISs to resolve the ambiguities.

739 ***Bound conflict on column <col>; Repairinfeas will not continue***

There is a bound conflict on column <col>. Please check the bounds on the column, and remove any conflicts before running the `XPRSrepairinfeas` procedure again (bound conflicts are trivial causes of infeasibility).

740 ***Unable to create relaxed problem***

The Optimizer is unable to create the relaxed problem. The relaxed problem may require significantly more memory than the base problem if many of the preferences are set to a positive value. Please check that there is enough free memory.

741 ***Relaxed problem is infeasible. Please increase freedom by introducing new nonzero preferences***

The relaxed problem remains infeasible. Zero preference values indicate constraints (or bounds) that will not be relaxed. Try introducing new nonzero preferences to allow the problem to become feasible.

742 ***Repairinfeas stability error: relaxed problem is infeasible. You may want to increase the value of delta***

The relaxed problem is reported to be infeasible by the Optimizer in the second phase of the repairinfeas procedure. Try increasing the value of the parameter delta to improve stability.

743 ***Optimization aborted, repairinfeas unfinished***

The optimization was aborted by CTRL-C or by hitting a time limit. The relaxed solution is not available.

744 ***Optimization aborted, MIP solution may be sub-optimal***

The MIP optimization was aborted by either CTRL-C or by hitting a time limit. The relaxed solution may not be optimal.

745 ***Optimization of the relaxed problem is sub-optimal***

The relaxed solution may not be optimal due to early termination.

746 ***All preferences are zero, repairinfeas will not continue.***
Use options `-a -b -r -lbp -ubp -lrp` or `-grp` to add nonzero preferences

Zero preference values indicate constraints (or bounds) that will not be relaxed. In case when all preferences are zero, the problem cannot be relaxed at all. Try introducing nonzero preferences and run `XPRSrepairinfeas` again.

748 ***Negative preference given for a <sense> bound on <row/column> <name>***

A negative preference value is set for constraint or bound <name>. Preference values should be non-negative. The preferences describe the modeler's willingness to relax a given constraint or bound, with zero preferences interpreted as the corresponding constraints or bounds not being allowed to be relaxed. Please provide a zero preference if the constraint or bound is not meant to be relaxed. Also note, that very small preferences lead to very large penalty values, and thus may increase the numerical difficulty of the problem.

749 ***Relaxed problem is infeasible due to `cutoff`***

A user defined `cutoff` value makes the relaxed problem infeasible. Please check the cutoff value `CURRMIPCUTOFF`.

750 ***Empty matrix file : <name>***

The MPS file <name> is empty. Please check the name of the file and the file itself.

751 ***Invalid column marker type found : <text>***

The marker type <text> is not supported by the MPS reader. Please refer to the Appendix A.2 for supported marker types.

752 ***Invalid floating point value : <text>***

The reader is unable to interpret the string <text> as a numerical value.

753 ***<num> lines ignored***

The MPS reader has ignored <num> number of lines. This may happen for example if an unidentified section was found (in which case warning 785 is also invoked).

754 ***Insufficient memory***

Insufficient memory was available while reading in an MPS file.

755 ***Column name is missing***

A column name field was expected while reading an MPS file. Please add a column name to the row. If the `MPSFORMAT` control is set to 0 (fixed format) then please check that the name field contains a column name, and is positioned correctly.

756 ***Row name is missing in section `OBJNAME`***

No row name is provided in the `OBJNAME` section. If no user defined objective name is provided, the reader uses the first neutral row (if any) as the objective row. However, to avoid ambiguity, if no user defined objective row was meant to be supplied, then please exclude the `OBJNAME` section from the MPS file.

757 ***Missing objective sense in section `OBJSENSE`***

No objective sense is provided in section `OBJSENSE`. If no user defined objective sense is provided, the reader sets the objective sense to minimization by default. However, to avoid ambiguity, if no user defined objective sense was meant to be supplied, then please exclude the `OBJSENSE` section from the MPS file.

758 ***No `SETS` and `SOS` sections are allowed in the same file***

The Optimizer expects special order sets to be defined in the `SETS` section. However, for compatibility considerations, the Optimizer can also interpret the `SOS` section. The two formats differ only in syntax, and feature the same expressive power. Both a `SETS` and a `SOS` section are not expected to be present in the same matrix file.

759 ***File not in fixed format : <file>***

The Optimizer control `MPSFORMAT` was set to 0 to indicate that the MPS file <file> being read is in fixed format, but it violates the MPS field specifications.

760 ***Objective row <row> defined in section `OBJNAME` or in `MPSOBJNAME` was not found***

The user supplied objective row <row> is not found in the MPS file. If the MPS file contains an `OBJNAME` section please check the row name provided, otherwise please check the value of the control `MPSOBJNAME`.

761 ***Problem name is not provided***

The `NAME` section is present in the MPS file, but contains no problem name (not even blanks), and the `MPSFORMAT` control is set to 0 (fixed format) preventing the reader to look for the problem name in the next line. Please make sure that a problem name is present, or if it's positioned in the next line (in which case the first column in the line should be a whitespace) then please set `MPSFORMAT` to 1 (free format) or -1 (autodetect format).

762 ***Missing problem name in section `NAME`***

Unexpected end of file while looking for the problem name in section `NAME`. The file is likely to be corrupted. Please check the file.

763 ***Ignoring range value for free row : <row>***

A range value is defined for free row <row>. Range values have no effect on free rows. Please make sure that the type of the row in the `ROWS` section and the row name in the `RANGE` section are both correct.

764 ***<sec> section is not yet supported in an MPS file, skipping section***

The section <sec> is not allowed in an MPS file. Sections like "`SOLUTION`" and "`BASIS`" must appear in separate "`.slx`" and "`.bas`" files.

765 ***Ignoring repeated specification for column : <col>***

Column <col> is defined more than once in the MPS file. Any repeated definitions are ignored. Please make sure to use unique column names. If the column names are unique, then please make sure that the `COLUMNS` section is organized in a contiguous order.

766 ***Ignoring repeated coefficients for row <row> found in `RANGE` <range>***

The range value for row <row> in range vector <range> in the `RANGE` section is defined more than once. Any repeated definitions are ignored. Please make sure that the row names in the `RANGE` section are correct.

767 ***Ignoring repeated coefficients for row <row> found in `RHS` <rhs>***

The value for row <row> in right hand side vector <rhs> is defined more than once in the `RHS` section. Any repeated definitions are ignored. Please make sure that the row names in the `RHS` section are correct.

768 ***Ignoring repeated specification for row : lt;rowgt;***

Row <row> is defined more than once in the MPS file. Any repeated definitions are ignored. Please make sure to use unique row names.

769 ***Unexpected bound type : <type>***

The `BOUNDS` section contains the unknown bound type <type>. If the `MPSFORMAT` control is set to 0 (fixed format) then please make sure that the type of the bound is correct and positioned properly.

770 ***Missing prerequisite section <sec1> for section <sec2>***

Section <sec2> must be defined before section <sec1> in the MPS file being read. Please check the order of the sections.

771 ***Unable to open file : <file>***

Please make sure that file <file> exists and is not locked.

772 ***Unexpected column type : <type> : <column>***

The `COLUMNS` section contains the unknown column type <type>. If the `MPSFORMAT` control is set to 0 (fixed format) then please make sure that the type of the column is correct and positioned properly.

773 ***Unexpected number of fields in section : <sec>***

Unexpected number of fields was read by the reader in section <sec>. Please check the format of the line. If the `MPSFORMAT` control is set to 0 (fixed format) then please make sure that the fields are positioned correctly. This error is often caused by names containing spaces in free format, or by name containing spaces in fixed format but positioned incorrectly.

774 ***Unexpected row type : <type>***

The `ROWS` section contains the unknown row type <type>. If the `MPSFORMAT` control is set to 0 (fixed format) then please make sure that the type of the row is correct and positioned properly.

775 ***Unexpected set type : <type>***

The `SETS` or `SOS` section contains the unknown set type <type>. If the `MPSFORMAT` control is set to 0 (fixed format) then please make sure that the type of the row is correct and positioned properly.

776 ***Ignoring unknown column name <col> found in `BOUNDS`***

Column <col> found in the `BOUNDS` section is not defined in the `COLUMNS` section. Please check the name of the column.

777 ***Ignoring quadratic coefficient for unknown column : <col>***

Column <col> found in the `QUADOBJ` section is not defined in the `COLUMNS` section. Please check the name of the column.

778 ***Ignoring unknown column name <col> found in set <set>***

Column <col> found in the definition of set <set> in the `SETS` or `SOS` section is not defined in the `COLUMNS` section. Please check the name of the column.

779 ***Wrong objective sense: <sense>***

The reader is unable to interpret the string <sense> in the `OBJSENSE` section as a valid objective sense. The objective sense should be either `MAXIMIZE` or `MINIMIZE`. The reader accepts sub-strings of these if they uniquely define the objective sense and are at least 3 characters long. Note that if no `OBJSENSE` section is present, the sense of the objective is set to minimization by default. Please provide a valid objective sense.

780 ***Ignoring unknown row name <row> found in column <column>***

Row <row> found in the column <column> in the `COLUMNS` section is not defined in the `ROWS` section. Please check the name of the row.

781 ***Ignoring unknown row name <row> found in `RANGE`***

Row <row> found in the `RANGE` section is not defined in the `ROWS` section. Please check the name of the row.

782 ***Ignoring unknown row name <row> found in `RHS`***

Row <row> found in the RHS section is not defined in the `ROWS` section. Please check the name of the row.

783 ***Expecting numerical value***

A numerical value field was expected while reading an MPS file. Please add the missing numerical entry. If the `MPSFORMAT` control is set to 0 (fixed format) then please check that the value field contains a numerical value and is positioned correctly.

784 ***Null char in text file***

A null char ('\0') encountered in the MPS file. An MPS file is designed to be a text file and a null char indicates possible errors. Null chars are treated as spaces ' ' by the reader, but please check the origin of the null char.

785 ***Unrecognized section <sec> skipped***

The section <sec> is not recognized as an MPS section. Please check the section identifier string in the MPS file. In such cases, the reader stops reading to avoid unexpected results after reading.

787 ***Empty set: <set>***

No set members are defined for set <set> in the MPS file. Please check if the set is empty by intention.

788 ***Repeated definition of section <sec> ignored***

Section <sec> is defined more than once in the MPS file. Any repeated definitions are ignored. Many sections may include various versions of the described part if the problem (like different `RHS` values, `BOUNDS` or `RANGES`), but please include those in the same section.

790 ***Wrong section in the basis file: <section>***

Unrecognized section <section> found in the basis file. Please check the format of the file.

791 ***`ENDATA` is missing. File is possibly corrupted***

The `ENDATA` section is missing from the end of the file. This possible indicates that part of the file is missing. Please check the file.

792 ***Ignoring `BS` field***

BS fields are not supported by the Optimizer, and are ignored. Basis files containing `BS` fields may be created by external software. Please convert `BS` fields to either `XU` or `XL` fields.

793 ***Superbasic variable outside bounds. Value moved to closest bound***

A superbasic variable in the basis file are outside its bounds. The value of the variable has been modified to satisfy its bounds. Please check that the value in the basis file is correct. In case the variable should be set to the value given by the basis file, please modify the bounds on the variable.

794 ***Value of fixed binary column <col> changed to <val>***

The lower and upper bound for binary variable <col> was to <val>. Binaries may only be fixed at level 0 or 1.

795 ***Xpress/Mosel extensions: number of opening and closing brackets mismatch***

The LP file appears to be created by Mosel, using the Xpress MPS extensions to include variable names with whitespaces, however the file seems to be broken due to a mismatch in opening and closing brackets.

796 ***Char <c> is not supported in a name by file format. It may not be possible to read such files back correctly. Please set FORCEOUTPUT to 1 to write the file anyway, or use scrambled names.***

Certain names in the problem object may be incompatible with different file formats (like names containing spaces for LP files). If the Optimizer might be unable to read back a problem because of non-standard names, it will give an error message and won't create the file. However, you may force output using control `FORCEOUTPUT` or change the names by using scrambled names (`-s` option for `XPRSwriteprob`).

797 ***Wrong section in the solution file: <sec>***

Section <sec> is not supported in `.slx` MPS solution files.

798 ***Empty <type> file : <file>***

File <file> of type <type> is empty.

799 ***Ignoring quadratic coefficients for unknown row name <row>***

No row with name <row> was defined in the `ROWS` sections. All rows having a `QCMATRIX` section must be defined as a row with type '`L`' or '`G`' in the `ROWS` section.

835 ***Given solution column count does not match given problem***

The given solution contains a different column count compared to the loaded problem.

843 ***Delayed row (lazy constraint) <row> is not allowed to be of type 'N'. Row ignored***

Delayed rows cannot be neutral. Please define all neutral rows as ordinary ones in the `ROWS` section.

847 ***Model cut (user cut) <row> is not allowed to be of type 'N'. Row ignored***

Model cuts cannot be neutral. Please define all neutral rows as ordinary ones in the `ROWS` section.

862 ***Quadratic constraint rows must be of type 'L' or 'G'. Wrong row type for row <row>***

All quadratic rows must be of type '`L`' or '`G`' in the `ROWS` section of the MPS file (and the corresponding quadratic matrix be positive semi-definite).

863 ***The current version of the Optimizer does not yet support MIQCQP problems***

The current version of the Optimizer does not yet support mixed integer quadratically constrained problems.

864 ***Quadratic constraint rows must be of type 'L' or 'G'. Wrong row type for row <row>***

A library function was trying to define (or change to) a row with type '`L`' having quadratic coefficients. All quadratic rows are required to be of type '`L`' (and the corresponding quadratic matrix be positive semi-definite).

865 ***Row <row> is already quadratic***

Cannot add quadratic constraint matrices together. To change an already existing matrix, either use the `XPRSchgqrowcoeff` library function, or delete the old matrix first.

866 ***The divider of the quadratic objective at <pos> must be 2 or omitted***

The LP file format expects, tough may be omitted, an "`/2`" after the each quadratic objective term defined between square brackets. No other divider is accepted. The role of the "`/2`" is to notify the user of the implied division in the quadratic objective (that does not apply to quadratic constraints).

867 ***Not enough memory for tree search***

There is not enough memory for one of the nodes in the tree search.

884 ***Fatal user error detected in callback***

An error occurred during a user callback.

898 ***Cannot define range for quadratic rows. Range for row <row> ignored***

Quadratic constraints are required to be convex, and thus it is not allowed to set a range on quadratic rows. Each quadratic row should have a type of '`L`' or '`G`'.

899 ***The quadratic objective is not convex. Please check your model or use nlpoptimize.***

The quadratic objective is not convex. Please check that the proper sense of optimization (minimization or maximization) is used or call `XPRSnlpoptimize` if the problem is supposed to be non-convex.

900 ***The quadratic part of row <row> defines a non-convex region. Please check your model or use nlpoptimize.***

The quadratic in <row> is not convex. Please check that the proper sense of constraint is defined (less or equal or greater or equal constraint) or call `XPRSnlpoptimize` if the problem is supposed to be non-convex.

901 ***Duplicated QCMATRIX section for row <row> ignored.***

The MPS file may contain one Q matrix for each row. In case of duplicates, only the first is loaded into the matrix

902 ***Calling function <func> is not supported from the current context.***

This XPRS function cannot be called from this callback.

903 ***Row <row> with right hand side value larger than infinity ignored.***

The matrix file being read contains a right hand side that is larger than the predefined infinity constant XPRS_PLUSINFINITY. Row is made neutral.

904 ***Function is not allowed outside optnode callback.***

The used function of the branching manager is not allowed to call outside optnode.

905 ***Bad index passed to function.***

The index passed to function is not in range of the attribute.

906 ***MIP entity cannot be branched further.***

The selected MIP entity is fixed and cannot be branched further.

907 ***Column is continuous and cannot be branched.***

The given column is of continuous type. The used function does not support branching on continuous columns.

909 ***Limit exceeded.***

The limit of a certain object is exceeded.

910 ***Empty branch or branching object.***

The given branch or branching object is empty.

911 ***Invalid information provided for branching object.***

The given branching object contains invalid information.

912 ***Branching object(s) cannot be changed/used at this time.***

The branching object is not fix yet. Hence, it cannot be changed/used.

913 ***Required data missing in function call for branching object.***

Data is missing in function call for branching object.

914 ***Unexpected error triggered for branching object.***

Unexpected error happened on a branching object.

915 ***Branching object (ID=<id>) rejected because it is empty or contains empty branches.***

Improper branching object.

918 ***Module error.***

Model can not be modified.

919 ***Column must be of type semi-continuous, semi-integer or partial integer to change its global bound.***

The global bound can be modified only for semi-continuous, semi-integer or partial integer column.

920 ***Semi-continuous lower bound for column <column> must be non-negative.***

Only non-negative lower bounds can be specified for semi-continuos columns.

921 ***Partial integer limit for column <column> is outside the allowed range of 0 to*** $2^28 - 1$***.***

The give limit for the column is out of the allowed range.

932 ***Unknown column name <column> found in piecewise linear term.***

One of the columns given in the PWLOBJ or PWLCON section of the MPS file was not defined in the COLUMNS section.

933 ***Breakpoints for piecewise linear constraint <index> not sorted and contradicting points (<x>,<y1>) and (<x>,<y2>) given.***

The breakpoints for this piecewise linear function were not given as a sorted list, and due to contradicting points this cannot be fixed by sorting (since it is unclear what the left- and right-limit of the discontinuity would be).

934 ***Breakpoints for piecewise linear constraint <index> not sorted, will be sorted internally.***

The breakpoints for the piecewise linear function were not given as a sorted list, these will be reordered internally.

935 ***Piecewise linear breakpoints not given consecutively for variable <column>, will reorder them.***

The breakpoints for this column were not given consecutively but with breakpoints for another variable in between, will be fixed internally.

936 ***Ignoring duplicate piecewise linear breakpoint (<x>,<y>) for piecewise linear constraint <index>.***

The same breakpoint was given twice for this piecewise linear constrant, the second one will be dropped.

937 ***Piecewise linear section for column <column> contained at most one breakpoint.***

The piecewise linear function over this column only contained a single breakpoint or no breakpoints at all, which is not enough to define the function.

938 ***Discontinuity at the <beginning/end> of piecewise linear constraint <index> without a matching bound.***

The first or last two breakpoints of the piecewise linear function shared the same value for the input variable. While this is allowed to model discontinuous functions, it is only allowed for the first or last points if they match the corresponding bound and do not leave the function undefined before or after them.

939 ***Non-convex piecewise linear function with unbounded domain for column <column>, may lead to unbounded LP relaxation for reformulation even if MIP is bounded.***

The piecewise linear function over the unbounded variable was non-convex and the sum of slopes for the lower and upper limit is negative, potentially leading to an unbounded LP relaxation for the reformulation even though the MIP might be bounded. Consider giving explicit bounds for this variable if possible.

940 ***General constraint type <type> for constraint <name> not supported, should be one of MAX, MIN, AND, OR, ABS, PWL***

The general constraint <name> was assigned the type <type>, which is not supported. The type should be "MAX", "MIN", "AND", "OR", "ABS", or "PWL".

941 ***Entry <entry> of general constraint <name> neither a column nor a constant***

One of the lines in general constraints section for constraint <name> consisted of <entry>, which is neither a column name nor a valid constant (and in case of free-format MPS-format also not a GENCONS-type).

942 ***<type> constraint <index> includes non-binary variable <name>. Only binaries allowed in AND/OR general constraints.***

The <index>-th general constraint, which is either an "AND" or an "OR" constraint, includes variable <name> which is not binary.

943 ***The absolute value constraint <index> includes <number> variables. ABS-constraints should consist of exactly 2 variables.***

The <index>-th general constraint, which is an "ABS" constraint, includes more than two variables.

944 ***The general constraint <index> consists of <number> elements. General constraints need to have at least two elements.***

The <index>-th general constraint consists of less than two elements. All general constraints must include at least two elements ("abs" exactly two).

947 ***PWLCON section refers to piecewise linear constraint <name>, which was not defined in the PWLNAM section.***

In the MPS format, all piecewise linear constraints must first be defined in the PWLNAM section before adding extreme points in the PWLCON section.

948 ***Missing bracket or comma in piecewise linear function or general constraint***

Piecewise linear functions should be given as (<val1>, <val2>), general constraints as <keyword> (<col1>, <col2>, ..., <colk>). In this case one of the commas or brackets was missing.

949 ***Invalid piecewise linear constraint number passed to <function>.Piecewise linear constraint number <val1> is invalid <val2>***

There has been an error in one of the arguments of function <function>. The <val1>th given piecewise linear constraint number <val2> is invalid. Please make sure that the piecewise linear constraint numbers are not smaller than 0 and not larger than the total number of piecewise linear constraints in the current problem, given by attribute `PWLCONS`.

950 ***Invalid general constraint number passed to <function>. <val1>-th General constraint number <val2> is invalid***

There has been an error in one of the arguments of function <function>. The <val1>-th given general constraint number <val2> is invalid. Please make sure that the general constraint numbers are not smaller than 0 and not larger than the total number of general constraints in the current problem, given by attribute `GENCONS`.

953 ***Column <index> appears as both resultant and input variable in piecewise linear/general constraint <index>.***

The resultant of the piecewise linear or general constraint is equal to one of the input columns. These need to be distinct.

954 ***__pwlobj() appearing outside of objective.***

The keyword __pwlobj(<col>) may be used inside the linear objective function to declare that the objective contribution of this variable will later be defined in the pwlobj section, but it should never appear in any other section.

955 ***Adjusting extremal breakpoints of piecewise linear constraint <index> to match bounds of input.***

The first/last breakpoint of piecewise linear constraint <index> (and potentially further ones afterwards) did not match the bound of the input variable and needed to be adjusted, even with presolve disabled. Note that this change will persist after the solve.

956 ***Given breakpoint (<val1>, <val2>) for piecewise linear constraint <index> with infinite values and non-constant slope.***

The piecewise linear constraint contained a breakpoint with infinite values. All breakpoints should have finite value, for piecewise linear constraints with unbounded rays a finite point along that ray should be given as the first/last breakpoint and the solver will extend that ray in a way that ensures numerical stability of the formulation.

957 ***Non-numerical value in problem definition as <name> of <name>.***

A non-numerical value (inf/nan) was given as a matrix or objective coefficient, right-hand side or column bound. The only allowed exceptions are lower bounds or right-hand sides of 'G' constraints with value -inf and upper bounds or right-hand sides of 'L' constraints at +inf.

958 ***Name missing from initial general constraint of type <type>.***

The first given general constraint of type <type> is missing a name, which is required in the MPS format.

959 ***Could not parse objective fields in row '<row>'.***

The MPS file contains a neutral row with additional data after the row name, but this data could not be parsed as valid objective fields. See Appendix A for information about how to specify multiple objectives in an MPS file.

960 ***Expected field '<field>' of objective row '<row>' to be an integer but found '<value>'. Using '<altvalue>' instead.***
Expected field '<field>' of objective row '<row>' to be non-negative but found '<value>'. Using '<altvalue>' instead.

The MPS file contains an objective row with an invalid field value. The displayed value will be used in place of the invalid value.

961 ***Objective row '<row>' has zero weight and will be disabled.***

The MPS file contains an objective row with zero weight. This objective is effectively disabled and will not be solved.

1001 ***Solution value redefined for column: <col>: <val1> -> <val2>***

Multiple definition of variable <col> is not allowed. Please use separate `SOLUTION` sections to define multiple solutions.

1002 ***Missing solution values in section <sec>. Only <val1> of <val2> defined***

Not all values were defined in the `SOLUTIONS` section. Variables with undefined values are set to 0.

1003 ***Please postsolve the problem first with `XPRSpostsolve` (postsolve).***

Not all values were defined in the `SOLUTIONS` section. Variables with undefined values are set to 0.

1004 ***Negative semi-continuous lower bound (<val>) for column <col> replaced with zero***

Wrong input parameter for the lower bound of a semi-continuous variable was modified to 0.

1005 ***Unrecognized column name : <col>***

No column with name <col> is present in the problem object while loading solution.

1006 ***Failed to capture solution information.***

Solution information is not available.

1020 ***Function <function> cannot be called here.***

The specified function can not be called.

1022 ***Error (<error>) while trying to run branching script.***

Branching script error.

1028 ***Unable to keep branch and bound tree memory usage below <val>Gb - currently using <val>Gb;***

The Optimizer was unable to keep the tree memory usage below the limit defined by the `TREEMEMORYLIMIT` control - the solve will continue but performance will be impaired.

1030 ***Duplicate names are not allowed - row/column/set/constraint <row1> would have same name as row/column/set/constraint <row2>: <name>***

Each row, column, special ordered set, piecewise linear constraint and general constraint should have unique name.

1034 ***Unknown column name <col> found in indicators***

Columns <col> found in the `INDICATORS` section is not defined in the `COLUMNS` section. Please check the name of the column.

1035 ***Unknown row name <row> found in indicators***

Row <row> found in the `INDICATORS` section is not defined in the `ROWS` section. Please check the name of the row.

1036 ***Unexpected indicator type : <type>***

Indicator type <type> found in the `INDICATORS` section is invalid. The type should be '`IF`'.

1037 ***Unexpected indicator active value : <value> for row <row>***

The value <value> found in the `INDICATORS` section is invalid. Values in this section should be either 0 and 1.

1038 ***Unsupported row type for indicator constraint <row>***

Rows configured as indicator constraints should have a type of '`L`' or '`G`'.

1039 ***Non-binary variable <col> used as an indicator binary***

The variable used in the condition part of an indicator constraint should be of type binary.

1054 ***Please use the FICO-SLP solver for general nonlinear problems. Contact " LICENSING_CONTACT_NAME " for a license.***

To solve general nonlinear problems the FICO-SLP solver can be used.

1055 ***Can not resume tree search - not currently solving a MIP.***

The tree search can not be resumed, the current problem is not mixed integer optimization.

1059 ***Unrecognized string identifier <id> passed to function <func>***

The string <id> given as input to the function <func> does not match any expected identifiers. Double-check spelling of <id> and consult documentation of function <func>, if available.

1071 ***Unable to dualize problems with quadratic coefficients***

The current version of the Optimizer only supports dualization of linear problems. Remove quadratic terms.

1074 ***Could not write tree to tree file.***

Branch-and-bound tree memory saving is disabled; re-enable this feature to allow the tree to be compressed and saved to the tree file.

1075 ***Message too long: message <mes> must be shorter than <maxlen> characters***

Could not write an error/info/warning message, due to a problem, row, column or set name being too long. Shorten all names s.t. they consist of at most <maxlen> characters.

1081 ***The Xpress-Optimizer license has been lost; this may be because of a connection failure with the license server.***

The connection to the Xpress license server has been lost and attempts to re-establish the connection failed. Check for network issues between client and license server.

1082 ***Tree heap create failed.***

Failed to create private heap for branch-and-bound tree; probably due to insufficient memory. If possible, try to free up memory on your system, reduce the problem size or set appropriate working limits.

1090 ***No license capacity.***

The FICO Xpress license file does not specify an Optimizer capacity; license has been incorrectly generated, please contact support@fico.com .

1092 ***Invalid scale factor***

An invalid scale factor was given for a row or a column. Scale factors are provided as the exponents of powers of two, and must be between 0 and `MAXSCALEFACTOR`

1093 ***Column scaling not allowed.***

Scaling of binary, integer and partial integer columns is not allowed.

1094 ***Invalid SOCP constraint detected.***

The Optimizer has detected an invalid SOCP constraint: a quadratic row has incorrectly been identified as a second order cone. Please contact support.

1097 ***The dependent variable of row (e.g. variable z in $z^2 \geqslant x^2 + y^2$) must be defined non-negative, otherwise the constraint is non-convex.***

The problem is not formulated in a standard second order conic formulation.

1098 ***Both dependent variables of row (e.g. variables u and v in $u * v \geqslant x^2 + y^2$) must be defined non-negative.***

The problem is not formulated in a standard second order conic formulation.

1100 ***Cut limit reached***

Failed to create cut, since the limit of storable cuts was reached. Please restart your solve with a less aggressive cutting strategy.

1101 ***Cannot call optimization routine recursively***

You cannot call `XPRSlpoptimize` or `XPRSmipoptimize` within a call of `XPRSlpoptimize` or `XPRSmipoptimize` on the same problem pointer. Either terminate the running optimization process or use a separate problem pointer. See also error code 1101.

1102 ***Presolve detected infeasibility on non-convex quadratic row***

All quadratic matrices in the quadratic constraints must be positive semi-definite or a second-order cone.

1103 ***Insufficient name buffer***

The supplied name buffer is too small. Allocate more memory for the buffer or use shorter names.

1104 ***No row/column/set names***

Tried to access a group of names that do not exist. Provide names to columns/rows/sets before doing so.

1105 ***Cannot create remote server session***

Creating the remote compute server session failed. This may have multiple reasons, the error message will specify the cause.

1106 ***Corrupt Xpress compute server interface***

There appears to be a mismatch between the libraries of Xpress. Please reinstall Xpress.

1107 ***Connection to remote server lost***

Connection to the compute Insight server was lost.

1108 ***Unable to create remote server job data package (disk full?)***

Solves using the compute server rely on creating temporary files.

1109 ***Remote server returned an error.***

This error forwards an error reported by the compute server.

1111 ***Remote server returned a warning***

This warning forwards a warning reported by the compute server.

1112 ***Remote server temporary directory provided creates a path too long.***

Names used during a compute server solve is directed from the probname and a unique identifier. Please shorten the probname.

1113 ***Limiting parallel MIP tasks***

When control RESOURCESTRATEGY is enabled, this warning notifies that the number of maximum MIP tasks have been limited to preserve memory.

1114 ***Error encountered during a sub-MIP solve. Please check the log for the cause.***

When the error message from the sub-MIP is available, it will be included. Otherwise, please check the solve log for details.

1115 ***Calling <function> is not supported during a solve.***

The <function> cannot be called during a solve.

1116 ***Paranthesis error in <name>.***

Names including parantheses are checked to include the same number of opening and closing parantheses.

1118 ***Ignoring unexpected flag '<flag>'.***

The flag <flag> passed into a command line call or API function was not understood and was ignored.

1119 ***Solves with Xpress compute server interface are not allowed in this application***

Integration with the Insight Compute interface has been enabled by setting the XPRESS_COMPUTE environment variable, but is not permitted for this application.

1120 ***Cannot disable compute mode when compute is enabled by the XPRS_COMPUTE environment variable.***

When Xpress is in compute mode as the environment variable XPRS_COMPUTE is set, it does not consume an Xpress license, hence local solves are not allowed.

1121 ***An error message produced when connecting to a compute server fails.***

This error is specific to when a connection fails when attempting to set the XPRS_COMPUTE control.

1122 ***Cannot use compute due to compute incompatible callbacks or controls set: 'reason'.***

When compute mode is enabled using the control XPRS_COMPUTE, problem setup compatiblity with copute mode is checked at the time a solve is initiated (as opposed to when enabled by environment variables, when the check is immediate).

1123 ***Cannot change the value of the COMPUTE control in a callback.***

It is not allowed to try to enter or leave compute mode while solving a problem, so in particular from a callback.

1124 ***An illegal escaped char is found: 'illegal sequence'.***

The MPS or LP file contained an escape sequence @HH, where HH is not a valid hexadecimal number. This error can only occur when the ESCAPE_NAMES control is enabled.

1125 ***Row <name> declared as both delayed row/indicator/model cut and delayed row/indicator/modelcut.***

A row was declared as more than one of delayed row / indicator constraint / model cut. These three options are mutually exclusive. An indicator constraint will always be part of the problem (with enforcement depending of the value of the indicator variable) and cannot be declared delayed or as a model cut. A row can also only be either delayed (in case it needs to be satisfied) or a model cut (if it is redundant).

1126 ***Row <name> flagged as indicator twice.***

A row was declared an indicator twice (for potentially different columns and values). A single row can only be implied by a single variable/term, if the same constraint should be implied by different terms, the row should be duplicated with each different term implying its own copy. If the indicator variable needs to be changed via the API, the old one first needs to be removed by a call to XPRSdelindicators before the new one can be added in XPRSsetindicators.

1127 ***The requested operation requires executing a subprocess, which is forbidden by security restrictions.***

In order to execute the command or function, the Optimizer needed to execute an subprocess, e.g., an external compression tool. This is forbidden when security restrictions have been enabled, which includes all FICO cloud environments.

1128 ***No Xpress Gobal Solver license found.***

You are not authorized to use the FICO Xpress Global solver. Please contact your local sales office to discuss upgrading to the FICO Xpress Global if you wish to use this command.

1130 ***Invalid objective number passed to <function>. Objective number <num> is invalid.***

Objective number <num> is invalid. There was an error in one of the arguments of function <function>. The objective number must be between 0 and `OBJECTIVES - 1`.

1132 ***Parameter <name> is invalid for function <function>. Must be non-negative.***

A negative absolute or relative tolerance was provided to function <function>. These parameters must be equal to or greater than zero.

1133 ***Ignoring additional objectives because an objective name was specified***

The MPS file contains a multi-objective problem, but only one objective will be read, because either the file contains an `OBJNAME` directive, or the `MPSOBJNAME` control was set to the name of an objective to read.

1134 ***Invalid solve number passed to <function>. Solve number <num> is less than the number of multi-objective solves that have been performed.***

An attribute of a multi-objective solve was queried, but fewer than <num> solves have been performed, so the attribute value is not available. The solve number must be between 0 and `SOLVEDOBJS - 1`. The solve may not have started yet or the multi-objective optimization may have been interrupted. Note that `SOLVEDOBJS` can be fewer than `OBJECTIVES` since all objectives with the same priority are combined into a single objective and solved at once.

1135 ***Expected one of priority, weight, reltol or abstol at <line.col>.***

An unknown objective attribute was found at column <col>, line <line> of the LP file. See A.3.8 for more information about how to specify multiple objectives in an LP file.

1136 ***Not writing neutral row '<row>' since its name conflicts with an objective name.***

The neutral row <row> has the same name as an objective, and they cannot both be written to the MPS file because they reside in the same namespace. The row will be omitted from the MPS file. This row is likely to be present because the `KEEPNROWS` control was set when the MPS file was read.

1137 ***The name of row '<row>' conflicts with an objective name.***

The MPS file cannot be written because the row <row> has the same name as an objective, and they both reside in the same namespace.

1138 ***Cannot add additional objectives to a problem containing quadratic coefficients. Remove the quadratic objective coefficients first.***
Cannot add quadratic objective coefficients to a problem containing multiple objectives. Remove the additional objectives first.

The current version of the Optimizer only supports multi-objective problems where all objectives are linear. To work around this limitation, reformulate the problem with a linear objective function and express the quadratic terms via transfer rows.

1139 ***Cannot add additional objectives to a problem containing nonlinear coefficients. Remove the nonlinear objective coefficients first.***
Cannot add nonlinear objective coefficients to a problem containing multiple objectives. Remove the additional objectives first.

The current version of the Optimizer only supports multi-objective problems where all objectives are linear. To work around this limitation, reformulate the problem with a linear objective function and express the nonlinear terms via transfer rows.

1140 ***'<feature>' is not supported in this release.***

This feature is not supported for global solves in the current release.

1141 ***Invalid value for control.***

An attempt was made to set a control to a value that is outside the range of valid values for that control.

1142 ***Global solver not supported in the presence of <feature>.***

The global solver cannot be used in combination with certain features, including user functions and multistart.

1143 ***Overflow when converting <value> to a 32-bit integer.***

A 64-bit integer must be converted to a 32-bit integer but cannot be represented as 32-bit integer.

1144 ***Invalid offset passed to <function>. Offset <value> at position <index> is invalid.***

When providing an array of offsets into other arrays, the offsets must be non-decreasing. The value <value> at position <index> was found to be smaller than the value at position <index>-1.

1145 ***Invalid argument in '<function>'***

An invalid argument was passed to the specified function.

1146 ***Cannot delete column <column index> because it is referenced by PWL <PWL index>***

A column cannot be deleted if it is still referenced by a piecewise linear function. This would render the definition of the function invalid. Delete the function first.

1147 ***Cannot delete column <column index> because it is referenced by general constraint <general constraint index>***

A column cannot be deleted if it is still referenced by a general constraint. This would render the definition of the general constraint invalid. Delete the general constraint first.

1148 ***Number of vcores detected on the system exceeds the maximum authorized by the license.***

The number of vcores detected on the system cannot exceed the maximum set by the license.

1149 ***Unknown <index> in save file ignored***

A control or attribute with the given index does not exist in this version of the Xpress optimizer. The value specified in the save file will be ignored.

1150 ***control/attribute with unknown or wrong type <type> in save file ignored***

The type in the save file for control/attribute <index> does not match the type in this version of the Xpress optimizer. The setting in the save file is ignored.

1151 ***Maximum length of nonlinear formulas exceeded. The maximum is <num>***

The Optimizer cannot solve your problem since the number of tokens in nonlinear formulas exceeds the maximum allowed. Please contact your local sales office to discuss upgrading your license to solve larger problems. Note that formulas containing only min/max/abs/pwl terms do not count towards this limit.

1152 ***Problem has too many nonlinear user functions. The maximum is <num>***

The Optimizer cannot solve your problem since the number of user functions exceeds the maximum allowed. Please contact your local sales office to discuss upgrading your license to solve larger problems.

1153 ***'<function>' is not supported for nonlinear solves.***

This function is only supported for linear and/or mixed-integer linear or convex quadratic solves but not for global, Knitro or SLP solves.

1155 ***A column with the protected name = has to be fixed to one.***

The name '=' is reserved for a column which is fixed to one to be multiplied with nonlinear formulas and may not be assigned to any other column.

1156 ***Implied division by zero of column '<column>' in row '<row>' detected in the problem.***

The solver detected that the denominator of a division has to be zero (either explicitly in the model or it has to be zero to satisfy other constraints). The column that is being divided may be an original column or an auxiliary column introduced by the global solver.

9999 ***Generic error message***

Please contact support@fico.com.

Appendix

Appendix A

Log and File Formats

A.1 File Types

The Optimizer generates or inputs a number of files of various types as part of the solution process. By default these all take file names governed by the problem name (*problem_name*), but distinguished by their three letter extension. The file types associated with the Optimizer are as follows:

Extension	Description	File Type
.alt	Matrix alteration file, input by `XPRSalter` (`ALTER`).	ASCII
.asc	CSV format solution file, output by `XPRSwritesol` (`WRITESOL`).	ASCII
.bss	Basis file, output by `XPRSwritebasis` (`WRITEBASIS`), input by `XPRSreadbasis` (`READBASIS`).	ASCII
.csv	Output file, output by `XPRSiiswrite`.	ASCII
.dir	Directives file (MIP only), input by `XPRSreaddirs` (`READDIRS`).	ASCII
.glb	Tree file (MIP only), used by `XPRSmipoptimize` (`MIPOPTIMIZE`).	Binary
.hdr	Solution header file, output by `XPRSwritesol` (`WRITESOL`).	ASCII
.lp	LP format matrix file, input by `XPRSreadprob` (`READPROB`).	ASCII
.mps	MPS / XMPS format matrix file, input by `XPRSreadprob` (`READPROB`).	ASCII
.prt	Fixed format solution file, output by `XPRSwriteprtsol` (`WRITEPRTSOL`).	ASCII
.sol	Solution file created by `XPRSwritebinsol` (`WRITEBINSOL`).	Binary
.slx	Solution file created by `XPRSwriteslxsol` (`WRITESLXSOL`).	ASCII
.xtm	Tuner method created by `XPRStunerwritemethod`.	ASCII
.xtr	Tuner result created by `XPRStune`.	XML
.json	Remote Solving Configuration file	JSON

In the following sections we describe the formats for a number of these.

Note that CSV stands for comma-separated-values text file format.

A.1.1 File Compression

Many Optimizer commands support reading and writing compressed files. Support for Gzip compression is built in, so commands like the following will work as expected:

```
readprob problem.mat.gz
writeprob problem.gz
```

The Optimizer can also support additional compression formats, as long as the necessary command-line tools are available in the `PATH` environment variable. The following table shows the supported compression formats along with the required tools:

Extension	Required Tool
`gz`	Built-in support
`zip`	`zip` / `unzip`
`tar, tgz`	`tar`
`bz2, bzip2`	`bzip2`
`7z`	`7z` / `7za` / `7zr`
`xz`	`xz`
`lz4`	`lz4`
`Z`	`compress` / `decompress`

In addition, the `7z` tool can be used to read and write most of the formats listed above.

NB: compression using external tools is not supported in cloud environments due to security restrictions.

A.2 XMPS Matrix Files

The FICO Xpress Optimizer accepts matrix files in LP or MPS format, and an extension of this, XMPS format. In that the latter represents a slight modification of the industry-standard, we provide details of it here.

XMPS format defines the following fields:

Field	1	2	3	4	5	6
Columns	2-3	5-12	15-22	25-36	40-47	50-61

The following sections are defined:

NAME	the matrix name;
ROWS	introduces the rows;
COLUMNS	introduces the columns;
QUADOBJ / QMATRIX	introduces a quadratic objective function;
QCMATRIX	introduces the quadratic constraints;
DELAYEDROWS	introduces the delayed rows;
MODELCUTS	introduces the model cuts;
INDICATORS	introduces the indicator contraints;
SETS	introduces SOS definitions;
RHS	introduces the right hand side(s);
RANGES	introduces the row ranges;
BOUNDS	introduces the bounds;
GENCONS	introduces general constraints;
ENDATA	signals the end of the matrix.
compatibility	additional sections for extensions of the MPS format that can be read but not written.

All section definitions start in column `1`.

A.2.1 NAME section

Format:	Cols 1-4	Field 3
	`NAME`	*model_name*

A.2.2 OBJSENSE section

Format:	Cols 1-8	Field 3
	`OBJSENSE`	*sense*

Alternatively, the sense can be provided on a separate line in Field 2:

Field 1	Field 2
blank	*sense*

The objective sense can be `MAXIMIZE` or `MINIMIZE`, or any prefix of those values of length at least 3, e.g., `MAX` or `MAXIM`. It is case insensitive.

If the `OBJSENSE` section is not present, the objective sense is assumed to be minimization.

A.2.3 ROWS section

Format:	Cols 1-4
	`ROWS`

followed by row definitions in the format:

Field 1	Field 2	Field 3	Field 4	Field 5	Field 6
type	*row_name*	*priority*	*weight*	*abstol*	*reltol*

The row types (Field 1) are:

N	unconstrained (for objective functions);
L	less than or equal to;
G	greater than or equal to;
E	equality.

Fields 3-6 are used for multi-objective problems, and they may only be populated for neutral rows (objectives). In a multi-objective problem, you must specify either the weight and priority of every objective, or you must specify the weight, priority, absolute and relative tolerances for every objective. No other combination of fields is valid, in order to remain compatible with older MPS files. For more information about multi-objective problems, please refer to 5.12.

A.2.4 COLUMNS section

Format:	Cols 1-7
	`COLUMNS`

followed by columns in the matrix in column order, i.e. all entries for one column must finish before those for another column start, where:

Field 1	Field 2	Field 3	Field 4	Field 5	Field 6
blank	*col*	*row1*	*value1*	*row2*	*value2*

specifies an entry of *value1* in column *col* and row *row1* (and *value2* in *col* and row *row2*). The Field 5/Field 6 pair is optional.

A.2.5 QUADOBJ / QMATRIX section (Quadratic Programming only)

A quadratic objective function can be specified in an MPS file by including a `QUADOBJ` or `QMATRIX` section. For fixed format XMPS files, the section format is as follows:

Format:	Cols 1-7
	`QUADOBJ`

or

Format:	Cols 1-7
	`QMATRIX`

followed by a description of the quadratic terms. For each quadratic term, we have:

Field 1	Field 2	Field 3	Field 4
blank	*col1*	*col2*	*value*

where *col1* is the first variable in the quadratic term, *col2* is the second variable and *value* is the associated coefficient from the `Q` matrix. In the `QMATRIX` section all nonzero `Q` elements must be specified. In the `QUADOBJ` section only the nonzero elements in the upper (or lower) triangular part of `Q` should be specified. In the `QMATRIX` section the user must ensure that the `Q` matrix is symmetric, whereas in the `QUADOBJ` section the symmetry of `Q` is assumed and the missing part is generated automatically.

Note that the Q matrix (i.e., the quadratic part of the objective) has an implicit factor of 0.5 when included in the objective function.
This means, for instance that an objective function of the form

$$5x^2 + 7xy + 9y^2$$

is represented in a `QUADOBJ` section as:

```
QUADOBJ
x      x      10
x      y       7
y      y      18
```

(The additional term '`y x 7`' is assumed which is why the coefficient is not doubled); and in a `QMATRIX` section as:

```
QMATRIX
x      x      10
x      y       7
y      x       7
y      y      18
```

The `QUADOBJ` and `QMATRIX` sections must appear somewhere after the `COLUMNS` section and must only contain columns previously defined in the columns section. Columns with no elements in the problem matrix must be defined in the `COLUMNS` section by specifying a (possibly zero) cost coefficient.

A.2.6 QCMATRIX section (Quadratic Constraint Programming only)

Quadratic constraints may be added using `QCMATRIX` sections.

Format:	Cols 1-8	Field 3
	`QCMATRIX`	`row_name`

Each constraint having quadratic terms should have it's own `QCMATRIX` section. The `QCMATRIX` section exactly follows the description of the `QMATRIX` section, i.e. for each quadratic term, we have:

Field 1	Field 2	Field 3	Field 4
blank	*col1*	*col2*	*value*

where `col1` is the first variable in the quadratic term, `col2` is the second variable and value is the associated coefficient from the Q matrix. All nonzero Q elements must be specified. The user must ensure that the Q matrix is symmetric. For instance a constraint of the form

$$qc1 : x + 5x^2 + 7xy + 9y^2 \leqslant 2$$

is represented as:

```
NAME example
ROWS
 L qc1
COLUMNS
   x      qc1     1
   y      qc1     0
QCMATRIX qc1
   x      x       5
   x      y       3.5
   y      x       3.5
   y      y       9
RHS
 RHS1     qc1     2
END
```

The `QCMATRIX` sections must appear somewhere after the `COLUMNS` section and must only contain columns previously defined in the columns section. Columns with no elements in the problem matrix must be defined in the `COLUMNS` section by specifying a (possibly zero) cost coefficient. `QCMATRICES` must be defined only for rows of type `L` or `G` and must have no range value defined in the `RANGE` section.

The FICO Xpress Optimizer can solve both convex and non-convex (MI)QPs and (MI)QCPs. For non-convex problems, some of the restrictions of FICO Xpress Global apply, e.g., callback usability is limited. A quadratic problem is convex, when the Q matrix of the objective is positive semi-definite for minimization and negative semi-definite for maximization. A quadratically constrained problem is convex, when all row Q matrices are positive semi-definite for $\leqslant$ rows and negative semi-definite for $\geqslant$ rows. For convex (MI)QCPs, equations must not contain a quadratic part.

NOTE: technically, there is one exception for the restriction on the row type being L or G. If the row is the first nonbinding row (type N) then the section is treated as a QMATRIX section instead. Please be aware, that this also means that the objective specific implied divider of 2 will be assumed (Q matrix has an implicit factors of 0.5 when included in the objective function, see the QMATRIX section). It's probably much better to use the QMATRIX or QUADOBJ sections to define quadratic objectives.

A.2.7 DELAYEDROWS section

This specifies a set of rows in the matrix that will be treated as delayed rows during a tree search. These are rows that must be satisfied for any integer solution, but will not be loaded into the active set of constraints until required.

This section should be placed between the ROWS and COLUMNS sections. A delayed row may be of type L, G or E. Each row should appear either in the ROWS or the DELAYEDROWS section, not in both. Otherwise, the format used is the same as of the ROWS section.

Format:	Cols 1-11
	DELAYEDROWS

followed by row definitions in the format:

Field 1	Field 2
type	*row_name*

NOTE: For compatibility reasons, section names DELAYEDROWS and LAZYCONS are treated as synonyms.

A.2.8 MODELCUTS section

This specifies a set of rows in the matrix that will be treated as model cuts during a tree search. During presolve the model cuts are removed from the matrix. Following optimization, the violated model cuts are added back into the matrix and the matrix is re-optimized. This continues until no violated cuts remain. This section should be placed between the ROWS and COLUMNS sections. Model cuts may be of type L, G or E. The model cuts must be "true" model cuts, in the sense that they are redundant at the optimal MIP solution. The Optimizer does not guarantee to add all violated model cuts, so they must not be required to define the optimal MIP solution.

Each row should appear either in the ROWS, DELAYEDROWS or in the MODELCUTS section, not in any two or three of them. Otherwise, the format used is the same as of the ROWS section.

Format:	Cols 1-9
	MODELCUTS

followed by row definitions in the format:

Field 1	Field 2
type	*row_name*

NOTE: A problem is not allowed to consists solely from model cuts. For compatibility reasons, section names MODELCUTS and USERCUTS are treated as synonyms.

A.2.9 INDICATORS section

This specifies that a set of rows in the matrix will be treated as indicator constraints during a tree search. These are constraints that must be satisfied only when their associated controlling binary variables have specified values (either 0 or 1).

This section should be placed after any `QUADOBJ`, `QMATRIX` or `QCMATRIX` sections. The section format is as follows:

Format:	Cols 1-10
	`INDICATORS`

Subsequent records give the associations between rows and the controlling binary columns, with the following form:

Field 1	Field 2	Field 3	Field 4
type	*row_name*	*col_name*	*value*

which specifies that the row `row_name` must be satisfied only when column `col_name` has value `value`. Here type must always be `IF` and value can be either 0 or 1. Also referenced rows must be of type `L` or `G` only, and referenced columns must be binary.

A.2.10 SETS section (Integer Programming only)

Format:	Cols 1-4
	`SETS`

This record introduces the section which specifies any Special Ordered Sets. If present it must appear after the `COLUMNS` section and before the `RHS` section. It is followed by a record which specifies the type and name of each set, as defined below.

Field 1	Field 2
type	*set*

Where *type* is `S1` for a Special Ordered Set of type 1 or `S2` for a Special Ordered Set of type 2 and *set* is the name of the set.

Subsequent records give the set members for the set and are of the form:

Field 1	Field 2	Field 3	Field 4	Field 5	Field 6
blank	*set*	*col1*	*value1*	*col2*	*value2*

which specifies a set member *col1* with reference value *value1* (and *col2* with reference value *value2*). The Field 5/Field 6 pair is optional.

A.2.11 RHS section

Format:	Col 1-3
	`RHS`

followed by the right hand side as defined below:

Field 1	Field 2	Field 3	Field 4	Field 5	Field 6
blank	*rhs*	*row1*	*value1*	*row2*	*value2*

specifies that the right hand side column is called *rhs* and has a value of *value1* in row *row1* (and a value of *value2* in row *row2*). The Field 5/Field 6 pair is optional.

A.2.12 RANGES section

Format:	Cols 1-6
	`RANGES`

followed by the right hand side ranges defined as follows:

Field 1	Field 2	Field 3	Field 4	Field 5	Field 6
blank	*rng*	*row1*	*value1*	*row2*	*value2*

specifies that the right hand side range column is called *rng* and has a value of *value1* in row *row1* (and a value of *value2* in row *row2*). The Field 5/Field 6 pair is optional.

For any row, if b is the value given in the `RHS` section and r the value given in the `RANGES` section, then the activity limits below are applied:

Row Type	**Sign of r**	**Upper Limit**	**Lower Limit**
G	+	b+r	b
L	+	b	b-r
E	+	b+r	b
E	-	b	b+r

A.2.13 BOUNDS section

Format:	Cols 1-6
	`BOUNDS`

followed by the bounds acting on the variables:

Field 1	Field 2	Field 3	Field 4
type	*blank*	*col*	*value*

The Linear Programming bound types are:

UP	for an upper bound;
LO	for a lower bound;
FX	for a fixed value of the variable;
FR	for a free variable;
MI	for a non-positive ('minus') variable;
PL	for a non-negative ('plus') variable (the default).

There are six additional bound types specific to Integer Programming:

UI	for an upper bounded general integer variable;
LI	for a lower bounded general integer variable;
BV	for a binary variable;
SC	for a semi-continuous variable;
SI	for a semi-continuous integer variable;
PI	for a partial integer variable.

The value specified is an upper bound on the largest value the variable can take for types `UP`, `FR`, `UI`, `SC` and `SI`; a lower bound for types `LO` and `LI`; a fixed value for type `FX`; and ignored for types `BV`, `MI` and `PL`. For type `PI` it is the switching value: below which the variable must be integer, and above which the variable is continuous. If a non-integer value is given with a `UI` or `LI` type, only the integer part of the value is used.

- **Integer variables** may only take integer values between `0` and the upper bound. Integer variables with an upper bound of unity are treated as binary variables.
- **Binary variables** may only take the values `0` and `1`. Sometimes called 0/1 variables.
- **Partial integer variables** must be integral when they lie below the stated value, above that value they are treated as continuous variables.
- **Semi-continuous variables** may take the value zero or any value between a lower bound and some finite upper bound. By default, this lower bound is `1.0`. Other positive values can be specified as an explicit lower bound. For example

```
BOUNDS
LO x 0.8
SC x 12.3
```

  means that `x` can take the value zero or any value between `0.8` and `12.3`.

- **Semi-continuous integer variables** may take the value zero or any integer value between a lower bound and some finite upper bound.

A.2.14 GENCONS section

Format:	Cols 1-7
	`GENCONS`

This record introduces the section which specifies any general constraints, namely min, max, and, or, abs, pwl constraints. If present it must appear after the `COLUMNS` section. It is followed by a record which specifies the type and name of each general constraint, as defined below.

Field 1	Field 2
type	*name*

Where *type* is `MAX` for a maximum-constraint, `MIN` for a minimum-constraint, `AND` for an and-constraint, `OR` for an or-constraint, `ABS` for an absolute-value-constraint or `PWL` for a piecewise linear constraint and *name* is the name of the general constraint.

Subsequent records for min/max/and/or/abs give the elements for the constraint and are of the form:

Field 1	Field 2
blank	*col/val*

For all general constraints, the first given element (which needs to be the name of a column) will be the so-called "resultant". For the max- and min-constraints, the resultant is followed by an arbitrary number of further column names or values, and the resultant should be the maximum/minimum of the remaining columns and values. For the and- and or-constraints the resultant is followed by an arbitrary number of further column names, where all the columns (including the resultant) need to be binary, and the resultant will be one if and only if all (and) or at least one (or) of the remaining variables are one. For the abs-constraint, the resultant should be followed by exactly one futher column name, and the resultant will take the absolute value of the other column.

As an example, the constraint `z = max {x, y, 5.0}` could be written as

```
GENCONS
MAX m1
z
x
y
5.0
```

For piecewise linear constraints the format is slightly different, consisting of exactly one line of the form:

Field 1	Field 2
col1	*col2*

and at least three lines of the form:

Field 1	Field 2
val1	*val2*

The first line defines the two variables that should be restricted by a piecewise linear relationship and the points given in the remaining lines will define the piecewise linear function, which is extended beyond the first and last point according to the slope of the previous ones. For instance the piecewise linear constraint

```
y = 0, if x < 0
y = x, if 0 <= x <= 5
y = 2x - 5, if 5 < x <= 10
y = 15, if x > 10
```

could be represented as:

```
GENCONS
PWL p1
   x  y
-1 0
0  0
5  5
10 15
11 15
```

A.2.15 ENDATA section

Format:	Cols 1-6
	`ENDATA`

is the last record of the file.

A.2.16 Compatibility

The optimizer is also able to read in some further sections defined by extensions of the LP format. This includes the SOS section, which is a different way of writing down special ordered sets, and the following sections that offer different ways of formulating piecewise linear constraints and objectives:

A.2.17 PWLOBJ section

Piecewise linear objective functions may be added using `PWLOBJ` sections.

Format:	Cols 1-6
	`PWLOBJ`

The piecewise linear objective function is defined via its extreme points, i.e., the function itself is given by all convex combinations of neighboring extreme points as well as the infinite rays defined by the first two and last two points. Each row of the `PWLOBJ` section defines one extreme point for one column.

Field 1	Field 2	Field 3	Field 4
blank	*col*	*value1*	*value2*

where `col` is the variable whose objective contribution is defined through the piecewise linear function and `value2` is the objective contribution if the variable takes `value1`. For instance the piecewise linear objective function

```
0, if x < 0
x, if 0 <= x <= 5
2x - 5, if 5 < x <= 10
15, if x > 10
```

could be represented as:

```
PWLOBJ
   x         -1 0
   x         0  0
   x         5  5
   x         10 15
   x         11 15
```

If there are piecewise linear objective functions for multiple variables, these should be given consecutively (i.e., first all extreme points for x, then all for y). Furthermore, for each variable, the extreme points should be sorted according to non-decreasing `value1`. The piecewise linear functions do not necessarily need to be continuous, in this case two extreme points with identical `value1` and different `value2` can be given and the first one will be used as the lefthand-limit and the second one as the righthand-limit. Note that for `value1` itself, both `value2` can appear in the solution due to tolerances.

A.2.18 PWLNAM section

`PWLNAM` is the first of the two sections defining piecewise linear constraints.

Format:	Cols 1-6
	`PWLNAM`

Similar to the piecewise linear objective, the piecewise linear constraints will mainly be defined through its extreme points, which happens in the `PWLCON` section. In addition to that, however, the two variables involved in the restriction y = f(x), with piecewise linear function f, need to be specified, and additionally a pre- and postslope are given, defining the slope of the piecewise linear function before the first and after the last extreme point. Each piecewise linear function needs to be named, to later refer to it in the `PWLCON` section when specifying the extreme points.

Field 1	Field 2	Field 3	Field 4	Field 5	Field 6
blank	*name*	*col1*	*col2*	*value1*	*value2*

where `col1` is the resulting variable (y above), `col2` is the input variable (x above), `value1` is the preslope defining the piecewise linear function up to the first extreme point and `value1` is the postslope defining it after the last extreme point.

A.2.19 PWLCON section

`PWLCON` is the second of the two sections defining piecewise linear constraints.

Format:	Cols 1-6
	`PWLCON`

Each piecewise linear constraint introduced in the `PWLNAM` section needs to be further specified through its extreme points, defining the behaviour between the pre- and postlope. Each line consists of the name of a piecewise linear constraint introduced in the `PWLNAM` followed by a list of extreme points. Similar to the PWLOBJ section, the functions can be discontinuous, in which case the extreme points have to be given in the correct order.

Field 1	Field 2	Field 3	Field 4
blank	*name*	*value1*	*value2*

where `name` is the name of a piecewise linear function introduced in the `PWLNAM` section and `value1` and `value2` define an extreme point, where `value1` is the value of the input variable and `value1` is the corresponding output value. For instance the piecewise linear constraint

```
y = 0, if x < 0
y = x, if 0 <= x <= 5
y = 2x - 5, if 5 < x <= 10
y = 15, if x > 10
```

could be represented as:

```
PWLNAM
 pwc1   y      x1      0    0
PWLCON
 pwc1    0    0
 pwc1    5    5
 pwc1    10   15
```

A.3 LP File Format

Matrices can be represented in text files using either the MPS file format (`.mps` or `.mat` files) or the LP file format (`.lp` files). The LP file format represents matrices more intuitively than the MPS format in that it expresses the constraints in a row-oriented, algebraic way. For this reason, matrices are often written to LP files to be examined and edited manually in a text editor. Note that because the variables are 'declared' as they appear in the constraints during file parsing the variables may not be stored in the FICO Xpress Optimizer memory in the way you would expect from your enumeration of the variable names. For example, the following file:

```
Minimize
obj: - 2 x3

Subject To
c1: x2 - x1 <= 10
c2: x1 + x2 + x3 <= 20

Bounds
x1 <= 30

End
```

after being read and rewritten to file would be:

```
\Problem name:
Minimize
 - 2 x3

Subject To
c1: x2 - x1 <= 10
c2: x3 + x2 + x1 <= 20

Bounds
x1 <= 30

End
```

Note that the last constraint in the output `.lp` file has the variables in reverse order to those in the input `.lp` file. The ordering of variables in the last constraint of the rewritten file is the order that the variables were encountered during file reading. Also note that although the optimal solution is unique for this particular problem in other problems with many equal optimal solutions the path taken by the solver may depend on the variable ordering and therefore by changing the ordering of your constraints in the .lp file may lead to different solution values for the variables.

A.3.1 Rules for the LP file format

The following rules can be used when you are writing your own `.lp` files to be read by the FICO Xpress Optimizer.

A.3.2 Comments and blank lines

Text following a backslash (\) and up to the subsequent carriage return is treated as a comment. Blank lines are ignored. Blank lines and comments may be inserted anywhere in an `.lp` file. For example, a common comment to put in LP files is the name of the problem:

```
\Problem name:  prob01
```

A.3.3 File lines, white space and identifiers

White space and carriage returns delimit variable names and keywords from other identifiers. Keywords are case insensitive. Variable names are case sensitive. Although it is not strictly necessary, for clarity of your LP files it is perhaps best to put your section keywords on their own lines starting at the first character position on the line. There is no maximum on the length names of on the length of input lines. Lines may be broken for continuation wherever you may use white space.

A.3.4 Sections

The LP file is broken up into sections separated by section keywords. The following are a list of section keywords you can use in your LP files. A section started by a keyword is terminated with another section keyword indicating the start of the subsequent section.

Section keywords	Synonyms	Section contents
`maximize` or `minimize`	`maximum max minimum min`	One linear expression describing the objective function.
`subject to`	`subject to:` `such that` `st` `s.t.` `st.` `subjectto` `suchthat` `subject` `such`	A list of constraint expressions.
`bounds`	`bound`	A list of bounds expressions for variables.
`integers`	`integer` `ints` `int`	A list of variable names of integer variables. Unless otherwise specified in the bounds section, the default relaxation interval of the variables is [0, 1].
`generals`	`general` `gens` `gen`	A list of variable names of integer variables. Unless otherwise specified in the bounds section, the default relaxation interval of the variables is [0, `XPRS_PLUSINFINITY`].
`binaries`	`binary` `bins` `bin`	A list of variable names of binary variables.
`semi-continuous`	`semi continuous` `semis` `semi` `s.c.`	A list of variable names of semi-continuous variables.
`semi integers`	`s.i.`	A list of semi-integer bound expressions for variables.
`partial integer`	`p.i.`	A list of variable names of partial integer variables.
`general constraints`	`general constraint` `gencons` `g.c.`	A list of min/max/and/or/abs and piecewise linear constraints.
`ranges`	`range`	A list of constraint ranges.

Variables that do not appear in any of the variable type registration sections (i.e., `integers`, `generals`, `binaries`, `semi-continuous`, `semi integer`, `partial integer`) are defined to be continuous variables by default. That is, there is no section defining variables to be continuous variables.

With the exception of the objective function section (`maximize` or `minimize`) and the constraints section (`subject to`), which must appear as the first and second sections respectively, the sections may appear in any order in the file. The only mandatory section is the objective function section. Note that you can define the objective function to be a constant in which case the problem is a so-called constraint satisfaction problem. The following two examples of LP file contents express empty problems with constant objective functions and no variables or constraints.

Empty problem 1:

```
Minimize

End
```

Empty problem 2:

```
Minimize

0

End
```

The end of a matrix description in an LP file can be indicated with the keyword `end` entered on a line by itself. This can be useful for allowing the remainder of the file for storage of comments, unused matrix definition information or other data that may be of interest to be kept together with the LP file.

A.3.5 Names

Variable, constraint and other names can use any of the alphanumeric characters (a-z, A-Z, 0-9) and any of the following symbols:

```
!"#$%&/,.;?_ `'()|~'
```

A name can not begin with a number, a period, a comma, a slash or parenthesis and should not be the same as any of the section or subsection names. Care should be taken using the characters E or e since these may be interpreted as exponential notation for numbers.

Curly braces (`{` and `}`) cannot be used in Xpress LP files since these are used as delimiters in non-linear formulas.

The `@` (at) character has a special meaning in Xpress: it is used to escape characters that are otherwise invalid in LP files. See also the `ESCAPENAMES` control.

A.3.6 Linear expressions

Linear expressions are used to define the objective function and constraints. Terms in a linear expression must be separated by either a + or a − indicating addition or subtraction of the following term in the expression. A term in a linear expression is either a variable name or a numerical coefficient followed by a variable name. It is not necessary to separate the coefficient and its variable with white space or a carriage return although it is advisable to do so since this can lead to confusion. For example, the string " 2e3x" in an LP file is interpreted using exponential notation as 2000 multiplied by variable x rather than 2 multiplied by variable e3x. Coefficients must precede their associated variable names. If a coefficient is omitted it is assumed to be 1.

A.3.7 Objective function

The objective function section can be written in a similar way to the following examples using either of the keywords `maximize` or `minimize`. Note that the keywords `maximize` and `minimize` are not used for anything other than to indicate the following linear expression to be the objective function. Note the following two examples of an LP file objective definition:

```
Maximize
 - 1 x1 + 2 x2 + 3x + 4y
```

or

```
Minimize
- 1 x1 + 2 x2 + 3x + 4y
```

No line continuation character is required to break the objective function across multiple lines and it can be broken wherever you may use white space.

The objective function can be named in the same way as for constraints (see later).

A.3.8 Multiple objectives

For multi-objective problems, the priority, weight, absolute and relative tolerances can be specified before the objective expression, for example:

```
Minimize
Obj1: Priority=3
x1 + x2
Obj2: Priority=2 Weight=-1
x3 - x4
Obj3: Priority=1 AbsTol=1 RelTol=0.01
x5 + 2 x6
```

For more information about multi-objective problems, please refer to 5.12.

A.3.9 Constraints

The section of the LP file defining the constraints is preceded by the keyword `subject to`. Each constraint definition must begin on a new line. A constraint may be named with an identifier followed by a colon before the constraint expression. Constraint names must follow the same rules as variable names. If no constraint name is specified for a constraint then a default name is assigned of the form `C0000001`, `C0000002`, `C0000003`, etc. Constraint names are trimmed of white space before being stored.

The constraints are defined as a linear expression in the variables followed by an indicator of the constraint's sense and a numerical right-hand side coefficient. The constraint sense is indicated intuitively using one of the tokens: >=, <=, or =. For example, here is a named constraint:

```
depot01: - x1 + 1.6 x2 - 1.7 x3 <= 40
```

Note that tokens > and < can be used, respectively, in place of the tokens >= and <=.

No line continuation character is required when breaking a constraint across multiple lines, and lines may be broken for continuation wherever you may use white space.

A.3.10 Delayed rows

Delayed rows are defined in the same way as general constraints, but after the `delayed rows` keyword. Note that delayed rows shall not include quadratic terms. The definition of constraints, delayed rows and model cuts should be sequentially after each other.

For example:

```
Minimize
obj: x1 + x2
subject to
x1 <= 10
x1 + x2 >= 1
delayed rows
x1 >= 2
```

```
end
```

For compatibility reasons, the term `lazy constraints` is used as a synonym to `delayed rows`.

A.3.11 Model cuts

Model cuts are defined in the same way as general constraints, but after the `model cuts` keyword. Note that model cuts shall not include quadratic terms. The definition of constraints, delayed rows and model cuts should be sequentially after each other.

For example:

```
Minimize
obj: x1 + x2
subject to
x1 <= 10
x1 + x2 >= 1
model cuts
x1 >= 2
end
```

For compatibility reasons, the term `user cuts` is used as a synonym to `model cuts`.

A.3.12 Indicator contraints

Indicator constraints are defined in the constraints section together with general constraints (that is, under the keyword `subject to`). The syntax is as follows:

```
constraint_name: col_name = value -> linear_inequality
```

which means that the constraint `linear_inequality` should be enforced only when the variable `col_name` has value `value`.

As for general constraints, the `constraint_name:` part is optional; `col_name` is the name of the controlling binary variable (it must be declared as binary in the `binaries` section); and `value` may be either 0 or 1. Finally the `linear_inequality` is defined in the same way as for general constraints.

For example:

```
Minimize
obj: x1 + x2
subject to
x1 + 2 x2 >= 2
x1 = 0 -> x2 >= 2
binary
x1
end
```

A.3.13 Bounds

The list of bounds in the bounds section are preceded by the keyword `bounds`. Each bound definition must begin on a new line. Single or double bounds can be defined for variables. Double bounds can be defined on the same line as `10 <= x <= 15` or on separate lines in the following ways:

```
10 <= x
15 >= x
```

or

```
x >= 10
x <= 15
```

If no bounds are defined for a variable the FICO Xpress Optimizer uses default lower and upper bounds. An important point to note is that the default bounds are different for different types of variables. For continuous variables and variables declared in the `generals` section, the interval defined by the default bounds is [0, `XPRS_PLUSINFINITY`], while for variables declared in the `integers` section (see later) the relaxation interval defined by the default bounds is [0, 1]. Note that the constant `XPRS_PLUSINFINITY` is defined in the FICO Xpress Optimizer header files in your FICO Xpress Optimizer libraries package.

If a single bound is defined for a variable the FICO Xpress Optimizer uses the appropriate default bound as the second bound. Note that negative upper bounds on variables must be declared together with an explicit definition of the lower bound for the variable. Also note that variables can not be declared in the bounds section. That is, a variable appearing in a bounds section that does not appear in the objective section or in the constraint section is ignored.

If multiple lower bounds or multiple upper bounds are defined for the same variable, only the last definition is applied. For example, if the bounds section contains the following bound definitions on a variable x

```
-1 <= x <= 1
x <= 2
```

then the first upper bound of `x <= 1` will be ignored due to the second definition of `x <= 2`. The resulting bounds will be `-1 <= x <= 2`.

Bounds that fix a variable can be entered as simple equalities. For example, x6 = 7.8 is equivalent to 7.8 <= x6 <= 7.8. The bounds $+\infty$ (positive infinity) and $-\infty$ (negative infinity) must be entered as strings (case insensitive):

```
+infinity, -infinity, +inf, -inf.
```

Note that the keywords `infinity` and `inf` may not be used as a right-hand side coefficient of a constraint.

A variable with a negative infinity lower bound and positive infinity upper bound may be entered as `free` (case insensitive). For example, `x9 free` in an LP file bounds section is equivalent to:

```
- infinity <= x9 <= + infinity
```

or

```
- infinity <= x9
```

In the last example here, which uses a single bound is used for `x9` (which is positive infinity for continuous example variable `x9`).

A.3.14 Generals, Integers and binaries

The `generals`, `integers` and `binaries` sections of an LP file is used to indicate the variables that must have integer values in a feasible solution. The difference between the variables registered in each of these sections is in the definition of the default bounds that the variables will have. For variables registered in the `generals` section the default bounds are 0 and `XPRS_PLUSINFINITY`. For variables registered in the `integers` section the default bounds are 0 and 1. The bounds for variables registered in the `binaries` section are 0 and 1.

The lines in the `generals`, `integers` and `binaries` sections are a list of white space or carriage return delimited variable names. Note that variables can not be declared in these sections. That is, a variable appearing in one of these sections that does not appear in the objective section or in a constraint in the constraint section is ignored.

A.3.15 Semi-continuous and semi-integer

The `semi-continuous` and `semi integers` sections of an LP file relate to two similar classes of variables and so their details are documented here simultaneously.

The `semi-continuous` (or `semi integers`) section of an LP file are used to specify variables as semi-continuous (or semi-integer) variables, that is, as variables that may take either (a) value 0 or (b) real (or integer) values from specified thresholds and up to the variables' upper bounds.

The lines in a `semi-continuous` (or `semi integers`) section are a list of white space or carriage return delimited entries that are variable name-number pair. For the `semi-continuous` secion it is also possible to provide a variable name only. The following example shows the format of entries in the `semi-continuous` section.

```
Semi-continuous
x7 >= 2.3
x8
x9 > 4.5
```

The following example shows the format of entries in the `semi integer` section.

```
Semi integers
x7 >= 3
x9 > 5
```

Note that it is possible to use either the >= token or the > token. The resulting threshold will be identical for both cases. It is not possible to use the <= token.

The threshold of the interval within which a variable may have real (or integer) values is defined in two ways depending on whether the entry for the variable is (i) a variable name or (ii) a variable name-number pair. If the entry is just a variable name, then the variable's threshold is the variable's lower bound, defined in the `bounds` section (see earlier). If the entry for a variable is a variable name-number pair, then the variable's threshold is the number value in the pair.

It is important to note that if (a) the threshold of a variable is defined by a variable name-number pair and (b) a lower bound on the variable is defined in the `bounds` section, then:

Case 1) If the lower bound is less then zero, then the lower bound is zero.

Case 2) If the lower bound is greater than zero but less than the threshold, then the value of zero is essentially cut off the domain of the semi-continuous (or semi-integer) variable and the variable becomes a simple bounded continuous (or integer) variable.

Case 3) If the lower bound is greater than the threshold, then the variable becomes a simple lower bounded continuous (or integer) variable.

If no upper bound is defined in the `bounds` section for a semi-continuous (or semi-integer) variable, then the default upper bound that is used is the same as for continuous variables, for semi-continuous variables, and `generals` section variables, for semi-integer variables.

A.3.16 Partial integers

The `partial integers` section of an LP file is used to specify variables as partial integer variables, that is, as variables that can only take integer values from their lower bounds up to specified thresholds and then take continuous values from the specified thresholds up to the variables' upper bounds.

The lines in a `partial integers` section are a list of white space or carriage return delimited variable name-integer pairs. The integer value in the pair is the threshold below which the variable must have integer values and above which the variable can have real values. Note that lower bounds and upper bounds can be defined in the `bounds` section (see earlier). If only one bound is defined in the `bounds` section for a variable or no bounds are defined then the default bounds that are used are the same as for continuous variables.

The following example shows the format of the variable name-integer pairs in the `partial integers` section.

```
Partial integers
x11 >= 8
x12 >= 9
```

Note that you can not use the <= token in place of the >= token.

A.3.17 Special ordered sets

Special ordered sets are defined as part of the `constraints` section of the LP file. The definition of each special ordered set looks the same as a constraint except that the sense is always = and the right hand side is either `S1` or `S2` (case sensitive) depending on whether the set is to be of type 1 or 2, respectively. Special ordered sets of type 1 require that, of the non-negative variables in the set, one at most may be non-zero. Special ordered sets of type 2 require that at most two variables in the set may be non-zero, and if there are two non-zeros, they must be adjacent. Adjacency is defined by the weights, which must be unique within a set given to the variables. The weights are defined as the coefficients on the variables in the set constraint. The sorted weights define the order of the special ordered set. It is perhaps best practice to keep the special order sets definitions together in the LP file to indicate (for your benefit) the start of the special ordered sets definition with the comment line `\Special Ordered Sets` as is done when a problem is written to an LP file by the FICO Xpress Optimizer. The following example shows the definition of a type 1 and type 2 special ordered set.

```
Sos101: 1.2 x1 + 1.3 x2 + 1.4 x4 = S1
Sos201: 1.2 x5 + 1.3 x6 + 1.4 x7 = S2
```

A.3.18 Quadratic programming problems

Quadratic programming problems (QPs) with quadratic objective functions are defined using a special format within the objective function description. The algebraic coefficients of the function `x'Qx` appearing in the objective for QP problems are specified inside square brackets `[]`. All quadratic coefficients must appear inside square brackets. Multiple square bracket sections may be used and quadratic terms in the same variable(s) may appear more than once in quadratic expressions.

Division by two of the QP objective is either implicit, or expressed by a `/2` after the square brackets, thus `[...]` and `[...]/2` are equivalent.

Within a square bracket pair, a quadratic term in two different variables is indicated by the two variable names separated by an asterisk (*). A squared quadratic term is indicated with the variable name followed by a carat (^) and then a 2.

For example, the LP file objective function section:

```
Minimize
obj: x1 + x2 + [ x1^2 + 4 x1 * x2 + 3 x2^2 ] /2
```

Note that if in a solution the variables `x1` and `x2` both have value 1 then value of the objective function is `1 + 1 + (1*1 + 4*1*1 + 3*1*1) / 2 = 2 + (8) / 2 = 6.`

A.3.19 Quadratic Constraints

Quadratic terms in constraints are introduced using the same format and rules as for the quadratic objective, but without the implied or explicit `/2` after the square brackets.

For example:

```
Minimize
obj: x1 + x2
s.t.
x1 + [ x1^2 + 4 x1 * x2 + 3 x2^2 ] <= 10
x1 >= 1
end
```

Please be aware of the differences of the default behaviour of the square brackets in the objective compared to the constraints. For example problem:

```
min y + [ x^2 ]
st.
 x >= 1
 y >= 1
end
```

Has an optimal objective function value of 1.5, while problem:

```
min t
 s.t.
 -t + y + [ x^2 ] <= 0
 x >= 1
 y >= 1
end
```

has an optimum of 2. The user is suggested to use the explicit `/2` in the objective function like:

```
min y + [ x^2 ] / 2
st.
 x >= 1
 y >= 1
end
```

to make sure that the model represents what the modeller meant.

The FICO Xpress Optimizer can solve both convex and non-convex (MI)QCPs. For non-convex problems, some of the restrictions of FICO Xpress Global apply, e.g., callback usability is limited. A quadratically constrained problem is convex, when all quadratic row matrices are positive semi-definite for $\leqslant$ rows and negative semi-definite for $\geqslant$ rows. For convex (MI)QCPs, equations must not contain a quadratic part.

A.3.20 General Constraints

The general constraints section started by the record `General Constraints` specifies min, max, and, or, abs and piecewise linear constraints. Each line defines one such constraint, beginning with a name, followed by a colon, a resultant variable, a sign, a keyword and further variables (or breakpoints for the piecewise linear constraints) in brackets with spaces and commas. The keywords are `MAX` for a maximum-constraint, `MIN` for a minimum-constraint, `AND` for an and-constraint, `OR` for an or-constraint, `ABS` for an absolute-value-constraint and `PWL` for a piecewise linear constraint. For the max- and min-constraints, the resultant is followed by an arbitrary number of further column names or values, and the resultant should be the maximum/minimum of the remaining columns and values. For the and- and or-constraints the resultant is followed by an arbitrary number of further column names, where all the columns (including the resultant) need to be binary, and the resultant will be one if and only if all (and) or at least one (or) of the remaining variables are one. For the abs-constraint, the resultant should be followed by exactly one further column name, and the resultant will take the absolute value of the other column. For the piecewise linear constraints, there needs to be exactly one input variable, followed by a colon and a list of breakpoints. Note that general constraints may only introduce new variables if they are placed immediately after the subject to (or delayed rows, model cuts or pwl) sections. An example for a max-constraint would be

```
myCons: m = MAX ( x , y , 0.0 )
```

An example for a piecewise linear constraint would be

```
myPwl: y = PWL ( x ): (-1,-1) (0,0) (10,20) (10,0) (11,0)
```

defining that `y = f(x)`, where f is a piecewise linear function with value x if x is negative, `y = 2x` if `0 <= x <= 10` and `y = 0` if x is larger than 10.

A.3.21 Constraint ranges

Constraint ranges may be defined in a section beginning with the keyword `ranges`. Each range definition must be on a separate line, and must specify an upper and lower range for a named constraint that has previously been defined in the `constraints` section.

For example:

```
Minimize
x1 + x2
subject to
c1: x + y <= 10
ranges
1 <= c1 <= 10
end
```

A.3.22 Extended naming convention

If the names in the problem do not comply with the LP file format, the optimizer will automatically check if uniqueness and reproducibility of the names could be preserved by prepending `x(` and appending `)` to all names, i.e. the parenthesis inside the original names are always presented in pairs. In these cases, the optimizer will create an LP file with the extended naming convention format. Use control `FORCEOUTPUT` to force the optimizer to write the names in the problem out as they are.

A.3.23 Compatibility to other extensions

The FICO Xpress Optimizer is also able to read (but not write) further sections defined by extensions of the LP format. These include the SOS section, as a different way of defining special orderred sets, and the PWLObj and PWL sections for piecewise linear objective and constraints.

The piecewise linear objective section is started by the `PWLObj` line. It is followed by lines consisting of one variable name and a list of extreme points defining the piecewise linear objective function for this variable. For example the line

```
x: (-1,-1) (0,0) (10,20) (10,0) (11,0)
```

defines that if x is negative, the objective contribution is x. If x is between 0 and 10, then the objective contribution is 2x and if x is larger than 10, then the objective contribution is 0. For each variable, the extreme points should be sorted according to non-decreasing variable value. The piecewise linear functions do not necessarily need to be continuous, in this case two extreme points with identical variable values and different function values can be given and the first one will be used as the lefthand-limit and the second one as the righthand-limit. Note that for the point where the discontinuity appears, both function values can appear in the solution due to tolerances.

The piecewise linear constraint section is started by the `PWL` keyword. Each piecewise linear constraint defines a restriction y = f(x), where f is a piecewise linear function. The lines in the input format consist of a name, the input variable, a preslope, the extreme points and a postslope. The preslope defines the function before the first extreme point and the postslope defines it after the last one. Discontinuities are possible as for the objective function. Note that pwl sections may only introduce new variables if they are placed immediately after the subject to (or delayed rows, model cuts or general constraints) sections. Above example would look as follows, assuming that instead of the objective it now defines the value of a variable y :

```
pwlc1: y = x 1 (0,0) (10,20) (10,0) 0
```

A.4 ASCII Solution Files

Solution information is available from the Optimizer in a number of different file formats depending on the intended use. The `XPRSwritesol` (`WRITESOL`) command produces two files, *problem_name*`.hdr` and *problem_name*`.asc`, whose output has comma separated fields and is primarily intended for input into another program. By contrast, the command `XPRSwriteprtsol` (`WRITEPRTSOL`) produces fixed format output intended to be sent directly to a printer, the file *problem_name*`.prt`. All three of these files are described below.

A.4.1 Solution Header .hdr Files

This file only contains one line of characters comprising header information which may be used for controlling the reading of the `.asc` file (which contains data on each row and column in the problem). The single line is divided into fourteen fields, separated by commas, as follows:

Field	Type	Width	Description
1	string	10	matrix name;
2	integer	4	number of rows in problem;
3	integer	6	number of structural columns in problem;
4	integer	4	sequence number of the objective row;
5	string	3	problem status (see notes below);
6	integer	4	direction of optimization (`0`=none, `1`=min, `2`=max);
7	integer	6	number of iterations taken;
8	integer	4	final number of infeasibilities;
9	real	12	final object function value;
10	real	12	final sum of infeasibilities;
11	string	10	objective row name;
12	string	10	right hand side row name;
13	integer	1	flag: integer solution found (`1`), otherwise `0`;
14	integer	4	matrix version number.

- Character fields contain character strings enclosed in double quotes.
- Integer fields contain right justified decimal digits.
- Fields of type real contain a decimal character representation of a real number, right justified, with six digits to the right of the decimal point.

- The status of the problem (field 5) is a single character as follows:

C	optimization interrupted (like ctrl-c);
O	optimal;
N	infeasible;
S	stability problems;
U	unbounded;
Z	unfinished.

A.4.2 CSV Format Solution .asc Files

The bulk of the solution information is contained in this file. One line of characters is used for each row and column in the problem, starting with the rows, ordered according to input sequence number. Each line contains ten fields, separated by commas, as follows:

Field	Type	Width	Description
1	integer	6	input sequence number of variable;
2	string	10	variable (row or column) name;
3	string	3	variable type (C=column; N, L, G, E for rows);
4	string	4	variable status (LL, **, BS, UL, EQ, SB or ??);
5	real	12	value of activity;
6	real	12	slack activity (rows) or input cost (columns;)
7	real	12	lower bound (−1000000000 if none);
8	real	12	upper bound (1000000000 if none);
9	real	12	dual activity (rows) or reduced cost (columns);
10	real	12	right hand side value (rows) or blank (columns).

- The field Type is as for the `.hdr` file.
- The variable type (field 3) is defined by:
  C structural column;
  N `N` type row;
  L `L` type row;
  G `G` type row;
  E `E` type row;
- The variable status (field 4) is defined by:
  `LL` non-basic at lower bound;
  `**` basic and infeasible;
  BS basic and feasible;
  UL non-basic at upper bound;
  EQ equality row;
  SB variable is super-basic;
  ?? unknown.

A.4.3 Fixed Format Solution (.prt) Files

This file is the output of the `XPRSwriteprtsol` (`WRITEPRTSOL`) command and has the same format as is displayed to the console by `PRINTSOL`. The format of the display is described below by way of an example, for which the simple example of the FICO Xpress Getting Started manual will be used.

The first section contains summary statistics about the solution process and the optimal solution that has been found. It gives the matrix (problem) name (`simple`) and the names of the objective function and right hand sides that have been used. Then follows the number of rows and columns, the fact that it was a maximization problem, that it took two iterations (simplex pivots) to solve and that the best solution has a value of `171.428571`.

```
Problem Statistics
Matrix simple
Objective *OBJ*
RHS *RHS*
Problem has      3 rows and      2 structural columns

Solution Statistics
Maximization performed
Optimal solution found after      3 iterations
Objective function value is     171.428571
```

Next, the *Rows Section* presents the solution for the rows, or constraints, of the problem.

```
Rows Section
Number Row    At    Value     Slack Value Dual Value  RHS
N   1    *OBJ*  BS  171.428571  -171.428571  .000000      .000000
L   2    second UL  200.000000      .000000  .571429   200.000000
L   3    first  UL  400.000000      .000000  .142857   400.000000
```

The first column shows the constraint type: `L` means a 'less than or equal to' constrain; `E` indicates an 'equality' constraint; `G` refers to a 'greater than or equal to' constraint; `N` means a 'nonbinding' constraint – this is the objective function.

The sequence numbers are in the next column, followed by the name of the constraint. The `At` column displays the status of the constraint. A `UL` indicator shows that the row is at its upper limit. In this case a $\leqslant$ row is hard up against the right hand side that is constraining it. `BS` means that the constraint is not active and could be removed from the problem without changing the optimal value. If there were $\geqslant$ constraints then we might see `LL` indicators, meaning that the constraint was at its lower limit. Other possible values include:

**	basic and infeasible;
EQ	equality row;
??	unknown.

The `RHS` column is the right hand side of the original constraint and the `Slack Value` is the amount by which the constraint is away from its right hand side. If we are tight up against a constraint (the status is `UL` or `LL`) then the slack will be `0`.

The `Dual Value` is a measure of how tightly a constraint is acting. If a row is hard up against a $\leqslant$ constraint then it might be expected that a greater profit would result if the constraint were relaxed a little. The dual value gives a precise numerical measure to this intuitive feeling. In general terms, if the right hand side of a $\leqslant$ row is increased by 1 then the profit will increase by the dual value of the row. More specifically, if the right hand side increases by a sufficiently small δ then the profit will increase by δ x dual value, since the dual value is a marginal concept. Dual values are sometimes known as *shadow prices*.

Finally, the *Columns Section* gives the solution for the columns, or variables.

```
Columns Section
Number Column At    Value    Input Cost   Reduced Cost
C   4     a     BS  114.285714    1.000000  .000000
C   5     b     BS   28.571429    2.000000  .000000
```

The first column contains a `C` meaning column (compare with the rows section above). The number is a sequence number. The name of the decision variable is given under the `Column` heading. Under `At` is the status of the column: `BS` means it is away from its lower or upper bound, `LL` means that it is at its lower bound and `UL` means that the column is limited by its upper bound. Other possible values include:

**	basic and infeasible;
EQ	equality row;
SB	variable is super-basic;
??	unknown.

The `Value` column gives the optimal value of the variable. For instance, the best value for the variable `a` is `114.285714` and for variable `b` it is `28.571429`. The `Input Cost` column tells you the coefficient of the variable in the objective function.

The final column in the solution print gives the `Reduced Cost` of the variable, which is always zero for variables that are away from their bounds – in this case, away from zero. For variables which are zero, it may be assumed that the per unit contribution is not high enough to make production viable. The reduced cost shows how much the per unit profitability of a variable would have to increase before it would become worthwhile to produce this product. Alternatively, and this is where the name *reduced cost* comes from, the cost of production would have to fall by this amount before any production could include this without reducing the best profit.

In case there is a basis available, sensitivity analysis values may be included in the output by passing the `s` flag.

The rows section is extended with two new columns, presenting the range of the RHS in which it can be changed before it affects the optimality or feasibiltiy of the current basis.

```
Rows Section
                                                                                    RHS Sensitivity
   Number     Row      At       Value       Slack Value   Dual Value       RHS        Lower Range   Upper Rang
 G       1  third      LL      15.000000        .000000    10.571428     15.000000     10.199999     25.00000
```

For the columns section, objective ranges are added.

```
Columns Section
                                                                     Objective Sensitivity
   Number    Column    At       Value       Input Cost   Reduced Cost  Lower Range   Upper Range
 C      58  c          BS       2.742857       5.849999        .000000      .999998      8.350002
```

A third section presents the sensitivity values for the column bounds.

```
Bound sensitivity
                                  Lower Bound Sensitivity                   Upper Bound Sensitivity
   Number    Column    At    Lower Range        Bound      Upper Range   Lower Range      Bound       Upper Rang
 C      58  d          BS   -1.00000E+20        .000000      2.742857      2.742857     10.000000    1.00000E+2
```

A.4.4 ASCII Solution (.slx) Files

These files provide an easy to read format for storing solutions. An `.slx` file has a header `NAME` containing the name of the matrix the solution belongs to. Each line contains three fields as follows:

Field	Type	Width	Description
1	char	1	variable type;
2	string	variable	name of variable;
3	real	variable	value of activity.

The variable type (field 1) is defined by:

`C` structural column;
`S` LP solution only: slack variables;
`D` LP solution only: dual variables;
`R` LP solution only: reduced costs.

The file is closed by `ENDATA`.

It is possible to store multiple solutions in the same `.slx` file by repeating the `NAME` field following by the additional solution information.

Example

```
NAME solution 1
 C x1 0
 C x2 1
NAME solution 2
 C x1 1
 C x2 0
ENDATA
```

A.5 The Directives (.dir) File

This consists of an unordered sequence of records which specify branching priorities, forced branching directions and pseudo costs, read into the Optimizer using the `XPRSreaddirs` (`READDIRS`) command. By default its name is of the form *problem_name*`.dir`.

Directive file records have the format:

type	*name*	*value*

type is one of:

PR	implying a priority entry (the value gives the priority, which must be an integer between 0 and 1000. Values greater than 1000 are rejected, and real values are rounded down to the next integer. A low value means that the entity is more likely to be selected for branching.)
UP	the entity is to be forced up (value is not used)
DN	the entity is to be forced down (value is not used)
PU	an up pseudo cost entry (the value gives the cost)
PD	a down pseudo cost entry (the value gives the cost)
MC	a model cut entry (value is not used)
DR	a delayed row entry (value is not used)
BR	force the optimizer to branch on the entity even if it is satisfied. If a node solution is MIP feasible, the Optimizer will first branch on any branchable entity flagged with BR before returning the solution.

name is the name of a MIP entity (column or special ordered set), a row (for types `MC` and `DR`), or a mask. A mask may comprise ordinary characters which match the given character; a `?`, which matches any single character; or a `*`, which matches any string or characters. A `*` can only appear at the end of a mask. Note that a whitespace character encountered within the name field will signal the end of the name field and the start of the value field. To match an entity or row whose name contains whitespace characters, convert the name into a mask by replacing each whitespace character with by a `?`.

value is the value to accompany the type.

For example:

PR x1* 2

gives MIP entities (integer variables etc.) whose names start with `x1` a priority of `2`. Note that the use of a mask: a `*` matches all possible strings after the initial `x1`.

A.6 IIS description file in CSV format

This file contains information on a single IIS of an infeasible LP.

Field	Description
Name	Name of a row or column in conflict.
Type	Type of conflicting variable (row or column).
Sense	Sense of conflicting variable: (`LE` or `GE`) to indicate or rows. (`LO` or `UP`) to indicate lower or upper bounds for columns.
Bound	Value associated with the variable, i.e. RHS for rows and bound values for columns.
Dual value	The dual multipliers corresponding to the contradiction deducible from the IIS. Summing up all the rows and columns in the IIS multiplied by these values yields a contradicting constraint. This value is negative for <= rows and upper bounds, and positive for >= rows and lower bounds.
In iso	Indicates if the row or column is in isolation.

Note that each IIS may contain a row or column with only one of its possible senses. This also means that for equality rows and columns with both lower and upper bounds only one side of the restriction may be present. Range constraints in an IIS are converted to greater than or equal constraints.

An IIS often contains other columns than those listed in the IIS. Such columns are free, and have no associated conflicting bounds.

The information contained in these files is the same as returned by the `XPRSgetiisdata` function, or displayed by (`IIS -p`).

A.7 The Matrix Alteration (.alt) File

The Alter File is an ASCII file containing matrix revision statements, read in by use of the `XPRSalter` (`ALTER`) command, and should be named *problem_name*`.alt` by default. Each statement occupies a separate line of the file and the final line is always empty. The statements consist of *identifiers* specifying the object to be altered and *actions* to be applied to the specified object. Typically the identifier may specify just a row, for example `R2`, specifying the second row if that name has been assigned to row 2. If a coefficient is to be altered, the associated variable must also be specified. For example:

```
RRRRRRRR
CCRider
2.087
```

changes the coefficient of `CCRider` in row `RRRRRRRR` to `2.087`. The *action* may be one of the following possibilities.

A.7.1 Changing Upper or Lower Bounds

An upper or lower bound of a column may be altered by specifying the special 'rows' `**LO` and `**UP` for lower and upper bounds respectively.To change the objective coefficient of a column use the string `**OBJ`. For example, to change the lower bound (to `1.234`), upper bound (to `5.678`) and the objective (to `1234.0`) of column `x____0305` would look like:-

```
**LO
x___0305
1.234
**UP
x___0305
5.678
**OBJ
x___0305
1234.0
```

A.7.2 Changing Right Hand Side Coefficients

Right hand side coefficients of a row may be altered by changing values in the 'column' with the name of the right hand side.

A.7.3 Changing Constraint Types

The direction of a constraint may be altered. The row name is given first, followed by an action of `**NTx`, where `x` is one of:

`N`	for the new row type to be unconstrained;
`L`	for the new row type to be 'less than or equal to';
`G`	for the new row type to be 'greater than or equal to';
`E`	for the new row type to be an equality.
`R`	for the new row type to be a range row.

Note that `N` type rows will not be present in the matrix in memory if the control `KEEPNROWS` has been set to zero before `XPRSreadprob` (`READPROB`).

When turning a row to ranged row, a third entry, the range, is expected to be defined following `**NTR`. The rules for changing a row to a ranged row follow that of `XPRSchgrhsrange`.

A.8 The Tuner Method (.xtm) File

The tuner method file is in a straightforward plain text format. For example, when the two controls `MAXTIME` and `THREADS` are set by the user on the current problem and then `XPRStunerwritemethod` is called, the generated xtm file will look similar to the following:

```
FIXED-CONTROLS
  MAXTIME               = 100
  THREADS               = 1
TUNABLE-CONTROLS
  SBEFFORT              = 0.25, 4
  HEURSEARCHEFFORT      = 0.5, 2
  CUTFACTOR             = 0.5, 1, 5
  SCALING               = 0
  PRESOLVE              = 0
  VARSELECTION          = 2, 7
  CUTFREQ               = 2
  SYMMETRY              = 0, 1, 2
  COVERCUTS             = 0, 2
  GOMCUTS               = 0, 2, 10
  TREECOVERCUTS         = 0
  TREEGOMCUTS           = 0
  HEUREMPHASIS          = 0,1
  SBESTIMATE            = 1, 2, 3, 4, 5, 6
  HEURSEARCHROOTCUTFREQ= 2, 5
  HEURSEARCHROOTSELECT = 0, 3, 5
```

```
HEURSEARCHTREESELECT = 0, 3, 5
ROOTPRESOLVE         = 1
PREPROBING           = 3
BRANCHDISJ           = 0
```

The tuner method file consists of a section of fixed controls and a section of tunable controls.

A.8.1 The fixed controls

The fixed controls section starts with `FIXED-CONTROLS`, followed by control setting lines in assignment form. Each control in this section can only be assigned to one value. If the same control appears several times in this section, its first appearance will be used.

When writting out the tuner method file, all the controls set for the current problem will be included in the fixed control section. When reading in the tuner method file using `XPRStunerreadmethod`, these controls won't be applied to the current problem immediately, they will only be applied to the worker problem used in the tuner.

This section can be empty.

A.8.2 The tunable controls

The tunable controls section starts with `TUNABLE-CONTROLS`, followed by control setting lines in assignment form. Each control in this section can be assigned to one value, or multiple values seperated by commas. A control may appear multiple times in this section.

When reading in a tuner method file and writing it out again, the tunable controls may appear in a different order. If there is a control appearing multiple times in the original tuner method file, when written out, it will be combined into a single line with multiple values.

For bit-vector controls (see Section 9.2), such as `PRESOLVEOPS`, `SCALING`, or `HEURSEARCHROOTSELECT`, it is possible to either specify concrete assignments to the control or to specify individual bits that should be used for tuning. In the latter case, one has to use a colon instead of an equals sign. For example, "SCALING = 24, 675" will tune using the two given concrete values 24 and 675, while "SCALING : 3, 4, 9" will individually toggle bits 3, 4, and 9 of the default value for `SCALING` and potentially also try combinations of those later during tuning.

This section can be empty. When both the fixed and the tunable secitons are empty, the tuner will then use a pre-defined factory tuner method.

A.9 The Simplex Log

During the simplex optimization, a summary log is displayed at regular time intervals, determined by an internal deterministic timer. This summary log has the form:

Its	The number of iterations or steps taken so far.
Obj Value	The objective function value.
S	The current solution method (`p` primal; `d` dual).
Ninf	The number of infeasibilities.
Nneg	The number of variables which may improve the current solution if assigned a value away from their current bounds.
Sum inf	The scaled sum of infeasibilities. For the dual algorithm this is the scaled sum of dual infeasibilities when the number of negative dj's is non-zero.
Time	The number of seconds spent iterating.

A more detailed log can be displayed every `n` iterations by setting `LPLOGSTYLE` to `0` and `LPLOG` to `-n`. The detailed log has the form:

Its	The number of iterations or steps taken so far.
S	The current solution method (`p` primal; `d` dual).
Ninf	The number of infeasibilities.
Obj Value	If the solution is infeasible, the scaled sum of infeasibilities, otherwise: the objective value.
In	The sequence number of the variable entering the basis (negative if from upper bound).
Out	The sequence number of the variable leaving the basis (negative if to upper bound).
Nneg	The number of variables which may prove the current solution if assigned a value away from their current bounds.
Dj	The scaled rate at which the most promising variable would improve the solution if assigned a value away from its current bound (reduced cost).
Neta	A measure of the size of the inverse.
Nelem	Another measure of the size of the inverse.
Time	The number of seconds spent iterating.

If `LPLOG` is set to 0, no log is displayed until the optimization finishes.

A.10 The Barrier Log

The first line of the barrier log displays statistics about the Cholesky decomposition needed by the barrier algorithm. This line contains the following values:

Dense cols	The number of dense columns identified in the factorization.
NZ(L)	The number of nonzero elements in the Cholesky factorization.
Flops	The number of floating point operations needed to perform one factorization.

During the barrier optimization, a summary log is displayed in every iteration. This summary log has the form:

Its	The number of iterations taken so far.
P.inf	Maximal violation of primal constraints.
D.inf	Maximal violation of dual constraints.
U.inf	Maximal violation of upper bounds.
Primal obj	Value of the primal objective function.
Dual obj	Value of the dual objective function.
Compl	Value of the average complementarity.

After the barrier algorithm a crossover procedure may be applied. This process prints at most 3 log lines about the different phases of the crossover procedure. The structure of these lines follows The Simplex Log described in the section above.

If `BAROUTPUT` is set to 0, no log is displayed until the barrier algorithm finishes.

A.11 The Hybrid Gradient Log

The content and format of the log messages of the Hybrid Gradient method are identical to those of the

Newton Barrier method, with a few minor differences. Since there is no Cholesky factorization that table is not printed. Upper bound infeasibility is always 0, simply by the design of the algorithm.

If BAROUTPUT is set to 0, no log is displayed until the hybrid gradient algorithm finishes.

A.12 The MIP Log

During the branch and bound tree search (see XPRSmipoptimize (MIPOPTIMIZE)), a summary log of nine columns of information is frequently printed. By default, the printing frequency increases over time. If MIPLOG is explicitly set to a negative value −n, a log line will be printed every *n* nodes. The nine columns consist of:

Node	A sequential node number.
BestSoln	The value of the best integer solution found so far.
BestBound	A bound on the value of the optimal integer solution.
Sols	The number of integer solutions that have been found.
Active	The number of active (unsolved) nodes in the branch and bound tree search.
Depth	The depth of the current node in the branch and bound tree.
Gap	The percentage gap between the best solution and the best bound.
GInf	The number of MIP infeasibilities at the current node.
Time	The time taken.

This log is also printed when an integer feasible solution is found. An asterisk (*) printed in front of the node number indicates that a solution has been found by an integral LP relaxation. Single characters indicate that a heuristic solution has been found. Lower case characters stand for different strategies of the Optimizer's diving heuristic: the letter a corresponds to strategy 1, the letter b to strategy 2, and so forth. Compare control HEURDIVESTRATEGY. By default, several strategies are applied. Upper case letters stand for special search heuristics. More precisely, R, L, M, C, U, and Z stand for the different modes of local search that can be selected by controls HEURSEARCHROOTSELECT and HEURSEARCHTREESELECT. For technical reasons, a U might also appear after a restart. The letter F represents the feasibility pump, T stands for a trivial heuristic. S, G, and B are reserved for special purpose heuristics for problems with set packing/partitioning constraints, GUBs, and branching on constraints, respectively. An E indicates that a solution has been found during the calculation of branching estimates. P stands for heuristics that run in parallel to the MIP root node. H stands for the Shift-and-propagate heuristic. When solving with Xpress Global, an N indicates that an integer feasible node solution has been extended to a MINLP solution using a local solver.

Note that heuristic solutions found before the tree search (and thereby the MIP log) starts, are reported with a special log line of the form:

```
*** Solution found:     <SolVal>   Time:   <Time>    Heuristic: <DispChar> ***
```

Here <SolVal> and <Time> are numeric values indicating the solution value of the found heuristic solution and the run time after which it was found, respectively. The single <DispChar> character indicates which heuristic found the solution.

During root node cutting, the column Node is replaced by two columns Its and Type, columns Active and Depth are replaced by Add and Del, respectively. These have the following meaning:

Its	A counter for the number of cutting plane separation loops.
Type	The type of cuts that have been generated this round: `G` – Gomory cuts, `M` – model cuts, `O` – outer approximation cuts (only for nonlinear problems), `N` – network-based cuts, `K` – any other type of cuts
Add	Number of cuts added to the LP relaxation in this iteration
Del	Number of cuts deleted from the LP relaxation in this iteration

If `MIPLOG` is set to `3`, a detailed log of eight columns of search information is printed for each node:

Branch	A sequential node number.
Parent	The node number of the parent of the current node. A `D` or a `U` marks whether the current node is the down child or the up child, respectively, of its parent.
Solution	The optimal value of the LP relaxation at the current node.
Entity	The MIP entity on which the Optimizer will branch after this node.
Value	The current value of the entity chosen for branching.
Active	The number of active nodes in the tree search.
GInf	The number of violated MIP entities
Time	The time taken.

Not all the information described above is present for all nodes. If the LP relaxation is cut off, only Branch and Parent (and possibly Solution) are displayed. If the LP relaxation is infeasible, only Branch and Parent appear. The rest of the line will consist of a text message `relaxation exceeds cutoff` or `relaxation infeasible`. If an integer solution is discovered, this is highlighted before the log line is printed.

If `MIPLOG` is set to `2`, the detailed log is printed at integer feasible solutions only. When `MIPLOG` is set to `1`, the tree node logs are surpressed, but cutting loop logs will still be displayed. If `MIPLOG` is set to `0`, neither cut nor node log wil be pritned. In any case, LP logs and intermediate status messages might still be printed.

A.13 The IIS Log

While the IIS procedure is searching for a minimal set of restrictions that cause the infeasibility of the instance it prints a summary log periodically. The IIS log consists of three columns of information. This is repeated for each restriction type (rows, columns, discrete entities). The restriction type is printed before the log.

Remain-ing	The number of restrictions of the given type that still need to be checked by the deletion filter.
In IIS	The number of restrictions of the given type that are verified to be necessary for the IIS.
Time (s)	The total time spent in the IIS procedure (in seconds).

The printing of the IIS log can be disabled using the `IISLOG` control. A similar log is printed while searching for isolations.

A.14 The Tuner Log

While the tuner evaluates various control settings, it prints a summary log for each finished run. When tuning a MIP problem, the summary log consists of eight columns of information:

RunID	A sequential tuner run number.
Stat	Status of a finished run: S - Solved, T - Timeout, U - Unsolved and C - Cancelled.
Solution	The best integer solution.
Bound	The best bound.
Integral	The primal dual integral.
Gap	The relative MIP gap.
RunTime	The time spent for solving with this control setting.
TotTime	The total time spent for the tuner.

When tuning an LP problem, the summary log consists of five columns of information:

RunID	A sequential tuner run number.
Stat	Status of a finished run: S - Solved, T - Timeout, U - Unsolved and C - Cancelled.
Solution	The LP objective.
RunTime	The time spent for solving with this control setting.
TotTime	The total time spent for the tuner.

When the tuner finds an improving control setting, it will highlight the run with an asterisk (*) at the beginning of the log line. The tuner will also specify the control parameters and the log file name for the improving run.

If a control setting has been evaluated in previous tuner runs, its result can be reused. In this case, the tuner will print an extra `H` in the `Stat` column.

A.15 The Remote Solving Configuration file

This configuration file allows the user to control some of the ways the Xpress solver interacts with the remote Insight Compute Interface. It contains advanced configuration settings; it is expected that most users will not need to use these configuration options.

To use the configuration file, set the environment variable `XPRESS_COMPUTE_CONFIG` to the full path of the file including the file name itself; you must also set the `XPRESS_COMPUTE` and `XPRESS_COMPUTE_URL` environment variables to activate remote solving in the usual way. Changes to the configuration are only read when the Optimizer is first initialized with XPRSinit.

The configuration file must be a valid JSON document, containing a single object with key-value pairs. All keys are optional and Xpress will use sensible defaults for anything you do not specify. For example:

```
{
  "logLevel": 101,
  "caCertsPath": "C:/xpressmp/ssl/ca-bundle.crt"
}
```

The remainder of this section details the individual keys that can be set.

A.15.1 caCertsPath

This field can be set to the absolute path of the certificate bundle file to use for authenticating SSL certificates when communicating with a remote server using HTTPS. If unspecified, Xpress will look for a file `ca-bundle.crt` in the path specified by the `MOSEL_SSL` environment variable (if set), or the `.mmssl` folder of the user's home directory (if not). If this file does not exist, it can be created in the default location by executing the command `mmssl setup`

For example:

```
{
  "caCertsPath": "C:/xpressmp/ssl/ca-bundle.crt"
}
```

A.15.2 cleanupJobs

This field controls whether Xpress should delete a compute job from the remote server when it successfully completes or is interrupted. The default is `false`. Set to `true` to automatically delete successfully completed jobs from the Insight Compute Interface.

For example:

```
{
  "cleanupJobs": true
}
```

A.15.3 executionService

This field contains the name of the Insight execution service that will be used for jobs. If unset, the Insight server's default execution service is used. Where the execution service name is set in the configuration file and the `COMPUTEEXECSERVICE` control, the value from the control will be used.

For example:

```
{
  "executionService": "SECONDARY"
}
```

A.15.4 logLevel

This field controls additional lines written to the problem's log that describe communication between the local Xpress application and the remote Insight Compute Interface. (Note it doesn't affect lines written by the remote Insight Compute Interface itself.) Supported levels are:

- `0` - write no extra log lines
- `1` - write error and warning messages only
- `2` - write error, warning and notification messages (the default)
- `101` - as `2`, but also write lines for every HTTP request made and event message received from the server
- `102` - as `101`, but include extra debugging output. Should be set on the advice of FICO product support only.

For example:

```
{
  "logLevel": 0
}
```

A.15.5 maxRetries

This field controls how many times a failed request to the remote server will be retried before we show an error to the user. There will be an exponentially increasing delay before each retry (0ms, 200ms, 400ms, 800ms, etc) - the default setting of `9` means we try each request for about 51 seconds before the user is informed of an error. Set to a lower value if you want to see errors quicker, or `0` to disable retries entirely.

For example:

```
{
  "maxRetries": 2
}
```

A.15.6 priority

This field controls the default priority of compute jobs created on the Xpress Insight Compute Interface. It should be a number between -100 and 100, with priority 0 used if not specified. If the priority is set to a nonzero value in both the configuration file and the `COMPUTEJOBPRIORITY` control, the value from the control will be used.

For example:

```
{
  "priority": 90
}
```

A.15.7 trustSrv

This field controls whether Xpress should trust the remote server without checking its certificate, when an HTTPS URL is used. If set to `false`, the authenticity of the remote server is checked using the list of trusted certification authorities and the operation will be aborted if the verification fails. Set to `true` if you want to use a server that has a known invalid or self-signed certificate, and acknowledge the security risks this brings. Default is `false`.

For example:

```
{
  "trustSrv": true
}
```

APPENDIX B

Contacting FICO

FICO provides clients with support and services for all our products.

FICO Customer Support

FICO Customer Support offers technical support and services ranging from self-help tools to direct assistance with a FICO technical support engineer. Support is available to all clients who have an active maintenance contract.

The FICO Customer Self-Service Portal (support.fico.com) is a secure web portal that allows users to open, review, and update their support cases; manage their organization's portal users; find solutions to common problems in the FICO Knowledge Base; and view the availability of their cloud applications 24 hours a day, 7 days a week.

You can find support contact information and a link to the FICO Customer Self-Service Portal (online support) on the Product Support home page (www.fico.com/en/product-support).

Please include *'Xpress'* in the subject line of your support queries.

FICO Community

The FICO Community is a great resource to find the experts and information you need to collaborate, support your business, and solve common business challenges. You can get informal technical support, build relationships with local and remote professionals, and improve your business practices. For additional information, visit the FICO Community (community.fico.com/welcome).

Documentation

FICO continually looks for new ways to improve and enhance the value of the products and services we provide.

If you have comments or suggestions regarding how we can improve this documentation, let us know by sending your suggestions to techpubs@fico.com. Please include your contact information (name, company, email address, and optionally, your phone number) so we may reach you if we have questions.

FICO Learning

FICO Learning is the principal provider of product training for our clients and partners. FICO Learning offers instructor-led classroom courses, web-based training, seminars, and training tools for both new user enablement and ongoing performance support.

For additional information, visit the FICO Learning home page at www.fico.com/en/product-training or email producteducation@fico.com.

Sales and maintenance

If you need information on other Xpress Optimization products, or you need to discuss maintenance contracts or other sales-related items, contact FICO by:

- Phone: +1 (408) 535-1500 or +44 207 940 8718
- Web: www.fico.com/optimization and use the available contact forms

About FICO

FICO (NYSE:FICO) is a leading analytics software company, helping businesses in 90+ countries make better decisions that drive higher levels of growth, profitability, and customer satisfaction. Learn more at www.fico.com or contact us at www.fico.com/en/contact-us.

Index

A

B