

FICO® Xpress Nonlinear

45.01

REFERENCE MANUAL

FICO® Xpress Optimization



©1983–2025 Fair Isaac Corporation. All rights reserved. This documentation is the property of Fair Isaac Corporation ("FICO"). Receipt or possession of this documentation does not convey rights to disclose, reproduce, make derivative works, use, or allow others to use it except solely for internal evaluation purposes to determine whether to purchase a license to the software described in this documentation, or as otherwise set forth in a written software license agreement between you and FICO (or a FICO affiliate). Use of this documentation and the software described in it must conform strictly to the foregoing permitted uses, and no other use is permitted.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, please report them to us in writing. Neither FICO nor its affiliates warrant that this documentation is error-free, nor are there any other warranties with respect to the documentation except as may be provided in the license agreement. FICO and its affiliates specifically disclaim any warranties, express or implied, including, but not limited to, non-infringement, merchantability and fitness for a particular purpose. Portions of this documentation and the software described in it may contain copyright of various authors and may be licensed under certain third-party licenses identified in the software, documentation, or both.

In no event shall FICO or its affiliates be liable to any person for direct, indirect, special, incidental, or consequential damages, including lost profits, arising out of the use of this documentation or the software described in it, even if FICO or its affiliates have been advised of the possibility of such damage. FICO and its affiliates have no obligation to provide maintenance, support, updates, enhancements, or modifications except as required to licensed users under a license agreement.

FICO is a registered trademark of Fair Isaac Corporation in the United States and may be a registered trademark of Fair Isaac Corporation in other countries. Other product and company names herein may be trademarks of their respective owners.

Patent(s): www.fico.com/en/patents

Xpress Nonlinear 45.01 (FICO® Xpress 9.7)

Deliverable Version: A

Last Revised: 29 July, 2025

Contents

I	Overview	1
1	Introduction	2
1.1	Mathematical programs	3
1.1.1	Linear programs	3
1.1.2	Convex quadratic programs	3
1.1.3	Convex quadratically constrained quadratic programs	3
1.1.4	Second order conic problems	3
1.1.5	General nonlinear optimization problems	4
1.1.6	Mixed integer programs	4
1.2	Technology Overview	4
1.2.1	The Simplex Method	4
1.2.2	The Logarithmic Barrier Method	5
1.2.3	Outer approximation schemes	5
1.2.4	Successive Linear Programming	5
1.2.5	Second Order Methods	5
1.2.6	Mixed Integer Solvers	5
1.3	API naming convention	6
2	An example problem	7
2.1	Problem Definition	7
2.2	Problem Formulation	7
3	Modeling in Mosel	9
3.1	Basic formulation	9
3.2	Setting up and solving the problem	10
3.3	Looking at the results	11
3.4	User functions	11
3.5	Parallel evaluation of Mosel user functions	12
4	The Xpress NonLinear API Functions	14
4.1	Header files	14
4.2	Initialization	14
4.3	Callbacks	14
4.4	Creating the linear part of the problem	15
4.5	Adding the non-linear part of the problem	17
4.6	Adding the non-linear part of the problem using character formulae	20
4.7	Checking the data	21
4.8	Solving and printing the solution	21
4.9	Closing the program	21
4.10	Adding initial values	22

5	The Nonlinear Console Program	23
5.1	The Console Nonlinear	23
5.1.1	The nonlinear console extensions	23
5.1.2	Common features of the Xpress Optimizer and the Xpress Nonlinear Optimizer console	24
II	Advanced	26
6	Nonlinear Problems	27
6.1	Coefficients and formulas	27
6.2	SLP variables	28
6.3	Local and global optimality	28
6.4	Convexity	28
6.5	Converged and practical solutions	29
6.6	The duals of general, nonlinear program	29
7	Extended MPS file format	31
7.1	Formulae	31
7.2	COLUMNS	32
7.3	BOUNDS	33
7.4	SLPDATA	33
7.4.1	DR (Determining row)	33
7.4.2	EC (Enforced constraint)	34
7.4.3	FR (Free variable)	34
7.4.4	FX (Fixed variable)	34
7.4.5	IV (Initial value)	34
7.4.6	LO (Lower bounded variable)	35
7.4.7	Rx, Tx (Relative and absolute convergence tolerances)	35
7.4.8	SB (Initial step bound)	35
7.4.9	UF (User function)	36
7.4.10	UP (Free variable)	36
7.4.11	WT (Explicit row weight)	36
7.4.12	DL (variable specific Determining row cascade iteration Limit)	36
8	Xpress-SLP Solution Process	38
8.1	Analyzing the solution process	39
8.2	The initial point	39
8.3	Derivatives	39
8.3.1	Finite Differences	39
8.3.2	Symbolic Differentiation	40
8.3.3	Automatic Differentiation	40
8.4	Points of inflection	40
8.5	Trust regions	41
9	Handling Infeasibilities	42
9.1	Infeasibility Analysis in the Xpress Optimizer	42
9.2	Managing Infeasibility with Xpress Knitro	42
9.3	Managing Infeasibility with Xpress-SLP	43
9.4	Penalty Infeasibility Breakers in XSLP	43

10 Cascading	45
10.1 Determining rows and determining columns	46
11 Convergence criteria	47
11.1 Convergence criteria	47
11.2 Convergence overview	47
11.2.1 Strict Convergence	47
11.2.2 Extended Convergence	47
11.2.3 Stopping Criterion	48
11.2.4 Step Bounding	49
11.3 Convergence: technical details	49
11.3.1 Closure tolerance (CTOL)	52
11.3.2 Delta tolerance (ATOL)	52
11.3.3 Matrix tolerance (MTOL)	52
11.3.4 Impact tolerance (ITOL)	53
11.3.5 Slack impact tolerance (STOL)	53
11.3.6 Fixed variables due to determining columns smaller than threshold (FX)	54
11.3.7 User-defined convergence	54
11.3.8 Static objective function (1) tolerance (VTOL)	54
11.3.9 Static objective function (2) tolerance (OTOL)	55
11.3.10 Static objective function (3) tolerance (XTOL)	55
11.3.11 Extended convergence continuation tolerance (WTOL)	56
12 Xpress-SLP Structures	57
12.1 SLP Matrix Structures	57
12.1.1 Augmentation of a nonlinear coefficient	57
12.1.2 Augmentation of a nonlinear term	59
12.1.3 Augmentation of a user-defined SLP variable	60
12.1.4 SLP penalty error vectors	61
12.2 Xpress-SLP Matrix Name Generation	61
12.3 Xpress-SLP Statistics	62
12.4 SLP Variable History	64
13 Xpress NonLinear Formulae	65
13.1 Parsed and unparsed formulae	65
13.2 Example of an arithmetic formula	66
13.3 Example of a formula involving a simple function	67
14 User Functions	68
14.1 Callbacks and user functions	68
14.2 User function interface	68
14.3 User Function declaration in native languages	69
14.3.1 User function declaration in C	69
14.4 Programming Techniques for User Functions	70
14.4.1 Deltas	70
14.4.2 Return values and ReturnArray	70
14.4.3 Returning Derivatives	70
14.4.4 Function Instances	71
15 Management of zero placeholder entries	73

15.1	The augmented matrix structure	73
15.2	Derivatives and zero derivatives	73
15.3	Placeholder management	74
16	Special Types of Problem	76
16.1	Nonlinear objectives	76
16.2	Convex Quadratic Programming	76
16.3	Mixed Integer Nonlinear Programming	77
16.3.1	Mixed Integer SLP	77
16.3.2	Heuristics for Mixed Integer SLP	77
16.3.3	Fixing or relaxing the values of the SLP variables	78
16.3.4	Iterating at each node	79
16.3.5	Termination criteria at each node	79
16.3.6	Callbacks	80
16.4	Integer and semi-continuous delta variables	80
17	Xpress NonLinear multistart	82
III	Reference	83
18	Reference Documentation by Topic/Functionality	84
18.1	Bit-vector	84
18.1.1	Bit-vector library functions	84
18.1.2	Bit-vector controls	84
18.2	Branching	85
18.2.1	Branching controls	85
18.3	Callback	85
18.3.1	Callback library functions	86
18.4	Cascading	87
18.4.1	Cascading library functions	87
18.4.2	Cascading controls	88
18.5	Controls and Attributes	88
18.5.1	Controls and Attributes library functions	89
18.6	Cuts	89
18.6.1	Cuts controls	89
18.7	Data Input	90
18.7.1	Data Input library functions	90
18.7.2	Data Input controls	90
18.8	Data Information	90
18.8.1	Data Information library functions	90
18.8.2	Data Information attributes	91
18.9	Derivatives	91
18.9.1	Derivatives controls	91
18.9.2	Derivatives attributes	92
18.10	File IO	92
18.10.1	File IO library functions	92
18.10.2	File IO controls	92
18.11	Heuristics	92
18.11.1	Heuristics controls	93
18.12	Knitro	93
18.12.1	Knitro controls	93
18.13	Knitro-MINLP	96

18.13.1 Knitro-MINLP controls	96
18.14 Licensing	97
18.14.1 Licensing library functions	98
18.15 Limits	98
18.15.1 Limits controls	98
18.16 Linearizations	100
18.16.1 Linearizations controls	100
18.16.2 Linearizations attributes	101
18.17 Logging	101
18.17.1 Logging library functions	101
18.17.2 Logging controls	101
18.18 Memory	103
18.18.1 Memory controls	104
18.19 MISLP	104
18.19.1 MISLP library functions	104
18.19.2 MISLP controls	104
18.19.3 MISLP attributes	105
18.20 Misc	105
18.20.1 Misc library functions	106
18.20.2 Misc controls	106
18.20.3 Misc attributes	106
18.21 Multistart	106
18.21.1 Multistart library functions	106
18.21.2 Multistart controls	107
18.21.3 Multistart attributes	107
18.22 Names Manager	107
18.22.1 Names Manager library functions	108
18.23 Numerics	108
18.23.1 Numerics library functions	108
18.23.2 Numerics controls	108
18.23.3 Numerics attributes	108
18.24 Parallel	108
18.24.1 Parallel controls	109
18.25 Presolve	109
18.25.1 Presolve library functions	109
18.25.2 Presolve controls	109
18.25.3 Presolve attributes	110
18.26 Problem Creation	110
18.26.1 Problem Creation library functions	110
18.27 Problem Information	111
18.27.1 Problem Information library functions	111
18.27.2 Problem Information attributes	111
18.28 Problem Modification	113
18.28.1 Problem Modification library functions	113
18.29 Save Restore	113
18.29.1 Save Restore library functions	113
18.30 SLP	114
18.30.1 SLP library functions	114
18.30.2 SLP controls	116
18.30.3 SLP attributes	124
18.31 SLP-convergence	126
18.31.1 SLP-convergence library functions	126
18.31.2 SLP-convergence controls	126
18.31.3 SLP-convergence attributes	128
18.32 Solution Process	128

18.32.1 Solution Process library functions	128
18.32.2 Solution Process controls	129
18.32.3 Solution Process attributes	129
18.33 Solution	130
18.33.1 Solution library functions	130
18.33.2 Solution controls	131
18.33.3 Solution attributes	131
18.34 Tolerances	132
18.34.1 Tolerances controls	132
18.35 User Functions	136
18.35.1 User Functions library functions	136
18.35.2 User Functions controls	136
18.35.3 User Functions attributes	136

19 Problem Attributes 137

19.1 Double problem attributes	142
XSLP_CURRENTDELTACOST, SLPCURRENTDELTACOST	142
XSLP_CURRENTERRORCOST, SLPCURRENTERRORCOST	142
XSLP_ERRORCOSTS, SLPEXPLORCOSTS	142
XSLP_OPTTIME, NLP OPTTIME	142
XSLP_OBJVAL, NLP OBJVAL	143
XSLP_PENALTYDELTATOTAL, SLPPENALTYDELTATOTAL	143
XSLP_PENALTYDELTAVALUE, SLPPENALTYDELTAVALUE	143
XSLP_PENALTYERRORTOTAL, SLPPENALTYERRORTOTAL	143
XSLP_PENALTYERRORVALUE, SLPPENALTYERRORVALUE	143
XSLP_PRIMALINTEGRAL, NLP PRIMALINTEGRAL	144
XSLP_VALIDATIONINDEX_A, NLP VALIDATIONINDEX_A	144
XSLP_VALIDATIONINDEX_K, NLP VALIDATIONINDEX_K	144
XSLP_VALIDATIONINDEX_R, NLP VALIDATIONINDEX_R	144
XSLP_VALIDATIONNETOBJ, NLP VALIDATIONNETOBJ	144
XSLP_VSOLINDEX	145
19.2 Integer problem attributes	146
XSLP_COEFFICIENTS, SLPCOEFFICIENTS	146
XSLP_DELTAS, SLPDELTAS	146
XSLP_ECFCOUNT, SLPECFCOUNT	146
XSLP_EXPLOREDELTAS, SLPEXPLOREDELTAS	146
XSLP_EQUALS COLUMN, NLP EQUALS COLUMN	147
XSLP_IFS, NLP IFS	147
XSLP_IMPLICITVARIABLES, NLP IMPLICITVARIABLES	147
XSLP_INTEGERDELTAS, SLP INTEGERDELTAS	147
XSLP_ITER, SLP ITER	148
XSLP_JOBID, NLP JOBID	148
XSLP_KEEPESTITER, NLP KEEPESTITER	148
XSLP_MINUSPENALTYERRORS, SLP MINUSPENALTYERRORS	148
XSLP_MIPITER, SLP MIPITER	149
XSLP_MIPNODES, SLP MIPNODES	149
XSLP_MIPSOLS, SLP MIPSOLS	149
XSLP_MODEL COLS, NLP MODEL COLS	149
XSLP_MODEL ROWS, NLP MODEL ROWS	150
XSLP_MSSTATUS	150
XSLP_NLP STATUS, NLP STATUS	150
XSLP_NONCONSTANTCOEFFS, SLP NONCONSTANTCOEFFS	151
XSLP_NONLINEARCONSTRAINTS, NLP NONLINEARCONSTRAINTS	151
XSLP_ORIGINAL COLS, NLP ORIGINAL COLS	151

XSLP_ORIGINALROWS, NLPORIGINALROWS	151
XSLP_PENALTYDELTACOLUMN, SLPPENALTYDELTACOLUMN	152
XSLP_PENALTYDELTAROW, SLPPENALTYDELTAROW	152
XSLP_PENALTYDELTAS, SLPPENALTYDELTAS	152
XSLP_PENALTYERRORCOLUMN, SLPPENALTYERRORCOLUMN	152
XSLP_PENALTYERRORROW, SLPPENALTYERRORROW	153
XSLP_PENALTYERRORS, SLPPENALTYERRORS	153
XSLP_PLUSPENALTYERRORS, SLPPPLUSPENALTYERRORS	153
XSLP_PRESOLVEELIMINATIONS, NLPRESOLVEELIMINATIONS	153
XSLP_PRESOLVSTATE	154
XSLP_SBXCONVERGED, SLPSBXCONVERGED	154
XSLP_SEMICONTDeltas, SLPSEMICONTDeltas	154
XSLP_SOLVERSELECTED, LOCALSOLVERSELECTED	154
XSLP_SOLSTATUS, NLPSTATUS	155
XSLP_STATUS, SLPSTATUS	155
XSLP_STOPSTATUS, NLPSTOPSTATUS	157
XSLP_TOLSETS, SLPTOLSETS	157
XSLP_TOTALEVALUATIONERRORS, NLPTOTALEVALUATIONERRORS	157
XSLP_UCCONSTRAINEDCOUNT, SLPUCCONSTRAINEDCOUNT	158
XSLP_UFINSTANCES	158
XSLP_UFS, NLPUS	158
XSLP_UNCONVERGED, SLPUNCONVERGED	158
XSLP_USEDERIVATIVES, NLPUSEDERIVATIVES	159
XSLP_USERFUNCCALLS, NLPUSERFUNCCALLS	159
XSLP_VALIDATIONSTATUS, NLPVALIDATIONSTATUS	159
XSLP_VARIABLES, NLPVARIABLES	159
XSLP_ZEROESRESET, SLPZEROESRESET	160
XSLP_ZEROESRETAINED, SLPZEROESRETAINED	160
XSLP_ZEROESTOTAL, SLPZEROESTOTAL	160
19.3 Reference (pointer) problem attributes	161
XSLP_MIPPROBLEM	161
XSLP_XPRSPROBLEM	161
XSLP_XSLPPROBLEM	161
19.4 String problem attributes	162
XSLP_VERSIONDATE	162

20 Control Parameters

163

20.1 Double control parameters	176
XSLP_ATOL_A, SLPATOL_A	176
XSLP_ATOL_R, SLPATOL_R	176
XSLP_BARSTALLINGTOL, SLPBARSTALLINGTOL	176
XSLP_BOUNDTHRESHOLD, SLPBOUNDTHRESHOLD	177
XSLP_CASCADETOL_PA, SLPCASCADETOL_PA	177
XSLP_CASCADETOL_PR, SLPCASCADETOL_PR	177
XSLP_CDTOL_A, SLPCTOL_A	178
XSLP_CDTOL_R, SLPCTOL_R	178
XSLP_CLAMP SHRINK, SLPCLAMP SHRINK	179
XSLP_CLAMPVALIDATIONTOL_A, SLPCLAMPVALIDATIONTOL_A	179
XSLP_CLAMPVALIDATIONTOL_R, SLPCLAMPVALIDATIONTOL_R	179
XSLP_CTOL, SLPCTOL	180
XSLP_DAMP, SLPDAMP	180
XSLP_DAMPEXPAND, SLPDAMPEXPAND	180
XSLP_DAMP MAX, SLPDAMP MAX	181
XSLP_DAMP MIN, SLPDAMP MIN	181

XSLP_DAMPSHRINK, SLPDAMPSHRINK	182
XSLP_DEFAULTIV, NLPDEFAULTIV	182
XSLP_DEFAULTSTEPBOUND, SLPDEFAULTSTEPBOUND	182
XSLP_DELTA_A, SLPDELTA_A	183
XSLP_DELTA_INFINITY, SLPDELTA_INFINITY	183
XSLP_DELTA_R, SLPDELTA_R	183
XSLP_DELTA_X, SLPDELTA_X	184
XSLP_DELTA_Z, SLPDELTA_Z	184
XSLP_DELTA_ZERO, SLPDELTA_ZERO	184
XSLP_DELTACOST, SLPDELTACOST	185
XSLP_DELTACOSTFACTOR, SLPDELTACOSTFACTOR	185
XSLP_DELTAMAXCOST, SLPDELTAMAXCOST	185
XSLP_DJTOL, SLPDJTOL	186
XSLP_DRCOLDJTOL, SLPDRCOLDJTOL	186
XSLP_DRCOLTOL, SLPDRCOLTOL	186
XSLP_DRFIXRANGE, SLPDRFIXRANGE	187
XSLP_ECFTOL_A, SLPECFTOL_A	187
XSLP_ECFTOL_R, SLPECFTOL_R	187
XSLP_ENFORCECOSTSHRINK, SLPENFORCECOSTSHRINK	188
XSLP_ENFORCEMAXCOST, SLPENFORCEMAXCOST	188
XSLP_ERRORCOST, SLPERRORCOST	189
XSLP_ERRORCOSTFACTOR, SLPERRORCOSTFACTOR	189
XSLP_ERRORMAXCOST, SLPERRORMAXCOST	189
XSLP_ERRORTOL_A, SLPERRORTOL_A	190
XSLP_ERRORTOL_P, SLPERRORTOL_P	190
XSLP_ESCALATION, SLPESCALATION	190
XSLP_ETOL_A, SLPETOL_A	191
XSLP_ETOL_R, SLPETOL_R	191
XSLP_EVTOL_A, SLPEVTOL_A	191
XSLP_EVTOL_R, SLPEVTOL_R	192
XSLP_EXPAND, SLPEXPAND	192
XSLP_FEASTOLTARGET, SLPFEASTOLTARGET	193
XSLP_GRANULARITY, SLPGRANULARITY	193
XSLP_INFINITY, NLPINFINITY	193
XSLP_ITOL_A, SLPITOL_A	194
XSLP_ITOL_R, SLPITOL_R	194
XSLP_MATRIXTOL, SLPMATRIXTOL	195
XSLP_MAXWEIGHT, SLPMAXWEIGHT	195
XSLP_MEMORYFACTOR	196
XSLP_MERITLAMBDA, NLPMERITLAMBDA	196
XSLP_MINSBFACTOR, SLPMINSBFACTOR	197
XSLP_MINWEIGHT, SLPMINWEIGHT	197
XSLP_MIPCUTOFF_A, SLPMIPCUTOFF_A	197
XSLP_MIPCUTOFF_R, SLPMIPCUTOFF_R	198
XSLP_MIPERRORTOL_A, SLPMIPERRORTOL_A	198
XSLP_MIPERRORTOL_R, SLPMIPERRORTOL_R	199
XSLP_MIPOTOL_A, SLPMIPOTOL_A	199
XSLP_MIPOTOL_R, SLPMIPOTOL_R	199
XSLP_MSMAXBOUNDRANGE, MSMAXBOUNDRANGE	200
XSLP_MTOL_A, SLPMTOL_A	200
XSLP_MTOL_R, SLPMTOL_R	201
XSLP_MVTOL, SLPMVTOL	202
XSLP_OBJSENSE	202
XSLP_OBJTHRESHOLD, SLPOBJTHRESHOLD	202
XSLP_OBJTOPENALTYCOST, SLPOBJTOPENALTYCOST	203

XSLP_OPTIMALITYTOLTARGET, SLPOPTIMALITYTOLTARGET	203
XSLP_OTOL_A, SLPOTOL_A	204
XSLP_OTOL_R, SLPOTOL_R	204
XSLP_PRESOLVE_ELIMTOL, NLPPRESOLVE_ELIMTOL	205
XSLP_PRESOLVEZERO, NLPPRESOLVEZERO	205
XSLP_PRIMALINTEGRALALPHA, NLPPRIMALINTEGRALALPHA	205
XSLP_PRIMALINTEGRALREF, NLPPRIMALINTEGRALREF	206
XSLP_SHRINK, SLP SHRINK	206
XSLP_SHRINKBIAS, SLP SHRINKBIAS	206
XSLP_STOL_A, SLPSTOL_A	207
XSLP_STOL_R, SLPSTOL_R	207
XSLP_VALIDATIONFACTOR, NLPVALIDATIONFACTOR	208
XSLP_VALIDATIONTARGET_R, NLPVALIDATIONTARGET_R	208
XSLP_VALIDATIONTARGET_K, NLPVALIDATIONTARGET_K	208
XSLP_VALIDATIONTOL_A, NLPVALIDATIONTOL_A	209
XSLP_VALIDATIONTOL_K, NLPVALIDATIONTOL_K	209
XSLP_VALIDATIONTOL_R, NLPVALIDATIONTOL_R	209
XSLP_VTOL_A, SLPVTOL_A	210
XSLP_VTOL_R, SLPVTOL_R	211
XSLP_WTOL_A, SLPWTOL_A	211
XSLP_WTOL_R, SLPWTOL_R	212
XSLP_XTOL_A, SLPXTOL_A	213
XSLP_XTOL_R, SLPXTOL_R	214
XSLP_ZERO, NLPZERO	215
20.2 Integer control parameters	216
XSLP_ALGORITHM, SLPALGORITHM	216
XSLP_ANALYZE, SLPANALYZE	218
XSLP_AUGMENTATION, SLPAUGMENTATION	219
XSLP_AUTOSAVE, SLPAUTOSAVE	220
XSLP_BARCROSSOVERSTART, SLPBARCROSSOVERSTART	221
XSLP_BARLIMIT, SLPBARLIMIT	221
XSLP_BARSTALLINGLIMIT, SLPBARSTALLINGLIMIT	222
XSLP_BARSTALLINGOBJLIMIT, SLPBARSTALLINGOBJLIMIT	222
XSLP_BARSTARTOPS, SLPBARSTARTOPS	222
XSLP_CALCTHREADS, NLPCALCTHREADS	223
XSLP_CASCADE, SLPCASCADE	223
XSLP_CASCADENLIMIT, SLPCASCADENLIMIT	224
XSLP_CONTROL	225
XSLP_CONVERGENCEOPS, SLP CONVERGENCEOPS	225
XSLP_CUTSTRATEGY, SLPCUTSTRATEGY	226
XSLP_DAMPSTART, SLPDAMPSTART	227
XSLP_DELAYUPDATEROWS, SLPDELAYUPDATEROWS	227
XSLP_DELTAOFFSET, SLPDELTAOFFSET	227
XSLP_DELTAZLIMIT, SLPDELTAZLIMIT	228
XSLP_DERIVATIVES, NLPDERIVATIVES	228
XSLP_DETERMINISTIC, NLPDETERMINISTIC	229
XSLP_ECFCHECK, SLPECFCHECK	229
XSLP_ECHOXPRSMESSAGES	230
XSLP_ERROROFFSET, SLPERROROFFSET	230
XSLP_EVALUATE, NLPEVALUATE	230
XSLP_FILTER, SLPFILTER	231
XSLP_FINDIV, NLPFINDIV	231
XSLP_FUNCEVAL, NLPFUNCEVAL	232
XSLP_GRIDHEURSELECT, SLPGRIDHEURSELECT	233
XSLP_HEURSTRATEGY, SLPHEURSTRATEGY	233

XSLP_HESSIAN, NLPHESSIAN	233
XSLP_INFEASLIMIT, SLPINFEASLIMIT	234
XSLP_ITERLIMIT, SLPITERLIMIT	234
XSLP_JACOBIAN, NLPJACOBIAN	234
XSLP_KEEPEQUALSCOLUMN, NLPKEEPEQUALSCOLUMN	235
XSLP_LINQUADBR, NLPLINQUADBR	235
XSLP_LOG, NLPLOG	236
XSLP_LSITERLIMIT, SLPLSITERLIMIT	236
XSLP_LSPATTERNLIMIT, SLPLSPATTERNLIMIT	236
XSLP_LSSTART, SLPLSSTART	237
XSLP_LSZEROLIMIT, SLPLSZEROLIMIT	237
XSLP_MAXTIME, NLPMAXTIME	237
XSLP_MIPALGORITHM, SLPMIPALGORITHM	238
XSLP_MIPCUTOFFCOUNT, SLPMIPCUTOFFCOUNT	239
XSLP_MIPCUTOFFLIMIT, SLPMIPCUTOFFLIMIT	239
XSLP_MIPDEFAULTALGORITHM, SLPMIPDEFAULTALGORITHM	240
XSLP_MIPFIXSTEPBOUNDS, SLPMIPFIXSTEPBOUNDS	240
XSLP_MIPITERLIMIT, SLPMIPITERLIMIT	241
XSLP_MIPLOG, SLPMIPLOG	241
XSLP_MIPOCOUNT, SLPMIPOCOUNT	241
XSLP_MIPRELAXSTEPBOUNDS, SLPMIPRELAXSTEPBOUNDS	242
XSLP_MULTISTART, MULTISTART	242
XSLP_MULTISTART_LOG, MULTISTART_LOG	243
XSLP_MULTISTART_MAXSOLVES, MULTISTART_MAXSOLVES	243
XSLP_MULTISTART_MAXTIME, MULTISTART_MAXTIME	243
XSLP_MULTISTART_POOLSIZE, MULTISTART_POOLSIZE	244
XSLP_MULTISTART_SEED, MULTISTART_SEED	244
XSLP_NLPSOLVER, NLPSOLVER	244
XSLP_MULTISTART_THREADS, MULTISTART_THREADS	245
XSLP_OCOUNT, SLPOCOUNT	245
XSLP_PENALTYINFOSTART, SLPPENALTYINFOSTART	246
XSLP_POSTSOLVE, NLPPOSTSOLVE	246
XSLP_PRESOLVE, NLPPRESOLVE	246
XSLP_PRESOLVELEVEL, NLPPRESOLVELEVEL	247
XSLP_PRESOLVEOPS, NLPPRESOLVEOPS	247
XSLP_PROBING, NLPPROBING	248
XSLP_REFORMULATE, NLPREFORMULATE	248
XSLP_SAMECOUNT, SLPSAMECOUNT	249
XSLP_SAMEDAMP, SLPSAMEDAMP	250
XSLP_SBROWOFFSET, SLPSBROWOFFSET	250
XSLP_SBSTART, SLPSBSTART	250
XSLP_SCALE	251
XSLP_SCALECOUNT	251
XSLP_SOLVER, LOCALSOLVER	252
XSLP_SLLOG	252
XSLP_STOPOUTOFRANGE, NLPSTOPOUTOFRANGE	252
XSLP_THREADS, NLPTHREADS	253
XSLP_THREADSAFEUSERFUNC, NLPTHREADSAFEUSERFUNC	253
XSLP_TRACEMASKOPS, SLPTRACEMASKOPS	253
XSLP_UNFINISHEDLIMIT, SLPUNFINISHEDLIMIT	254
XSLP_UPDATEOFFSET, SLPUPDATEOFFSET	255
XSLP_VCOUNT, SLPVCOUNT	255
XSLP_VLIMIT, SLPVLIMIT	256
XSLP_WCOUNT, SLPWCOUNT	256
XSLP_XCOUNT, SLPXCOUNT	257

XSLP_XLIMIT, SLPXLIMIT	258
XSLP_ZEROCRITERION, SLPZEROCRITERION	259
XSLP_ZEROCRITERIONCOUNT, SLPZEROCRITERIONCOUNT	259
XSLP_ZEROCRITERIONSTART, SLPZEROCRITERIONSTART	260
20.3 String control parameters	261
XSLP_DELTAFORMAT, SLPDELTAFORMAT	261
XSLP_ITERFALLBACKOPS, SLPITERFALLBACKOPS	261
XSLP_IVNAME, NLPIVNAME	262
XSLP_MINUSDELTAFORMAT, SLPMINUSDELTAFORMAT	262
XSLP_MINUSERERRORFORMAT, SLPMINUSERERRORFORMAT	262
XSLP_PENALTYCOLFORMAT, SLPDENALTYCOLFORMAT	263
XSLP_PENALTYROWFORMAT, SLPDENALTYROWFORMAT	263
XSLP_PLUSDELTAFORMAT, SLPPLUSDELTAFORMAT	263
XSLP_PLUSERERRORFORMAT, SLPPLUSERERRORFORMAT	264
XSLP_SBLOROWFORMAT, SLP_SBLOROWFORMAT	264
XSLP_SBNAME, SLP_SBNAM	265
XSLP_SBUPROWFORMAT, SLP_SBPUPROWFORMAT	265
XSLP_TOLNAME, SLP_TOLNAME	265
XSLP_TRACEMASK, SLPTRACEMASK	266
XSLP_UPDATEFORMAT, SLPUPDATEFORMAT	266
20.4 Knitro controls	267

21 Library functions and the programming interface 268

21.1 Counting	268
21.2 The Xpress NonLinear problem pointer	268
21.3 The load functions	269
21.4 Library functions	269
XSLPaddcoefs, XPRSslpaddcoefs	277
XSLPaddformulas, XPRSslpaddformulas	279
XSLPadddfs	281
XSLPaddtolsets	282
XSLPadduserfunction, XPRSslpadduserfunction	283
XSLPaddvars	285
XSLPcalclacks, XPRSslpcalclacks	287
XSLPcascade, XPRSslpcascade	288
XSLPcascadeorder, XPRSslpcascadeorder	289
XSLPchgascadenlimit, XPRSslpchgascadenlimit	290
XPRSslpchgcoefstr	291
XSLPchgcoef, XPRSslpchgcoef	292
XSLPchgcoef, XPRSslpchgcoef	293
XSLPchgdelatetype, XPRSslpchgdelatetype	295
XSLPchgdf	296
XSLPchgformulastring, XPRSslpchgformulastring	297
XPRSslpchgformulastr	298
XSLPchgformula, XPRSslpchgformula	299
XSLPchgrowstatus, XPRSslpchgrowstatus	300
XSLPchgrowwt, XPRSslpchgrowwt	301
XSLPchgtolset	302
XSLPchgvar	304
XSLPconstruct, XPRSslpconstruct	306
XSLPcopycallbacks	307
XSLPcopycontrols	308
XSLPcopyprob	309
XSLPcreateprob	310

XSLPdelcoefs, XPRSSlpedelcoefs	311
XSLPdelformulas, XPRSnlpdelformulas	312
XSLPdeltolsets	313
XSLPdeluserfunction, XPRSnlpdeluserfunction	314
XSLPdelvars	315
XSLPdestroyprob	316
XSLPevaluatecoef, XPRSSlpevaluatecoef	317
XSLPevaluateformula, XPRSnlpevaluateformula	318
XSLPfixpenalties, XPRSSlpxfixpenalties	319
XSLPfree	320
XSLPgetbanner	321
XSLPgetcccoef, XPRSSlpgetcccoef	322
XPRSSlpgetcoefstr	323
XSLPgetcoeffformula, XPRSSlpgetcoeffformula	324
XSLPgetcoefs, XPRSSlpgetcoefs	325
XSLPgetcolinfo, XPRSSlpgetcolinfo	326
XSLPgetdblattrib	327
XSLPgetdblcontrol	328
XSLPgetdf	329
XSLPgetformula, XPRSnlpgetformula	330
XPRSnlpgetformulastr	331
XSLPgetformulastring, XPRSnlpgetformulastring	332
XSLPgetformularows, XPRSnlpgetformularows	333
XSLPgetindex	334
XSLPgetintattrib	335
XSLPgetintcontrol	336
XSLPgetlasterror	337
XSLPgetptrattrib	338
XSLPgetrowinfo, XPRSSlpgetrowinfo	339
XSLPgetrowstatus, XPRSSlpgetrowstatus	340
XSLPgetrowwt, XPRSSlpgetrowwt	341
XSLPgetslpsol, XPRSSgetslpsol	342
XSLPgetstrattrib	343
XSLPgetstrcontrol	344
XSLPgettolset	345
XSLPgetvar	346
XSLPimportlibfunc, XPRSnlpimportlibfunc	348
XSLPinit	349
XSLPinterrupt	350
XSLPitemname	351
XSLPloadcoefs, XPRSSlploadcoefs	352
XSLPloaddfs	354
XSLPloadformulas, XPRSnlploadformulas	355
XSLPloadtolsets	357
XSLPloadvars	358
XSLPmaxim	360
XSLPminim	361
XSLPmsaddcustompreset, XPRSSmsaddcustompreset	362
XSLPmsaddjob, XPRSSmsaddjob	363
XSLPmsaddpreset, XPRSSmsaddpreset	364
XSLPmsclear, XPRSSmsclear	365
XSLPnlpoptimize, XPRSnlpoptimize	366
XSLPpostsolve, XPRSnlppostsolve	367
XSLPpresolve	368
XSLPprintmemory	369

XSLPpprintevalinfo, XPRSnlpprintevalinfo	370
XSLPreadprob	371
XSLPmaxim	372
XSLPminim	373
XSLPstore	374
XSLPpreinitialize, XPRSnlppreinitialize	375
XSLPsave	376
XSLPsaveas	377
XSLPscaling	378
XSLPsetcbcascadeend, XPRSnlpsetcbcascadeend	379
XSLPsetcbcascadestart, XPRSnlpsetcbcascadestart	380
XSLPsetcbcascadevar, XPRSnlpsetcbcascadevar	381
XSLPsetcbcascadevarfail, XPRSnlpsetcbcascadevarfail	382
XSLPsetcbcofevalerror, XPRSnlpsetcbnlpcfevalerror	383
XSLPsetcbconstruct, XPRSnlpsetcbconstruct	384
XSLPsetcbdestroy	386
XSLPsetcbdrcol, XPRSnlpsetcbldrcol	387
XSLPsetcbintsol, XPRSnlpsetcbintsol	388
XSLPsetcbiterend, XPRSnlpsetcbpiterend	389
XSLPsetcbiterstart, XPRSnlpsetcbpiterstart	390
XSLPsetcbitervar, XPRSnlpsetcbpitervar	391
XSLPsetcbmessage	392
XSLPsetcbmsjobend, XPRSnlpsetcbmsjobend	394
XSLPsetcbmsjobstart, XPRSnlpsetcbmsjobstart	395
XSLPsetcbmswinner, XPRSnlpsetcbmswinner	396
XSLPsetcboptnode	397
XSLPsetcbprenode	398
XSLPsetcbpreupdatelinearization, XPRSnlpsetcbpreupdatelinearization	399
XSLPsetcbpresolved	401
XSLPsetcbslpend	402
XSLPsetcbslpnode	403
XSLPsetcbslpstart	404
XSLPsetcurrentiv, XPRSnlpsetcurrentiv	405
XSLPsetdblcontrol	406
XSLPsetdefaultcontrol	407
XSLPsetdefaults	408
XSLPsetdetrow, XPRSnlpsetdetrow	409
XSLPsetfunctionerror, XPRSnlpsetfunctionerror	410
XSLPsetinitval, XPRSnlpsetinitval	411
XSLPsetintcontrol	412
XSLPsetlogfile	413
XSLPsetparam	414
XSLPsetstrcontrol	415
XSLPunconstruct, XPRSnlpunconstruct	416
XSLPupdatelinearization, XPRSnlpupdatelinearization	417
XSLPvalidate, XPRSnlpvalidate	418
XSLPvalidatekkt, XPRSnlpvalidatekkt	419
XSLPvalidateprob, XPRSnlpvalidateprob	420
XSLPvalidaterow, XPRSnlpvalidaterow	421
XSLPvalidatevector, XPRSnlpvalidatevector	422
XSLPwriteprob	423
XSLPwriteslsol	424

22 Internal Functions

425

22.1	Trigonometric functions	426
	ARCCOS	427
	ARCSIN	428
	ARCTAN	429
	COS	430
	SIN	431
	TAN	432
22.2	Other mathematical functions	433
	ABS	434
	ERF	435
	ERFC	436
	EXP	437
	LN	438
	LOG10	439
	MAX	440
	MIN	441
	PWL	442
	SIGN	443
	SQRT	444
23	Error Messages	445
24	Xpress Knitro Control Parameters	452
24.1	Double control parameters	457
	XKTR_PARAM_BAR_FEASMODETOL, KNITRO_PARAM_BAR_FEASMODETOL	457
	XKTR_PARAM_BAR_INITMU, KNITRO_PARAM_BAR_INITMU	457
	XKTR_PARAM_DELTA, KNITRO_PARAM_DELTA	457
	XKTR_PARAM_FEASTOL, KNITRO_PARAM_FEASTOL	457
	XKTR_PARAM_FEASTOLABS, KNITRO_PARAM_FEASTOLABS	458
	XKTR_PARAM_INFEASTOL, KNITRO_PARAM_INFEASTOL	458
	XKTR_PARAM_MIP_INTEGERTOL, KNITRO_PARAM_INTEGERTOL	458
	XKTR_PARAM_MIP_INTGAPABS, KNITRO_PARAM_INTGAPABS	458
	XKTR_PARAM_MIP_INTGAPREL, KNITRO_PARAM_INTGAPREL	459
	XKTR_PARAM_OBJRANGE, KNITRO_PARAM_OBJRANGE	459
	XKTR_PARAM_OPTTOL, KNITRO_PARAM_OPTTOL	459
	XKTR_PARAM_OPTTOLABS, KNITRO_PARAM_OPTTOLABS	459
	XKTR_PARAM_PRESOLVE_TOL, KNITRO_PARAM_PRESOLVE_TOL	460
	XKTR_PARAM_XTOL, KNITRO_PARAM_XTOL	460
24.2	Integer control parameters	461
	XKTR_PARAM_ALGORITHM, KNITRO_PARAM_ALGORITHM	461
	XKTR_PARAM_BAR_DIRECTINTERVAL, KNITRO_PARAM_BAR_DIRECTINTERVAL	461
	XKTR_PARAM_BAR_FEASIBLE, KNITRO_PARAM_BAR_FEASIBLE	461
	XKTR_PARAM_BAR_INITPT, KNITRO_PARAM_BAR_INITPT	462
	XKTR_PARAM_BAR_MAXBACKTRACK, KNITRO_PARAM_BAR_MAXBACKTRACK	462
	XKTR_PARAM_BAR_MAXCROSSIT, KNITRO_PARAM_BAR_MAXCROSSIT	463
	XKTR_PARAM_BAR_MAXREFACTOR, KNITRO_PARAM_BAR_MAXREFACTOR	463
	XKTR_PARAM_BAR_MURULE, KNITRO_PARAM_BAR_MURULE	463
	XKTR_PARAM_BAR_PENCONS, KNITRO_PARAM_BAR_PENCONS	464
	XKTR_PARAM_BAR_PENRULE, KNITRO_PARAM_BAR_PENRULE	464
	XKTR_PARAM_BAR_SWITCHRULE, KNITRO_PARAM_BAR_SWITCHRULE	465
	XKTR_PARAM_GRADOPT, KNITRO_PARAM_GRADOPT	465
	XKTR_PARAM_HESSOPT, KNITRO_PARAM_HESSOPT	465
	XKTR_PARAM_HONORBND, KNITRO_PARAM_HONORBND	466

XKTR_PARAM_LMSIZE, KNITRO_PARAM_LMSIZE	466
XKTR_PARAM_MAXCGIT, KNITRO_PARAM_MAXCGIT	467
XKTR_PARAM_MAXIT, KNITRO_PARAM_MAXIT	467
XKTR_PARAM_MIP_BRANCHRULE, KNITRO_PARAM_MIP_BRANCHRULE	467
XKTR_PARAM_MIP_GUB_BRANCH, KNITRO_PARAM_MIP_GUB_BRANCH	468
XKTR_PARAM_MIP_HEURISTIC, KNITRO_PARAM_MIP_HEURISTIC	468
XKTR_PARAM_MIP_HEURISTIC_MAXIT, KNITRO_PARAM_MIP_HEURISTIC_MAXIT	468
XKTR_PARAM_MIP_IMPLICATNS, KNITRO_PARAM_MIP_IMPLICATNS	468
XKTR_PARAM_MIP_KNAPSACK, KNITRO_PARAM_MIP_KNAPSACK	469
XKTR_PARAM_MIP_LPALG, KNITRO_PARAM_MIP_LPALG	469
XKTR_PARAM_MIP_MAXNODES, KNITRO_PARAM_MIP_MAXNODES	469
XKTR_PARAM_MIP_MAXSOLVES, KNITRO_PARAM_MIP_MAXSOLVES	470
XKTR_PARAM_MIP_METHOD, KNITRO_PARAM_MIP_METHOD	470
XKTR_PARAM_MIP_OUTINTERVAL, KNITRO_PARAM_MIP_OUTINTERVAL	470
XKTR_PARAM_MIP_OUTLEVEL, KNITRO_PARAM_MIP_OUTLEVEL	470
XKTR_PARAM_MIP_PSEUDOINIT, KNITRO_PARAM_MIP_PSEUDOINIT	471
XKTR_PARAM_MIP_ROOTALG, KNITRO_PARAM_MIP_ROOTALG	471
XKTR_PARAM_MIP_ROUNDING, KNITRO_PARAM_MIP_ROUNDING	471
XKTR_PARAM_MIP_SELECTRULE, KNITRO_PARAM_MIP_SELECTRULE	471
XKTR_PARAM_MIP_STRONG_CANDLIM, KNITRO_PARAM_MIP_STRONG_CANDLIM	472
XKTR_PARAM_MIP_STRONG_LEVEL, KNITRO_PARAM_MIP_STRONG_LEVEL	472
XKTR_PARAM_MIP_STRONG_MAXIT, KNITRO_PARAM_MIP_STRONG_MAXIT	472
XKTR_PARAM_MIP_TERMINATE, KNITRO_PARAM_MIP_TERMINATE	472
XKTR_PARAM_OUTLEV, KNITRO_PARAM_OUTLEV	473
XKTR_PARAM_PRESOLVE, KNITRO_PARAM_PRESOLVE	473
XKTR_PARAM_SCALE, KNITRO_PARAM_SCALE	473
XKTR_PARAM_SOC, KNITRO_PARAM_SOC	474

Appendix 475

A The Xpress-SLP Log 476

A.0.1 Logging controls	476
A.0.2 The structure of the log	476

B Selecting the right algorithm for a nonlinear problem - when to use the XPRS library instead of XSLP 479

B.0.1 Convex Quadratic Programs (QPs)	479
B.0.2 Convex Quadratically Constrained Quadratic Programs (QCQPs)	479
B.0.3 Convexity	480
B.0.4 Characterizing Convexity in Quadratic Constraints	480

C Files used by Xpress NonLinear 482

D Contacting FICO 483

FICO Customer Support	483
FICO Community	483
Documentation	483
FICO Learning	484
Sales and maintenance	484
About FICO	484

Index

485

I. Overview

CHAPTER 1

Introduction

This part of the manual is intended to provide a general description of the facilities available for modeling with Xpress NonLinear. It is not an exhaustive list of possibilities, and it does not go into very great depth on some of the more advanced topics. All the functions and formats are given in more detail in the second part of this manual and the Xpress-Mosel Reference Manual (Module mmxnlp section).

Xpress Nonlinear consists of:

- the Xpress Optimizer to solve linear, mixed integer linear, and convex quadratic problems,
- Xpress-SLP which uses Successive Linear Programming to solve non-linear models, and
- Artelys Knitro, which is used as a plugin to solve highly nonlinear models.

Note that FICO Xpress Nonlinear is different from FICO Xpress Global, whose documentation can be found in the [Global Solver User Guide](#). FICO Xpress Nonlinear targets finding locally optimal solutions for nonconvex problems. FICO Xpress Global finds globally optimal solutions for nonconvex problems (which might be orders of magnitude harder). For convex problems, both will find globally optimal solutions. It cannot be easily determined a priori which of the two solvers will be superior in performance, since both tackle the same problem with rather different approaches. Note that if both solvers are licensed, Xpress will by default call the global solver, unless user functions are present. For a particular convex problem, it is worth checking both independently.

The functionalities of Xpress Nonlinear extend those of the Xpress Optimizer. Almost any problem that fits into the problem types supported by the Xpress Optimizer are automatically detected and converted into the appropriate format to take advantage of the power of the optimizer's purpose written algorithms.

Xpress-SLP is in essence, a technique which involves making a linear approximation of the original problem at a chosen point, solving the linear approximation and seeing how "far away" the solution point is from the original chosen point. If it is "sufficiently close" then the solution is said to have converged and the process stops. Otherwise, a new point is chosen, based on the solution, and a new linear approximation is made. This process repeats (iterates) until the solution converges. Although this process will find a solution which is the optimum for the linear approximation, there is no guarantee that the solution will be the optimum for the original non-linear problem (that is to say: it may not be the best possible solution to the original problem). Such a solution is called a "local optimum", because it is a better solution than any others in the immediate neighborhood, but may not be better than one a long way away.

The problem of local optima can be thought of as being like trying to find the deepest valley in a range of mountains. You can find a valley relatively easily (just keep going downhill). However, when you reach it, you have no idea whether there is a deeper valley somewhere else, because the mountains block your view. You have found a local optimum, but you do not know whether it is a global optimum. Indeed, in general, finding a global optimum requires a rigorous search.

While Xpress-SLP is most powerful for large or integer nonlinear problems, Knitro which can take advantage of using second order partial derivative information can be more beneficial for highly nonlinear models.

1.1 Mathematical programs

There are many specialised forms of models in mathematical programming, and if such a form can be identified, there are usually much more efficient solution techniques available. This section describes some of the major types of problem that Xpress Nonlinear can identify automatically.

1.1.1 Linear programs

Linear programming (LP) involves solving problems of the form

$$\begin{array}{ll} \text{minimize} & c^T x \\ \text{subject to} & Ax \leq b \end{array}$$

and in practice this encompasses, via transformations, any problem whose objective and constraints are linear functions.

Such problems were traditionally solved with the simplex method, although recently interior point methods have come to be favoured for larger instances. Linear programs can be solved quickly, and solution techniques scale to enormous sizes of the matrix A . However, few applications are genuinely linear. It was common in the past, however, to approximate general functions by linear counterparts when LPs were the only class of problem with efficient solution techniques.

1.1.2 Convex quadratic programs

Convex quadratic programming (QP) involves solving problems of the form

$$\begin{array}{ll} \text{minimize} & c^T x + x^T Q x \\ \text{subject to} & Ax \leq b \end{array}$$

for which the matrix Q is symmetric and positive semi-definite (that is, $x^T Q x \geq 0$ for all x). This encompasses, via transformations, all problems with a positive semi-definite Q and linear constraints. Such problems can be solved efficiently by interior point methods, and also by quadratic variants of the simplex method.

1.1.3 Convex quadratically constrained quadratic programs

Convex quadratically constrained quadratic programming (QCQP) involves solving problems of the form

$$\begin{array}{ll} \text{minimize} & c^T x + x^T Q x \\ \text{subject to} & Ax \leq b \\ & q_j^T x + x^T P_j x \leq d_j, \forall j \end{array}$$

for which the matrix Q and all matrices P_j are positive semi-definite. The most efficient solution techniques are based on interior point methods.

1.1.4 Second order conic problems

Second order conic problems is a special form of a convex quadratically constrained quadratic program, where although the quadratic matrix is not positive semi-definite, the feasible range of the

problem is convex, and there are specialized algorithm to solve them.

$$\begin{array}{ll}\text{minimize} & c^T x + x^T Q x \\ \text{subject to} & A x \leq b \\ & x \text{ is in } C_j, \forall j\end{array}$$

for which the matrix C_j is a convex second order cone and Q is positive semi-definite. The standard form of a second order cone is $x^T I x \leq y * y$ where y is non-negative, or (a rotated second order cone) $x^T I x \leq y * z$ where y and z are non-negative. Many quadratic problems can be formulated as a second order convex conic problem, including any convex quadratically constrained quadratic programs. Transformation happens automatically for most convertible problems.

1.1.5 General nonlinear optimization problems

Nonlinear programming (NLP) involves solving problems of the form

$$\begin{array}{ll}\text{minimize} & f(x) \\ \text{subject to} & g_j(x) \leq b, \forall j\end{array}$$

where $f(x)$ is an arbitrary function, and $g_j(x)$ are a set of arbitrary functions. This is the most general type of problem, and any constrained model can be realised in this form via simple transformations.

Until recently, few practical techniques existed for tackling such problems, but it is now possible to solve even large instances using Successive Linear Programming solvers (SLP) or second-order methods.

1.1.6 Mixed integer programs

Mixed-integer programming (MIP), in the most general case, involves solving problems of the form

$$\begin{array}{ll}\text{minimize} & f(x) \\ \text{subject to} & g_j(x) \leq b, \forall j \\ & x_k \text{ integral}\end{array}$$

It can be combined with any of the previous problem types, giving Mixed-Integer Linear Programming (MILP), Mixed-Integer Quadratic Programming (MIQP), Mixed-Integer Quadratically Constrained Quadratic Programming (MIQCQP), Mixed-Integer Second Order Conic Problems (MISOCP) and Mixed-Integer Nonlinear Programming (MINLP). Efficient solution techniques now exist for all of these classes of problem.

1.2 Technology Overview

In real-world applications, it is vital to match the right optimization technology to your problem. The FICO Xpress libraries provide dedicated, high performance implementations of optimization technologies for the many model classes commonly appearing in practical applications. This includes solvers for linear programming (LP), mixed integer programming (MIP), convex quadratic programming (QP), and convex quadratically constrained programming (QCQP), general nonlinear programming (NLP), and general mixed-integer nonlinear programming (MINLP).

1.2.1 The Simplex Method

The simplex method is one of the most well-developed and highly studied mathematical programming tools. The solvers in the FICO Xpress Optimizer are the product of over 30 years of research, and include high quality, competitive implementations of the primal and dual simplex methods for both linear and quadratic programs. A key advantage of the simplex method is that it can very quickly reoptimize a problem after it has been modified, which is an important step in solving mixed integer programs.

1.2.2 The Logarithmic Barrier Method

The interior point method of the FICO Xpress Optimizer is a state of the art implementation, with leading performance across a variety of large models. It is capable of solving not only the largest and most difficult linear and convex quadratic programs, but also convex quadratically constrained quadratic and second order conic programs. It includes optimized versions of both infeasible logarithmic barrier methods, and also homogeneous self-dual methods.

1.2.3 Outer approximation schemes

A drawback of the barrier methods is that they are not efficiently warm-started. This makes these methods unattractive for solving several related problems, like the ones arising from a branch and bound search. While for linear and convex quadratic problems the simplex methods can be used, there is no immediate such alternative for convex quadratic constrained and second order methods. To bridge the gap, outer approximation cutting schemes are used, which themselves may be warm started by a barrier solution.

1.2.4 Successive Linear Programming

For general nonlinear programs which are very large, highly structured, or contain a significant linear part, the FICO Xpress Sequential Linear Programming solver (XSLP) offers exceptional performance. Successive linear programming is a first order, iterative approach for solving nonlinear models. At each iteration, a linear approximation to the original problem is solved at the current point, and the distance of the result from the selected point is examined. When the two points are sufficiently close, the solution is said to have converged and the result is returned. This technique is thus based upon solving a sequence of linear programming problems and benefits from the advanced algorithmic and presolving techniques available for linear problems. This makes XSLP scalable, as well as efficient for large problems. In addition, the relatively simple core concepts make understanding the solution process and subsequent tuning comparatively straightforward.

1.2.5 Second Order Methods

Also integrated into the Xpress suite is Knitro from Artelys, a second-order method which is particularly suited to large-scale continuous problems containing high levels of nonlinearity. Second order methods approximate a problem by examining quadratic programs fitted to a local region. This can provide information about the curvature of the solution space to the solver, which first-order methods do not have. Advanced implementations of such methods, like Knitro, may as a result be able to produce more resilient solutions. This can be especially noticeable when the initial point is close to a local optimum.

1.2.6 Mixed Integer Solvers

The FICO Xpress MIP solver is a highly scalable parallel branch and bound framework for all classes of mixed integer programs. It is based on a branch and bound search utilizing continuous solvers, advanced cutting planes, in-tree presolving and multiple heuristics, for discovering primal solutions and tightening best bounds. The search is guided by advanced methods for selecting branching variables and estimating sub-tree sizes/efforts.

Mixed integer programming forms the basis of many important applications, and the implementation in the FICO Xpress Suite has proven itself in operation for some of the world's largest organizations. For mixed integer nonlinear problems (MINLPs), we offer [FICO Xpress Global](#) to solve convex and non-convex, continuous and mixed-integer, problems to proven global optimality. Furthermore, [FICO Xpress Nonlinear](#) offers black-box optimization and some enhanced nonlinear features in a local solver context.

1.3 API naming convention

Xpress Nonlinear has been developed as an extension to the XPRS library building on the SLP solver technology, which is reflected in the naming convention. All XPRS API functions are used the same way as normal to build the linear part of the problem, while the API functions prefixed with XSLP are used for all nonlinear aspects, independently of how the problem is solved afterwards (convex quadratic problems by a dedicated solver or Knitro instead of SLP). Some controls have both an XPRS and an XSLP counterpart, for example "XPRS_PREOLVE" and "XSLP_PREOLVE". In such cases, "XSLP_PREOLVE" refers to the nonlinear presolver (even if another solver than SLP is used to solve the problem afterwards) and "XPRS_PREOLVE" refers to problems that are not deemed as general nonlinear (LP, MIP or convex quadratic); in such cases, if SLP solves one of such problems as part of its iterative process, the XPRS control is respected for such sub-solves.

CHAPTER 2

An example problem

2.1 Problem Definition

The diameter of a two-dimensional shape is the greatest distance between any two of its points. For a circle, this definition corresponds to the normal meaning of "diameter". For a polygon (with straight sides), it is equivalent to the greatest distance between any two vertices.

What is the greatest area of a polygon with N sides and a diameter of 1?

2.2 Problem Formulation

This formulation is one of two described by Prieto [1]. It is easy to visualize, and has advantages in later examples. The pentagon is about the smallest model which can reasonably be used – it is non-trivial but is still just about small enough to be written out in full.

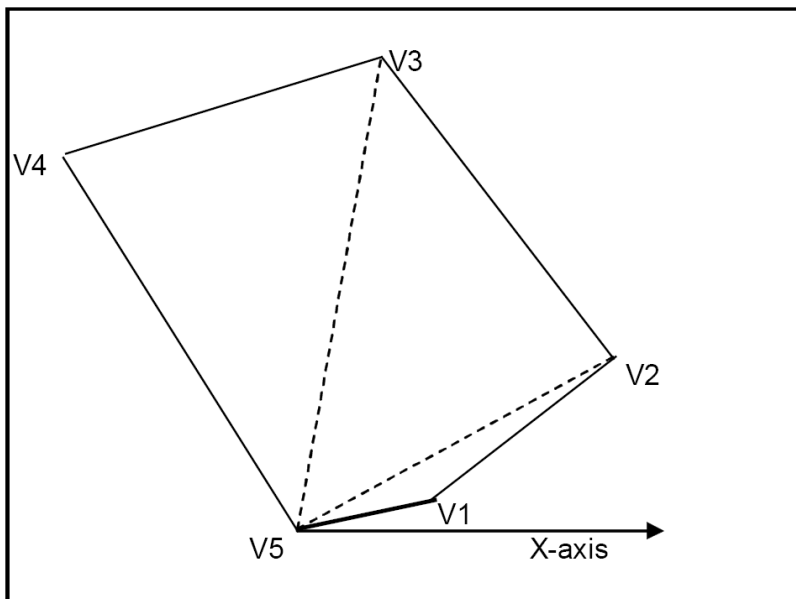


Figure 2.1: Polygon Example

One vertex (the highest-numbered, V_N) is chosen as the "base" point, and all the other vertices are measured from it, using (r, θ) coordinates – that is, the distance (" r ") is measured from the vertex, and the angle or bearing of the vertex (" θ ") is measured from the X-axis.

We shall use r_i and θ_i as the coordinates of vertex V_i . Then simple geometry and trigonometry gives:

- The area of the triangle $V_N V_i V_j$: $\text{area}(V_N V_i V_j) = \frac{1}{2} \cdot r_i \cdot r_j \cdot \sin(\theta_j - \theta_i)$
- The side $V_i V_j$ is given by: $(V_i V_j)^2 = r_i^2 + r_j^2 - 2 \cdot r_i \cdot r_j \cdot \cos(\theta_j - \theta_i)$
- The total area of the polygon is: $\sum_{i=2}^{N-1} \text{area}(V_N V_i V_{i-1})$
- The maximum diameter of 1 requires that all the sides of all the triangles are ≤ 1 – that is:
 $r_i \leq 1$ for $i = 1, \dots, N-1$
and
 $V_i V_j \leq 1$ for $i = 1, \dots, N-2, j = i+1, \dots, N-1$

We have assumed in the diagram 2.1 and in the formulation that $\theta_i \leq \theta_{i+1}$ – in other words, the vertices are in order anti-clockwise. In fact, this is not just an assumption, and we need to include these constraints as well.

In the diagram, we have assumed that the first angle θ_1 is ≥ 0 . This is not an additional restriction if we use the normal modeling convention that all variables are non-negative. We also assumed that the last vertex is still "above" the X-axis – that is, θ_{N-1} is $\leq 180^\circ$ (or π radians).

The requirement is therefore:

$$\begin{aligned}
 \text{maximize} \quad & \sum_{i=2}^{N-1} (r_i \cdot r_{i-1} \cdot \sin(\theta_i - \theta_{i-1})) * 0.5 && \text{(area of the polygon)} \\
 \text{subject to:} \quad & r_i \leq 1 \text{ for } i = 1, \dots, N-1 && \text{(distances between } V_N \text{ and other vertices)} \\
 & r_i^2 + r_j^2 - 2 \cdot r_i \cdot r_j \cdot \cos(\theta_j - \theta_i) \leq 1 \text{ for } i = 1, \dots, N-2, j = i+1, \dots, N-1 && \text{(distances between other pairs of vertices)} \\
 & \theta_1 \geq 0 && \text{(first bearing is non-negative)} \\
 & \theta_{i+1} - \theta_i \geq 0 \text{ for } i = 1, \dots, N-2 && \text{(bearings are in order)} \\
 & \theta_{N-1} \leq \pi && \text{(last vertex is above X-axis)}
 \end{aligned}$$

Reference:

(1) F.J. Prieto. *Maximum area for unit-diameter polygon of N sides, first model and second model* (Netlib AMPL programs in <ftp://netlib.bell-labs.com/netlib/ampl/models>).

CHAPTER 3

Modeling in Mosel

3.1 Basic formulation

Nonlinear capabilities in Mosel are provided by the `mmxnlp` module. Please refer to the module documentation for more details. This chapter provides a short introduction only. Note that Mosel modeling for FICO Xpress Nonlinear and [FICO Xpress Global](#) work exactly the same, s.t. you can easily switch between both solvers.

The model uses the Mosel module `mmxnlp` which contains the extensions required for modeling general non-linear expressions. This automatically loads the `mmxprs` module, so there is no need to include this explicitly as well.

```
model "Polygon"
uses "mmxnlp"
```

We can design the model to work for any number of sides, so one way to do this is to set the number of sides of the polygon as a parameter.

```
parameters
  N=5
end-parameters
```

The meanings of most of these declarations will become apparent as the modeling progresses.

```
declarations
  area: nlctr
  rho: array(1..N) of mpvar
  theta: array(1..N) of mpvar
  objdef: mpvar
  D: array(1..N,1..N) of nlctr
end-declarations
```

- The distances are described as "rho", to distinguish them from the default names for the rows in the generated matrix (which are R1, R2, etc).
- The types `nlctr` (nonlinear constraint) are defined by the `mmxnlp` module.

```
area := sum(i in 2..N-1) (rho(i) * rho(i-1) * sin(theta(i)-theta(i-1))) * 0.5
```

This uses the normal Mosel sum function to calculate the area. Notice that the formula is written in essentially the same way as normal, including the use of the `sin` function. Because the argument to the function is not a constant, Mosel will not try to evaluate the function yet; instead, it will be evaluated as part of the optimization process.

`area` is a Mosel object of type `nlctr`.

```
objdef = area
objdef is_free
```

What we really want to do is to maximize `area`. However, although Xpress NonLinear is happy in principle with a non-linear objective function, the Xpress Optimizer is not, unless it is handled in a special way. Xpress NonLinear therefore imposes the requirement that the objective function itself must be linear. This is not really a restriction, because – as in this case – it is easy to reformulate a non-linear objective function as an apparently linear one. Simply replace the function by a new `mpvar` and then maximize the value of the `mpvar`. In general, because the objective could have a positive or negative value, we make the variable free, so that it can take any value. In this example, we say:

```
objdef = area      defining the variable objdef to be equal to the non-linear ex-
                    pression area
objdef is_free     defining objdef to be a free variable
maximize (objdef)  maximizing the linear objective
```

This is firstly setting the standard bounds on the variables `rho` and `theta`. To reduce problems with sides of zero length, we impose a minimum of 0.1 on `rho(i)` instead of the default minimum of zero.

```
forall (i in 1..N-1) do
  rho(i) >= 0.1
  rho(i) <= 1
  setinitval(rho(i), 4*i*(N+1-i)/((N+1)^2))
  setinitval(theta(i), M_PI*i/N)
end-do
```

We also give Xpress NonLinear initial values by using the `setinitval` procedure. The first argument is the name of the variable, and the second is the initial value to be used. The initial values for `theta` are divided equally between 0 and π . The initial values for `rho` are designed to go from 0 (when $i = 0$ or N) to 1 (when i is about half way) and back.

```
forall (i in 1..N-2, j in i+1..N-1) do
  D(i,j) := rho(i)^2 + rho(j)^2 - rho(i)*rho(j)*2*cos(theta(j)-theta(i)) <= 1
end-do
```

This is creating the constraints `D(i, j)` which constrain the other sides of the triangles to be ≤ 1 .

These constraints could be made anonymous – that is, the assignment to an object of type `nlctr` could be omitted – but then it would not be possible to report the values.

```
forall (i in 2..N-1) do
  theta(i) >= theta(i-1) + 0.01
end-do
```

These anonymous constraints put the values of the `theta` variables in non-decreasing order. To avoid problems with triangles which have zero angles, we make each bearing at least 0.01 greater than its predecessor.

This is the boundary condition on the bearing of the final vertex.

```
theta(N-1) <= M_PI
```

3.2 Setting up and solving the problem

```
loadprob(objdef)
```

This procedure loads the currently-defined non-linear problem into the Xpress NonLinear optimization

framework. This includes any purely linear part. Where a constraint has a linear expression as its left or right hand side, that linear expression will be retained as linear relationships (constant coefficients) in the matrix. Thus, for example, in the anonymous constraint defining `objdef`, the `objdef` coefficient will be identified as a linear term and will appear as a separate item in the problem.

`maximise`

Optimization is carried out with the `maximise` or `minimise` procedures. They can take a string parameter – for example `maximise("b")` – as described in the Xpress NonLinear and Xpress Optimizer reference manuals.

With the default settings of the parameters, you will see usually nothing from the optimizer. The following parameters affect what is produced:

<code>xnlp_verbose</code>	Normally set to false. If set to true, it produces standard Xpress NonLinear iteration logging.
<code>xprs_verbose</code>	Normally set to false. If set to true, then information from the optimizer will also be output.
<code>xslp_log</code>	Normally set to -1. If set to 0, limited information is output from the SLP iterations. Settings of 1 or greater produce progressively more information for each SLP iteration.
<code>xslp_slplog</code>	If <code>xslp_log</code> is set to 0, this determines the frequency with which SLP progress is reported. The default is 10, which means that it prints every 10 SLP iterations.

3.3 Looking at the results

Within Mosel, the values of the variables and named constraints can be obtained using the `getsol`, `getslack` and similar functions. A simple report lists just the area and the positions of the vertices:

```
writeln("Area = ", getobjval)
forall (i in 1..N-1) do
  writeln("V", i, ": r=", getsol(rho(i)), " theta=", getsol(theta(i)))
end-do
```

This produces the following result for the case `N=5`:

```
Area = 0.657166
V1: r=0.616416 theta=0.703301
V2: r=1 theta=1.33111
V3: r=1 theta=1.96079
V4: r=0.620439 theta=2.58648
```

3.4 User functions

If an analytic description of the model is not available, it is possible to use black box functions implemented either directly in Mosel or by any external application.

The area calculation of the example could be implemented in Mosel as

```
public function MoselArea(I: array(Indices: range, Types: set of string) of real): real
  returned := (sum (i in 2..N-1) (I(i,"rho")*I(i-1,"rho")*sin(I(i,"theta")-I(i-1,"theta")))) * 0.5
end-function
```

The user function is linked to the model using a user function object

```

declarations
  AreaFunction : userfunc
end-declarations
AreaFunction := userfuncMosel("MoselArea")

```

The arguments, which can be any expression, are passed down using an array of expressions

```

declarations
  FunctionArg: array(RN,{ "rho", "theta" }) of nlctr
end-declarations
forall (i in 1..N) do
  FunctionArg(i, "rho") := rho(i)
  FunctionArg(i, "theta") := theta(i)
end-do

```

Once the user function is declared and the arguments built, the user function is added to the model using F:

```
Area := F(AreaFunction, FunctionArg)
```

The function arguments are copied at the point when the F function is used, any later changes to the arrays holding the arguments are ignored.

3.5 Parallel evaluation of Mosel user functions

It is possible to use parallel evaluations of simple Mosel functions that return a single real value. These functions may take an arbitrary array of nlctr expressions as input. It is the modeler's responsibility to ensure that the user functions to be called in parallel are thread-safe (i.e., they do not depend upon shared resources). Assuming the name of the user function is `MyFunc`, the user function before enabling the parallel version is expected to be declared as `usefuncMosel('MyFunc')`.

In order for `mmxnlpr` to be able to utilize parallel user function evaluations, the user function must be implemented as a public function in a Mosel package. Any initialization necessary to enable the evaluation of the user function should be performed as part of the package initialization (which is the code in in the main body of the package).

To enable parallel evaluations, a parallel enabled version of the user function needs to be generated using the `mmxnlpr` procedure `generateUFparallel`, which takes two arguments: the compiled package `.bim` name implementing the user function and the name of the user function within the package. It is good practice to use a separate Mosel model to perform this generation, keeping it separate from the main model. Multiple generated parallel user functions may be used within a single model.

The generator will produce a single Mosel file, the Mosel package `MyFunc_master`. This package also includes the worker model which will be responsible for the user function evaluations and will be resident in memory during the execution. The package also implements the parallel version of the user function, called `MyFunc_parallel`.

After compiling and including the master package into your model, it is this function that should be used in the actual model as `userfuncMosel('MyFunc_parallel', XSLP_DELTAS)`. In most cases, no other modifications are necessary, as the parallel function will detect the number of threads in the system and will start that many worker threads automatically. These will be shut down when your model finishes. Each worker's initialization code is performed only once, at the time of its first execution.

It may be necessary to explicitly start the worker threads, either to control the number of threads used, or to pass specific parameter settings to the user function package. This can be done by the procedure `MyFunc_StartWorkers(ThreadCount : integer, UfPackageParameters : string)`. In case it is necessary to stop the workers, the procedure `MyFunc_StopWorkers` may be used.

In case the user functions are computationally very expensive, by modifying the connection string in the generated module it is possible to utilize distributed/cloud-based computation of the user functions.

The worker model will only be compiled into memory during execution, but may be modified as necessary within the master model. For debugging purposes, it may be practical to redirect the worker to a file.

CHAPTER 4

The Xpress NonLinear API Functions

Instead of writing an extended MPS file and reading in the model from the file, it is possible to embed Xpress NonLinear directly into your application, and to create the problem, solve it and analyze the solution entirely by using the Xpress NonLinear API functions. This example uses the C header files and API calls. We shall assume you have some familiarity with the Xpress Optimizer API functions.

The structure of the model and the naming system will follow that used in the previous section, so you should read the chapter ?? first.

4.1 Header files

The header file containing the Xpress NonLinear definitions is `xslp.h`. This must be included together with the Xpress Optimizer header `xprs.h`, where `xprs.h` must come first.

```
#include "xprs.h"
#include "xslp.h"
```

4.2 Initialization

Xpress NonLinear and Xpress Optimizer both need to be initialized, and an empty problem created. All Xpress NonLinear functions return a code indicating whether the function completed successfully. A non-zero value indicates an error. For ease of reading, we have for the most part omitted the tests on the return codes, but a well-written program should always test the values.

```
XPRSprob mprob;
XSLPprob sprob;

if (ReturnValue=XPRSinit(NULL)) goto ErrorReturn;
if (ReturnValue=XSLPinit()) goto ErrorReturn;
if (ReturnValue=XPRScreateprob(&mprob)) goto ErrorReturn;
if (ReturnValue=XSLPcreateprob(&sprob, &mprob)) goto ErrorReturn;
```

4.3 Callbacks

It is good practice to set up at least a message callback, so that any messages produced by the system appear on the screen or in a file. The `XSLPsetcbmessage` function sets both the Xpress NonLinear and Xpress Optimizer callbacks, so that all messages appear in the same place.

```
XSLPsetcbmessage(sprob, XSLPMessage, NULL);
```

```
void XPRS_CC XSLPMessage(XSLPprob my_prob, void *my_object, char *msg, int len,
    int msg_type)
{
```



```

switch (msg_type) {
case 4: /* error */
case 3: /* warning */
case 2: /* dialogue */
case 1: /* information */
    printf("%s\n", msg);
    break;
default: /* exiting */
    fflush(stdout);
    break;
}
}

```

This is a simple callback routine, which prints any message to standard output.

4.4 Creating the linear part of the problem

The linear part of the problem, and the definitions of the rows and columns of the problem are carried out using the normal Xpress Optimizer functions.

```

#define MAXROW 20
#define MAXCOL 20
#define MAXELT 50
int nRow, nCol, nSide, nRowName, nColName;
int Sin, Cos;
char RowType[MAXROW];
double RHS[MAXROW], OBJ[MAXCOL], Element[MAXELT];
double Lower[MAXCOL], Upper[MAXCOL];
int ColStart[MAXCOL+1], RowIndex[MAXELT];
char RowNames[500], ColNames[500];

```

In this example, we have set the dimensions by using `#define` statements, rather than working out the actual sizes required from the number of sides and then allocating the space dynamically.

```

nSide = 5;
nRowName = 0;
nColName = 0;

```

By making the number of sides a variable (`nSide`) we can create other polygons by changing its value.

It is useful – at least while building a model – to be able to see what has been created. We will therefore create meaningful names for the rows and columns. `nRowName` and `nColName` count along the character buffers `RowNames` and `ColNames`.

```

nRow = nSide-2 + (nSide-1)*(nSide-2)/2 + 1;
nCol = (nSide-1)*2 + 2;
for (i=0; i<nRow; i++) RHS[i] = 0;

```

The number of constraints is:

<code>nSide-2</code>	for the relationships between adjacent thetas.
<code>(nSide-1) * (nSide-2) / 2</code>	for the distances between pairs of vertices.
<code>1</code>	for the OBJEQ non-linear "objective function".

The number of columns is:

<code>nSide-1</code>	for the thetas.
<code>nSide-1</code>	for the rhos.
<code>1</code>	for the OBJX objective function column.

We are using "C"-style numbering for rows and columns, so the counting starts from zero.

```
nRow = 0;
RowType[nRow++] = 'E'; /* OBJEQ */
nRowName = nRowName + 1 + sprintf(&RowNames[nRowName], "OBJEQ");
for (i=1; i<nSide-1; i++) {
    RowType[nRow++] = 'G'; /* T2T1 .. T4T3 */
    RHS[i] = 0.001;
    nRowName = nRowName + 1 + sprintf(&RowNames[nRowName], "T%dT%d", i+1, i);
}
```

This sets the row type indicator for OBJEQ and the theta relationships, with a right hand side of 0.001. We also create row names in the RowNames buffer. Each name is terminated by a NULL character (automatically placed there by the sprintf function). sprintf returns the length of the string written, excluding the terminating NULL character.

```
for (i=1; i<nSide-1; i++) {
    for (j=i+1; j<nSide; j++) {
        RowType[nRow] = 'L';
        RHS[nRow++] = 1.0;
        nRowName = nRowName + 1 + sprintf(&RowNames[nRowName], "V%dV%d", i, j);
    }
}
```

This defines the L-type rows which constrain the distances between pairs of vertices. The right hand side is 1.0 (the maximum value) and the names are of the form ViVj.

```
for (i=0; i<nCol; i++) {
    OBJ[i] = 0; /* objective function */
    Lower[i] = 0; /* lower bound normally zero */
    Upper[i] = XPRS_PLUSINFINITY; /* upper bound = infinity */
}
```

This sets up the standard column data, with objective function entries of zero, and default bounds of zero to plus infinity. We shall change these for the individual items as required.

```
nCol = 0;
nElement = 0;
ColStart[nCol] = nElement;
OBJ[nCol] = 1.0;
Lower[nCol++] = XPRS_MINUSINFINITY; /* free column */
Element[nElement] = -1.0;
RowIndex[nElement++] = 0;
nColName = nColName + 1 + sprintf(&ColNames[nColName], "OBJX");
```

This starts the construction of the matrix elements. nElement counts through the Element and RowIndex arrays, nCol counts through the ColStart, OBJ, Lower and Upper arrays. The first column, OBJX, has the objective function value of +1 and a value of -1 in the OBJEQ row. It is also defined to be "free", by making its lower bound equal to minus infinity.

```
iRow = 0
for (i=1; i<nSide; i++) {
    nColName = nColName + 1 + sprintf(&ColNames[nColName], "THETA%d", i);
    ColStart[nCol++] = nElement;
    if (i < nSide-1) {
        Element[nElement] = -1;
        RowIndex[nElement++] = iRow+1;
    }
    if (i > 1) {
        Element[nElement] = 1;
        RowIndex[nElement++] = iRow;
    }
    iRow++;
}
```

```
}
```

This creates the relationships between adjacent thetas. The tests on `i` are to deal with the first and last thetas which do not have relationships with both their predecessor and successor.

```
Upper[nCol-1] = 3.1415926;
```

This sets the bound on the final theta to be π . The column index is `nCol-1` because `nCol` has already been incremented.

```
for (i=1; i<nSide; i++) {
    Lower[nCol] = 0.01;          /* lower bound */
    Upper[nCol] = 1;
    ColStart[nCol++] = nElement;
    nColName = nColName + 1 + sprintf(&ColNames[nColName], "RHO%d", i);
}
ColStart[nCol] = nElement;
```

The remaining columns – the rho variables – have only non-linear formulas and so they do not appear in the linear section except as empty columns. They are bounded between 0.01 and 1.0 but have no entries. The final entry in `ColStart` is one after the end of the last column.

```
XPRSsetintcontrol(mprob, XPRS_MPSNAMELENGTH, 16);
```

If you are creating your own names – as we are here – then you need to make sure that Xpress Optimizer can handle both the names you have created and the names that will be created by Xpress NonLinear. Typically, Xpress NonLinear will create names which are three characters longer than the names you have used. If the longest name would be more than 8 characters, you should set the Xpress Optimizer name length to be larger – it comes in multiples of 8, so we have used 16 here. If you do not make the name length sufficiently large, then the `XPRSaddnames` function will return an error either here or during the Xpress NonLinear “construct” phase.

```
XPRSloadlp(mprob, "Polygon", nCol, nRow, RowType, RHS, NULL,
OBJ, ColStart, NULL,RowIndex, Element, Lower, Upper);
```

This actually loads the model into Xpress Optimizer. We are not using ranges or column element counts, which is why the two arguments are `NULL`.

```
XPRSaddnames(mprob, 1, RowNames, 0, nRow-1);
XPRSaddnames(mprob, 2, ColNames, 0, nCol-1);
```

The row and column names can now be added.

4.5 Adding the non-linear part of the problem

Be warned – this section is complicated, but it is the most efficient way – from SLP’s point of view – to input formulae. See the next section for a much easier (but less efficient) way of inputting the formulae directly.

```
#define MAXTOKEN 200
#define MAXFORM 20
...
int Sin, Cos;
FormulaStart[MAXFORM];
Type[MAXTOKEN];
double Value[MAXTOKEN];
```

The arrays for the non-linear part can often be re-used from the linear part. The new arrays are `FormulaStart` for the nonlinear formulas and `Type` and `Value` to hold the internal forms of the formulae.

```
XSLPgetindex(sprob, XSLP_INTERNALFUNCNAMES, "SIN", &Sin);
XSLPgetindex(sprob, XSLP_INTERNALFUNCNAMES, "COS", &Cos);
```

We will be using the Xpress NonLinear internal functions `SIN` and `COS`. The `XSLPgetindex` function finds the index of an Xpress NonLinear entity (character variable, internal or user function).

```
nToken = 0;
nForm = 0;
RowIndex[nForm] = 0;
FormulaStart[nForm++] = nToken;
```

For each nonlinear formula, the following information is required:

<code>RowIndex</code>	the index of the row.
<code>FormulaStart</code>	the beginning of the internal formula array for the nonlinear formula.

```
for (i=1; i<nSide-1; i++) {
    Type[nToken] = XSLP_COL;
    Value[nToken++] = nSide+i+1;
    Type[nToken] = XSLP_COL;
    Value[nToken++] = nSide+i;
    Type[nToken] = XSLP_OP;
    Value[nToken++] = XSLP_MULTIPLY;
    Type[nToken] = XSLP_RB;
    Value[nToken++] = 0;
    Type[nToken] = XSLP_COL;
    Value[nToken++] = i+1;
    Type[nToken] = XSLP_COL;
    Value[nToken++] = i;
    Type[nToken] = XSLP_OP;
    Value[nToken++] = XSLP_MINUS;
    Type[nToken] = XSLP_IFUN;
    Value[nToken++] = Sin;
    Type[nToken] = XSLP_OP;
    Value[nToken++] = XSLP_MULTIPLY;
    if (i>1) {
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_PLUS;
    }
}
```

This looks very complicated, but it is really just rather large. We are using the "reverse Polish" or "parsed" form of the formula for area. The original formula, written in the normal way, would look like this:

`RHO2 * RHO1 * SIN (THETA2 - THETA1) +`

In reverse Polish notation, tokens are pushed onto the stack or popped from it. Typically, this means that a binary operation `A x B` is written as `A B x` (push A, push B, pop A and B and push the result). The first term of our area formula then becomes:

`RHO2 RHO1 *) THETA2 THETA1 - SIN *`

Notice that the right hand bracket appears as an explicit token. This allows the `SIN` function to identify where its argument list starts – and incidentally allows functions to have varying numbers of arguments.

Each token of the formula is written as two items – `Type` and `Value`.

`Type` is an integer and is one of the defined types of token, as given in the `xslp.h` header file.

`XSLP_CON`, for example, is a constant; `XSLP_COL` is a column.

`Value` is a double precision value, and its meaning depends on the corresponding `Type`. For a `Type` of

XSLP_CON, Value is the constant value; for XSLP_COL, Value is the column number; for XSLP_OP (arithmetic operation), Value is the operand number as defined in `xslp.h`; for a function (type XSLP_IFUN for internal functions, XSLP_FUN for user functions), Value is the function number. A list of tokens for a formula is always terminated by a token of type XSLP_EOF.

The loop writes each term in order, and adds terms (using the XSLP_PLUS operator) after the first pass through the loop.

```
for (i=1; i<nSide-1; i++) {
    for (j=i+1; j<nSide; j++) {
        RowIndex[nForm] = iRow++;
        FormulaStart[nForm++] = nToken;

        Type[nToken] = XSLP_COL;
        Value[nToken++] = nSide+i;
        Type[nToken] = XSLP_CON;
        Value[nToken++] = 2;
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_EXPONENT;
        Type[nToken] = XSLP_COL;
        Value[nToken++] = nSide+j;
        Type[nToken] = XSLP_CON;
        Value[nToken++] = 2;
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_PLUS;
        Type[nToken] = XSLP_CON;
        Value[nToken++] = 2;
        Type[nToken] = XSLP_COL;
        Value[nToken++] = nSide+i;
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_MULTIPLY;
        Type[nToken] = XSLP_COL;
        Value[nToken++] = nSide+j;
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_MULTIPLY;
        Type[nToken] = XSLP_RB;
        Value[nToken++] = 0;
        Type[nToken] = XSLP_COL;
        Value[nToken++] = j;
        Type[nToken] = XSLP_COL;
        Value[nToken++] = i;
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_MINUS;
        Type[nToken] = XSLP_IFUN;
        Value[nToken++] = Cos;
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_MULTIPLY;
        Type[nToken] = XSLP_OP;
        Value[nToken++] = XSLP_MINUS;
        Type[nToken] = XSLP_EOF;
        Value[nToken++] = 0;
    }
}
```

This writes the formula for the distances between pairs of vertices. It follows the same principle as the previous formula, writing the formula in parsed form as:

$$\text{RHO}_i^2 + \text{RHO}_j^2 + 2 \text{RHO}_i * \text{RHO}_j * \cos(\text{THETA}_j - \text{THETA}_i)$$

```
XSLPloadformulass(sprob, nForm, RowIndex, FormulaStart, 1, Type,
Value);
```

The XSLPloadformulas is the most efficient way of loading non-linear formulas into a problem. There is an XSLPaddformulas function which is identical except that it does not delete any existing nonlinear formulas first. There is also an XSLPchgformula function, which can be used to change

individual nonlinear formulas one at a time. Because we are using internal parsed format, the "Parsed" flag in the argument list is set to 1.

4.6 Adding the non-linear part of the problem using character formulae

Provided that all entities – in particular columns and user functions – have explicit and unique names, the non-linear part can be input by writing the formulae as character strings. This is not as efficient as using the `XSLPloadformulas()` function but is generally easier to understand.

```
/* Build up nonlinear formulas */
/* Allow space for largest formula - approx 50 characters per side for area */
FormBuffer = (char *) malloc(50*nSide);
```

We shall be using large formulae, so we need a character buffer large enough to hold the largest formula we are using. The estimate here is 50 characters per side of the polygon for the area formula, which is the largest we are using.

```
/* Area */

BufferPos = 0;
for (i=1; i<nSide-1; i++) {
    if (i > 1) {
        BufferPos = BufferPos + sprintf(&FormBuffer[BufferPos], " + ");
    }
    BufferPos = BufferPos + sprintf(&FormBuffer[BufferPos], "RHO%d * RHO%d *
        SIN ( THETA%d - THETA%d )", i+1, i, i+1, i);
}
XSLPchgformulatext(sprob, 0, nSide, FormBuffer);
```

The area formula is of the form:

$$(RHO2 * RHO1 * SIN(THETA2 - THETA1) + RHO3 * RHO2 * SIN(THETA3 - THETA2) + \dots) / 2$$

The loop writes the product for each consecutive pair of vertices and also puts in the "+" sign after the first one.

The `XSLPchgformulatext` function is a variation of `XSLPchgformula` but uses a character string for the formula instead of passing it as arrays of tokens. The arguments to the function are:

RowIndex the index of the row.
FormBuffer the formula, written in character form.

In this case, RowIndex is zero.

```
/* Distances */
for (i=1; i<nSide-1; i++) {
    for (j=i+1; j<nSide; j++) {
        sprintf(FormBuffer, "RHO%d ^ 2 + RHO%d ^ 2 - 2 * RHO%d * RHO%d *
            COS ( THETA%d - THETA%d )", j, i, j, i, j, i);
        XSLPchgformulatext(sprob, iRow, FormBuffer);
        iRow++;
    }
}
```

This creates the formula for the distance between pairs of vertices and writes each into a new row.

Provided you have given names to any user functions in your program, you can use them in a formula in exactly the same way as SIN and COS have been used above.

4.7 Checking the data

Xpress NonLinear includes the function `XSLPwriteprob` which writes out a non-linear problem in text form which can then be checked manually. Indeed, the problem can then be run using the XSLP console program, provided there are no user functions which refer back into your compiled program. In particular, this facility does allow small versions of a problem to be checked before moving on to the full size ones.

```
XSLPwriteprob(sprob, "testmat", "");
```

The first argument is the Xpress NonLinear problem pointer; the second is the name of the matrix to be produced (the suffix ".mat" will be added automatically). The last argument allows various different types of output including "scrambled" names – that is, internally-generated names will be used rather than those you have provided. For checking purposes, this is obviously not a good idea.

4.8 Solving and printing the solution

```
XSLPmaxim(sprob, "");
```

The `XSLPmaxim` and `XSLPminim` functions perform a non-linear maximization or minimization on the current problem. The second argument can be used to pass flags as defined in the Xpress NonLinear Reference Manual.

```
XPRSwriteprtsol(mprob);
```

The standard Xpress Optimizer solution print can be obtained by using the `XPRSwriteprtsol` function. The row and column activities and dual values can be obtained using the `XPRSgetlpsol` function.

In addition, you can use the `XSLPgetvar` function to obtain the values of SLP variables – that is, of variables which are in nonlinear formulas, or which have nonlinear formulas. If you are using cascading (see the Xpress NonLinear reference manual for more details) so that Xpress NonLinear recalculates the values of the dependent SLP variables at each SLP iteration, then the value from `XSLPgetvar` will be the recalculated value, whereas the value from `XPRSgetlpsol` will be the value from the LP solution (before recalculation).

4.9 Closing the program

The `XSLPdestroyprob` function frees any system resources allocated by Xpress NonLinear for the specific problem. The problem pointer is then no longer valid. `XPRSdestroyprob` performs a similar function for the underlying linear problem `mprob`. The `XSLPfree` function frees any system resources allocated by Xpress NonLinear. You must then call `XPRSfree` to perform a similar operation for the optimizer.

```
XSLPdestroyprob(sprob);
XPRSdestroyprob(mprob);
XSLPfree();
XPRSfree();
```

If these functions are not called, the program may appear to have worked and terminated correctly. However, in such a case there may be areas of memory which are not returned to the system when the program terminates and so repeated executions of the program will result in progressive loss of available memory to the system, which will manifest itself in poorer performance and could ultimately produce a system crash.

4.10 Adding initial values

So far, Xpress NonLinear has started by using values which it estimates for itself. Because most of the variables are bounded, these initial values are fairly reasonable, and the model will solve. However, in general, you will need to provide initial values for at least some of the variables. Initial values are provided using the `XSLPsetinitval` function.

CHAPTER 5

The Nonlinear Console Program

5.1 The Console Nonlinear

The nonlinear optimizer is an extension to the FICO Xpress Optimizer interactive console.

The console for nonlinear is started from the command line using the following syntax:

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

5.1.1 The nonlinear console extensions

The nonlinear console is an extension of the Xpress optimizer console. The optimizer automatically switches to nonlinear mode if a nonlinear license is detected. All the optimizer console commands work the same way as in the normal optimizer console. The active working problem for those commands is the actual linearization after augmentation, and the linear part of the problem before augmentation.

Optimizer console commands with an extended effect:

readprob	Read in an MPS/MAT or LP file
minim	Minimize an LP, a MIP or an SLP problem
maxim	Maximize an LP, a MIP or an SLP problem
lpoptimize	Minimize or maximize a problem
mipoptimize	Solve the problem to MIP optimality
writeprob	Export the problem into file
dumpcontrols	Display controls which are at a non default value

The MPS file can be an extended MPS file containing a nonlinear model. The `lpoptimize`, `nloptimize`, `mipoptimize`, and `optimize` commands will call the right algorithm depending on the presence of nonlinearities. Note that if a license for Xpress Global is present, the global solver will be called by default. You need to pass the "-s" flag to `nloptimize`, or `optimize` to enforce a local solve. In general, all commands accept the same flags as the corresponding library function.

New commands:

cascade	Perform cascading
cascadeorder	Recalculate the cascading order
construct	Construct the augmented problem
dumpattributes	Display problem attributes
reinitialize	Reinitialize an augmented problem
setcurrentiv	Copy the current solution as initial value
slp_save	XSLPsave
slp_scaling	Display scaling statistics
unconstruct	Remove the augmentation
validate	Validate the current solution
validatekkt	Validate the kkt conditions for the current solution

In order to separate XSLP controls and attributes for the XPRS ones, all XSLP controls and attributes are pretagged as XSLP_ or SLP_, for example XSLP_ALGORITHM.

5.1.2 Common features of the Xpress Optimizer and the Xpress Nonlinear Optimizer console

All features of the Xpress optimizer console program is supported. For a full description, please refer to the Xpress optimizer reference manual.

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 Console XSLP can also be started by typing xslp into the "Run ..." dialog box in the Start menu.

The Console XSLP provides a quick and convenient interface for operating on a single problem loaded into XSLP. The Console XSLP 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.

The Console SLP 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 "CONST" followed by the Tab key Console Xpress will complete to the "CONSTRUCT". Note that once you have finished inputting the command name portion of your command line, Console Xpress can also auto-complete on file names. Note that the auto-completion of file names is case-sensitive.

Console XSLP 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 XSLP is first started it will attempt to read in an initialization file named optimizer.ini from the current working directory. This is an ASCII "script" file that may contain commands to be run at start up, which are intended to setup a customized default Console Xpress environment for the user (e.g., defining custom controls settings on the Console Xpress problem).

The Console XSLP interactive command line hosts a TCL script parser (<http://www.tcl.tk>). With TCL scripting the user can program flow control into their 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 Xpress 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 XSLP 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 and XSLP. The following example shows how to proceed with write and read access to the XSLP_ALGORITHM control:

```
[xpress C:\] xslp_algorithm=166
[xpress C:\] xslp_algorithm
166
```

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

```
[xpress C:\] set xslp_algorithm 166
166
[xpress C:\] $miplog
166
```

For examples on how TCL can be used for scripting, tuning and testing models, please refer to the Xpress Optimizer reference manual.

Console XSLP users may interrupt the running of the commands (e.g., `minim`) by typing Ctrl-C. Once interrupted Console Xpress 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.

When Console XSLP is being run with script input then Ctrl-C will not return to the command prompt and the Console Xpress process will simply stop.

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

The XSLP console program can be used as a direct substitute for the Xpress Optimizer console program. The one exception is the fixed format MPS files, which is not supported by XSLP and thus neither by the XSLP console.

II. Advanced

CHAPTER 6

Nonlinear Problems

Xpress NonLinear will solve nonlinear problems. In this context, a nonlinear problem is one in which there are nonlinear relationships between variables or where there are nonlinear terms in the objective function. There is no such thing as a nonlinear variable — all variables are effectively the same — but there are nonlinear constraints and formulae. A nonlinear *constraint* contains terms which are not linear. A nonlinear *term* is one which is not a constant and is not a variable with a constant coefficient. A nonlinear constraint can contain any number of nonlinear terms.

Xpress NonLinear will also solve linear problems — that is, if the problem presented to Xpress NonLinear does not contain any nonlinear terms, then Xpress NonLinear will still solve it, using the normal optimizer library.

The solution mechanism used by Xpress-SLP is *Successive* (or *Sequential*) *Linear Programming*. This involves building a linear approximation to the original nonlinear problem, solving this approximation (to an optimal solution) and attempting to validate the result against the original problem. If the linear optimal solution is sufficiently close to a solution to the original problem, then the SLP is said to have *converged*, and the procedure stops. Otherwise, a new approximation is created and the process is repeated. Xpress-SLP has a number of features which help to create good approximations to the original problem and therefore help to produce a rapid solution.

When licensed, Xpress NonLinear may also utilize Knitro to solve nonlinear problems.

Note that although the solution is the result of an optimization of the linear approximation, there is no guarantee that it will be an optimal solution to the original nonlinear problem. It may be a local optimum — that is, it is a better solution than any points in its immediate neighborhood, but there is a better solution rather further away. However, a converged SLP solution will always be (to within defined tolerances) a self-consistent — and therefore practical — solution to the original problem.

6.1 Coefficients and formulas

Later in this manual, it will be helpful to distinguish between nonlinear expressions written as coefficients and those written as formulas.

If X is a variable, then in the nonlinear expression $X * f(Y)$, $f(Y)$ is the *coefficient* of X .

If $f(X)$ appears in a nonlinear constraint, then $f(X)$ is a *nonlinear formula* in the nonlinear constraint.

If $X * f(Y)$ appears in a nonlinear constraint, then the entity $X * f(Y)$ is a *term* in the nonlinear constraint.

As this implies, a formula written as a variable multiplied by a coefficient can always be viewed as a term, but there are terms which cannot be viewed as variables multiplied by coefficients. For example, in the constraint

$$X - \text{SIN}(Y) = 0,$$

$\text{SIN}(Y)$ is a *formula* and cannot be written as a coefficient.

6.2 SLP variables

A variable which appears in a nonlinear coefficient or formula is described as an *SLP variable*.

Normally, any variable which has a nonlinear coefficient will also be treated as an SLP variable. However, it is possible to set options so that variables which do not appear in nonlinear coefficients or formulas are not treated as SLP variables.

Any variable, whether it is related to a nonlinear formula or not, can be defined by the user as an SLP variable. This is most easily achieved by setting an initial value for the variable.

6.3 Local and global optimality

A globally optimal solution is a feasible solution with the best objective value among all feasible solutions for an optimization problem. In contrast, a locally optimal solution only has the best objective value within a limited neighborhood surrounding it. It's important to note that better solutions may exist in different regions of the problem space. While for convex problems every local optimum is a global optimum, this is not the case for general nonlinear problems.

The FICO Xpress library offers solvers that provide global optima for both convex and nonconvex problems. The base offering includes solving linear, quadratic, and quadratically constrained problems, as well as their mixed-integer counterparts. To handle more general nonlinear problems, a license for [FICO Xpress Global](#) is required.

For black-box optimization or obtaining local optima with duals for continuous nonlinear problems, a license for [FICO Xpress Nonlinear](#) is necessary. FICO Xpress Nonlinear can also be used as a heuristic for mixed-integer nonlinear problems. It's important to note that finding a local optimum is usually much faster than finding a global optimum, but using a local solver doesn't provide guarantees on the global quality of the solution and may encounter issues such as local infeasibilities.

Lastly, neither local nor global optima are typically unique. The solution returned by a solver depends on the control settings and, especially for non-convex problems, the initial values provided. A connected set of initial points that yield the same locally optimal solutions is referred to as a region of attraction. These regions are algorithm and setting dependent.

6.4 Convexity

Convex problems have many desirable characteristics from the perspective of mathematical optimization. Perhaps the most significant of these is that should both the objective and the feasible region be convex, any local optimally solutions found are also known immediately to be globally optimal.

A constraint $f(x) \leq 0$ is convex if the matrix of second derivatives of f , that is to say its Hessian, is positive semi-definite at every point at which it exists. This requirement can be understood geometrically as requiring every point on every line segment which connects two points satisfying the constraint to also satisfy the constraint. It follows trivially that linear functions always lead to convex constraints, and that a nonlinear equality constraint is never convex.

For regions, a similar property must hold. If any two points of the region can be connected by a line segment which lies fully in the region itself, the region is convex. This extension is straightforward when the the properties of convex functions are considered.

It is important to note that convexity is necessary for some solution techniques and not for others. In particular, some solvers require convexity of the constraints and objective function to hold only in the feasible region, whilst others may require convexity to hold across the entire space, including infeasible points. In the special case of quadratic and quadratically constrained programs, Xpress NonLinear




Figure 6.1: Two convex functions on the left, and two non-convex functions on the right.

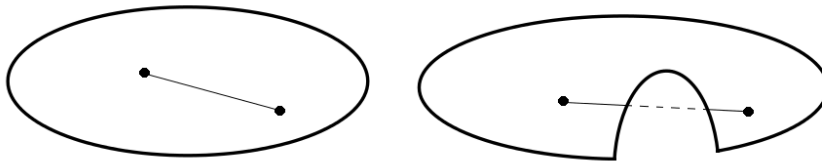


Figure 6.2: A convex region on the left and a non-convex region on the right.

seamlessly migrates problems to solvers whose convexity requirements match the convexity of the problem.

6.5 Converged and practical solutions

In a strict mathematical sense, an algorithm is said to have converged if repeated iterations do not alter the coordinates of its solution significantly. A more practical view of convergence, as used in the nonlinear solvers of the Xpress suite, is to also consider the algorithm to have converged if repeated iterations have no significant effect on either the objective value or upon feasibility. This will be called extended convergence to distinguish it from the strict sense.

For some problems, a solver may visit points at which the local neighborhood is very complex, or even malformed due to numerical issues. In this situation, the best results may be obtained when convergence of some of the variables is forced. This leads to practical solutions, which are feasible and converged in most variables, but the remaining variables have had their convergence forced by the solver, for example by means of a trust region. Although these solutions are not locally optimal in a strict sense, they provide meaningful, useful results for difficult problems in practice.

6.6 The duals of general, nonlinear program

The dual of a mathematical program plays a fundamental role in the theory of continuous optimization. Each variable in a problem has a corresponding partner in that problem's dual, and the values of those variables are called the reduced costs and dual multipliers (shadow prices). Xpress NonLinear makes estimates of these values available. These are normally defined in a similar way to the usual linear programming case, so that each value represents the rate of change of the objective when either increasing the corresponding primal variable or relaxing the corresponding primal constraint.

From an algorithmic perspective, one of the most important roles of the dual variables is to characterize local optimality. In this context, the dual multipliers and reduced costs are called Lagrange multipliers, and a solution with both primal and dual feasible variables satisfies the Karush-Kuhn-Tucker conditions.

However, it is important to note that for general nonlinear problems, there exist situations in which there are no such multipliers. Geometrically, this means that the slope of the objective function is orthogonal to the linearization of the active constraints, but that their curvature still prevents any movement in the improving direction.

As a simple example, consider:

$$\begin{array}{ll}\text{minimize} & y \\ \text{subject to} & x^2 + y^2 \leq 1 \\ & (x - 2)^2 + y^2 \leq 1\end{array}$$

which is shown graphically in figure 6.3.

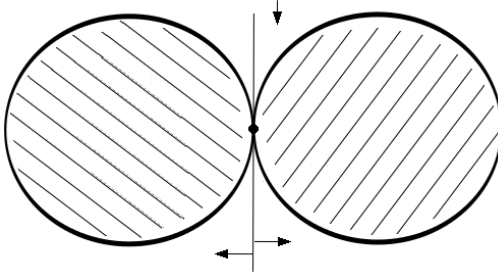


Figure 6.3: A problem admitting no dual values

This problem has a single feasible solution at (1,0). Reduced costs and dual multipliers could never be meaningful indicators of optimality, and indeed are not well-defined for this problem. Intuitively, this arises because the feasible region lacks an interior, and the existence of an interior (also referred to as the Slater condition) is one of several alternative conditions which can be enforced to ensure that such situations do not occur. The other common condition for well-defined duals is that the gradients of the active constraints are linearly independent.

Problems without valid duals do not often arise in practice, but it is important to be aware of the possibility. Analytic detection of such issues is difficult, and they manifest instead in the form of unexpectedly large or otherwise implausible dual values.

CHAPTER 7

Extended MPS file format

One method of inputting a problem to Xpress NonLinear is from a text file which is similar to the normal MPS format matrix file. The Xpress NonLinear file uses *free format* MPS-style data. All the features of normal free-format MPS are supported. There are no changes to the sections except as indicated below.

Note: the use of free-format requires that no name in the matrix contains any leading or embedded spaces and that no name could be interpreted as a number. Therefore, the following names are invalid:

B 02 because it contains an embedded space;

1E02 because it could be interpreted as 100 (the scientific or floating-point format number, 1.0E02).

It is possible to use column and row names including mathematical operators. A variable name **a+b** is valid. However, as an expression **a + b** would be interpreted as the addition of variables **a** and **b** - note the spaces between the variable names - it is best practice to avoid such names when possible. SLP will produce a warning if such names are encountered in the MPS file.

7.1 Formulae

One new feature of the Extended MPS format is the *formula*. A formula is written in much the same way as it would be in any programming language or spreadsheet. It is made up of (for example) constants, functions, the names of variables, and mathematical operators. The formula always starts with an equals sign, and each item (or *token*) is separated from its neighbors by one or more spaces.

Tokens may be one of the following:

- A constant;
- The name of a variable;
- An arithmetic operator "+", "-", "*", "/";
- The exponentiation operator "**" or "^";
- An opening or closing bracket "(" or ")";
- A comma "," separating a list of function arguments;
- The name of a supported internal function such as LOG, SIN, EXP;
- The name of a user-supplied function;
- A colon ":" preceding the return argument indicator of a multi-valued function;

- The name of a return argument from a multi-valued function.

The following are valid formulae:

- = SIN (A / B) SIN is a recognized internal function which takes one argument and returns one result (the sin of its argument).
- = A ^ B ^ is the exponentiation symbol. Note that the *formula* may have valid syntax but it still may not be possible to evaluate it (for example if $A = -1$ and $B = 0.5$).
- = MyFunc1 (C1 , - C2 , C3 : 1) MyFunc1 must be a function which can take three arguments and which returns an array of results. This formula is asking for the first item in the array.

The following are not valid formulae:

- SIN (A) Missing the equals sign at the start
- = SIN(A) No spaces between adjacent tokens
- = A * * B "***" is exponentiation, "*" (with an embedded space) is not a recognized operation.
- = MyFunc1 (C1 , - C2 , C3 , 1) If MyFunc1 is as shown in the previous set of examples, it returns an array of results. The last argument to the function must be delimited by a colon, not a comma, and is the name or number of the item to be returned as the value of the function.

There is no limit to the length of a formula. However, parsing very long records can be slow, and consideration should be given to pre-parsing them and passing the parsed formula to Xpress NonLinear rather than asking it to parse the formula itself.

7.2 COLUMNS

Normal MPS-style records of the form

column row1 value1 [row2 value2]

are supported. Non-linear relationships are modeled by using a formula instead of a constant in the *value1* field. If a formula is used, then only one coefficient can be described in the record (that is, there can be no *row2 value2*). The formula begins with an equals sign ("=") and is as described in the previous section.

A formula must be contained entirely on one record.

Variables used in formulae may be included in the COLUMNS section as variables, or may exist only as items within formulae. A variable which exists only within formulae is called an *implicit variable*.

Sometimes the non-linearity cannot be written as a coefficient. For example, in the constraint $Y - \text{LOG}(X) = 0$,

$\text{LOG}(X)$ cannot be written in the form of a coefficient. In such a case, the reserved column name "=" may be used in the first field of the record as shown:

Y	MyRow	1
=	MyRow	= - LOG (X)

Effectively, "=" is a column with a fixed activity of 1.0 .

When a file is read by `XSLPreadprob`, more than one coefficient can be defined for the same column/row intersection. As long as there is at most one constant coefficient (one not written as a formula), the coefficients will be added together. If there are two or more constant coefficients for the same intersection, they will be handled by the Optimizer according to its own rules (normally additive, but the objective function retains only the last coefficient).

7.3 BOUNDS

Bounds can be included for variables which are not defined explicitly in the `COLUMNS` section of the matrix. If they are not in the `COLUMNS` section, they must appear as variables within formulae (*implicit variables*). A `BOUNDS` entry for an item which is not a column or a variable will produce a warning message and will be ignored.

MIP entities (such as integer variables and members of Special Ordered Sets) must be defined explicitly in the `COLUMNS` section of the matrix. If a variable would otherwise appear only in formulae in coefficients, then it should be included in the `COLUMNS` section with a zero entry in a row (for example, the objective function) which will not affect the result.

7.4 SLPDATA

`SLPDATA` is a new section which holds additional information for solving the non-linear problem using SLP.

Many of the data items have a *setname*. This works in the same way as the `BOUND`, `RANGE` or `RHS` name, in that a number of different values can be given, each with a different set name, and the one which is actually used is then selected by specifying the appropriate setname before reading the problem.

Record type `IV` and the tolerance records `Tx`, `Rx` can have "=" as the variable name. This provides a default value for the record type, which will be used if no specific information is given for a particular variable.

Note that only linear `BOUND` types can be included in the `SLPDATA` section. Bound types for MIP entities (discrete variables and special ordered sets) must be provided in the normal `BOUNDS` section and the variables must also appear explicitly in the `COLUMNS` section.

All of the items in the `SLPDATA` section can be loaded into a model using Xpress NonLinear function calls.

7.4.1 DR (Determining row)

`DR variable rowname [weighting] [limit]`

The `DR` record defines the *determining row* for a variable.

In most non-linear problems, there are some variables which are effectively defined by means of an equation in terms of other variables. Such an equation is called a *determining row*. If Xpress NonLinear knows the determining rows for the variables which appear in coefficients, then it can provide better linear approximations for the problem and can then solve it more quickly. Optionally, a non-zero integer value can be included in the *weighting* field. Variables which have weights will generally be evaluated in order of increasing weight. Variables without weights will generally be evaluated after those which do have weights. However, if a variable `A` (with or without a weight) is dependent through its determining row on another variable `B`, then `B` will always be evaluated first. The optional *limit* field provides a variable specific value for `XSLP_CASCADENLIMIT`.

Example:

DR X Row1

This defines Row1 as the determining row for the variable X. If Row1 is

$X - Y * Z = 6$

then Y and Z will be recalculated first before X is recalculated as $Y * Z + 6$.

7.4.2 EC (Enforced constraint)

EC rowname

The EC record defines an *enforced constraint*. Penalty error vectors are never added to enforced constraints, so the effect of such constraints is maintained at all times.

Note that this means the *linearized* version of the enforced constraint will be active, so it is important to appreciate that enforcing too many constraints can easily lead to infeasible linearizations which will make it hard to solve the original nonlinear problem.

Example:

EC Row1

This defines Row1 as an enforced constraint. When the SLP is augmented, no penalty error vectors will be added to the constraint, so the linearized version of Row1 will constrain the linearized problem in the same sense (L, G or E) as the nonlinear version of Row1 constrains the original nonlinear problem.

7.4.3 FR (Free variable)

FR boundname variable

An FR record performs the same function in the SLPDATA section as it does in the BOUNDS section. It can be used for bounding variables which do not appear as explicit columns in the matrix.

7.4.4 FX (Fixed variable)

FX boundname variable value

An FX record performs the same function in the SLPDATA section as it does in the BOUNDS section. It can be used for bounding variables which do not appear as explicit columns in the matrix.

7.4.5 IV (Initial value)

IV setname variable [value | = formula]

An IV record specifies the initial value for a variable. All variables which appear in coefficients or terms, or which have non-linear coefficients, should have an IV record.

A formula provided as the initial value for a variable can contain references to other variables. It will be evaluated based on the initial values of those variables (which may themselves be calculated by formula). It is the user's responsibility to ensure that there are no circular references within the formulae. Formulae are typically used to calculate consistent initial values for dependent variables based on the values of independent variables.

If an IV record is provided for the *equals column* (the column whose name is "=" and which has a fixed value of 1.0), the value provided will be used for all SLP variables which do not have an explicit initial value of their own.

If there is no explicit or implied initial value for an SLP variable, the value of control parameter XSLP_DEFAULTIV will be used.

If the initial value is greater than the upper bound of the variable, the upper bound will be used; if the

initial value is less than the lower bound of the variable, the lower bound will be used.

If both a formula and a value are provided, then the explicit value will be used.

Example:

```
IV IVSET1 Col99 1.4971
```

```
IV IVSET2 Col99 2.5793
```

This sets the initial value of column Col99. The initial value to be used is selected using control parameter XSLP_IVNAME. If no selection is made, the first initial value set found will be used.

If Col99 is bounded in the range $1 \leq \text{Col99} \leq 2$ then in the second case (when IVSET2 is selected), an initial value of 2 will be used because the value given is greater than the upper bound.

```
IV IVSET2 Col98 = Col99 * 2
```

This sets the value of Col98 to twice the initial value of Col99 when IVSET2 is the selected initial value set.

7.4.6 LO (Lower bounded variable)

```
LO boundname variable value
```

A LO record performs the same function in the SLPDATA section as it does in the BOUNDS section. It can be used for bounding variables which do not appear as explicit columns in the matrix.

7.4.7 Rx, Tx (Relative and absolute convergence tolerances)

```
Rx setname variable value
```

```
Tx setname variable value
```

The Tx and Rx records (where "x" is one of the defined tolerance types) define specific tolerances for convergence of the variable. See the section "convergence criteria" for a list of convergence tolerances. The same tolerance set name (*setname*) is used for all the tolerance records.

Example:

```
RA TOLSET1 Col99 0.005
```

```
TA TOLSET1 Col99 0.05
```

```
RI TOLSET1 Col99 0.015
```

```
RA TOLSET1 Col101 0.01
```

```
RA TOLSET2 Col101 0.015
```

These records set convergence tolerances for variables Col99 and Col101. Tolerances RA (relative convergence tolerance), TA (absolute convergence tolerance) and RI (relative impact tolerance) are set for Col99 using the tolerance set named TOLSET1.

Tolerance RA is set for variable Col101 using tolerance sets named TOLSET1 and TOLSET2.

If control parameter XSLP_TOLNAME is set to the name of a tolerance set before the problem is read using XSLPreadprob, then only the tolerances on records with that tolerance set will be used. If XSLP_TOLNAME is blank or not set, then the name of the set on the first tolerance record will be used.

7.4.8 SB (Initial step bound)

```
SB setname variable value
```

An SB record defines the initial step bounds for a variable. Step bounds are symmetric (i.e. the bounds on the delta are $-SB \leq \text{delta} \leq +SB$). If a value of 1.0E+20 is used (equivalent to XPRS_PLUSINFINITY in programming), the delta will never have step bounds applied, and will almost always be regarded as converged.

If there is no explicit initial step bound for an SLP variable, a value will be estimated either from the size of the coefficients in the initial linearization, or from the values of the variable during the early SLP iterations. The value of control parameter `XSLP_DEFAULTSTEPBOUND` provides a lower limit for the step bounds in such cases.

If there is no explicit initial step bound, then the closure convergence tolerance cannot be applied to the variable.

Example:

```
SB SBSET1 Col199 1.5
SB SBSET2 Col199 7.5
```

This sets the initial step bound of column `Col199`. The value to be used is selected using control parameter `XSLP_SBNAME`. If no selection is made, the first step bound set found will be used.

7.4.9 UF (User function)

```
UF funcname = libraryname (functiontype) linkage = library
```

A UF record defines a user function.

The definition includes the function's type which matches the parameter supplied to the `adduserfunction` call.

Example:

```
UF MyFunc ( VECMAPDELTA ) DLL = UserLib
```

This defines a user function called `MyFunc` that takes multiple input arguments and supplies its own derivatives. The linkage is `DLL` (free-standing user library or DLL) and the function is in file `UserLib`.

7.4.10 UP (Free variable)

```
UP boundname variable value
```

An UP record performs the same function in the `SLPDATA` section as it does in the `BOUNDS` section. It can be used for bounding variables which do not appear as explicit columns in the matrix.

7.4.11 WT (Explicit row weight)

```
WT rowname value
```

The WT record is a way of setting the initial penalty weighting for a row. If `value` is positive, then the default initial weight is multiplied by the value given. If `value` is negative, then the absolute value will be used instead of the default weight.

Increasing the penalty weighting of a row makes it less attractive to violate the constraint during the SLP iterations.

Examples:

```
WT Row1 3
```

This changes the initial weighting on `Row1` by multiplying by 3 the default weight calculated by Xpress-SLP during problem augmentation.

```
WT Row1 -3
```

This sets the initial weighting on `Row1` to 3.

7.4.12 DL (variable specific Determining row cascade iteration Limit)

```
DL columnname limit
```

A `DL` record specifies a variable specific iteration limit to be imposed on the number of iterations when cascading the variable. This can be used to overwrite the setting of `XSLP_CASCADENLIMIT` for a specific variable.

CHAPTER 8

Xpress-SLP Solution Process

This section gives a brief overview of the sequence of operations within Xpress-SLP once the data has been set up. The positions of the possible user callbacks are also shown.

```
Check if problem is an SLP problem or not. Call the appropriate XPRS library function if not, and DONE.
[Call out to user callback if set by XSLPsetcbstart]
Augment the matrix (create the linearized structure) if not already done
If determining row data supplied, calculate cascading order and detect determining columns
DO
    [Call out to user callback if set by XSLPsetcbiterstart]
    If previous solution available, pre-process solution
        Execute line search
        [Call out to user callback if set by XSLPsetcbstart]
        Sequentially update values of SLP variables (cascading) and re-calculate coefficients
        For each variable (in a suitable evaluation order):
            Update solution value (cascading) and re-calculate coefficients
            [Call out to user callback if set by XSLPsetcbstart]
        [Call out to user callback if set by XSLPsetcbstart]
    Update penalties
    Update coefficients, bounds and RHS in linearized matrix
    Solve linearized problem using the Xpress Optimizer
    Recover SLP variable and delta solution values
    Test convergence against specified tolerances and other criteria
    For each variable:
        Test convergence against specified tolerances
        [Call out to user callback if set by XSLPsetcbstart]
    For each variable with a determining column:
        Check value of determining column and fix variable when necessary, or
        [Call out to user callback if set by XSLPsetcbstart]
        Reset variable convergence status if a change is made to a variable
    If not all variables have converged, check for other extended convergence criteria
    If the solution has converged, then BREAK
    For each SLP variable:
        Update history
        Reset step bounds
    [Call out to user callback if set by XSLPsetcbstart]
    Change row types for DC rows as required
    If SLP iteration limit is reached, then BREAK
ENDDO
[Call out to user callback if set by XSLPsetcbend]
```

For MISLP (mixed-integer SLP) problems, the above solution process is normally repeated at each node.

The standard procedure for each node is as follows:

```

Initialize node
[Call out to user callback if set by XSLPsetcbprenode]
Solve node using SLP procedure
If an optimal solution is obtained for the node then
    [Call out to user callback if set by XSLPsetcboptnode]
If an integer optimal solution is obtained for the node then
    [Call out to user callback if set by XSLPsetcbintsol]
When node is completed
    [Call out to user callback if set by XSLPsetcbslpnode]
  
```

When a problem is destroyed, there is a call out to the user callback set by `XSLPsetcbdestroy`.

8.1 Analyzing the solution process

Xpress-SLP provides a comprehensive set of callbacks to interact with, and to analyze the solution process. However, there are a set of purpose build options that are intended to assist and make the analysis more efficient.

For infeasible problems, it often helps to identify the source of conflict by running XPRESS' Irreducible Infeasibility Set (IIS) finder tool. The set found by IIS often helps to either point to a problem in the original model formulation, or if the infeasibility is a result of conflicting step bounds or linearization updates; please see control `XSLP_ANALYZE`.

It is often advantageous to trace a certain variable, constraint or a certain property through the solution process. `XSLP_TRACEMASK` and `XSLP_TRACEMASKOPS` allows for collecting detailed information during the solution process, without the need to stop XSLP between iterations.

For in depth debugging purposes or support requests, it is possible to create XSLP save files and linearizations at various iterations, controlled by `XSLP_AUTOSAVE` and `XSLP_ANALYZE`.

8.2 The initial point

The solution process is sensitive to the initial values which are selected for variables in the problem, and particularly so for non-convex problems. It is not uncommon for a general nonlinear problem to have a feasible region which is not connected, and in this case the starting point may largely determine which region, connected set, or basin of attraction the final solution belongs to.

Note that it may not always be beneficial to completely specify an initial point, as the solvers themselves may be able to detect suitable starting values for some or all of the variables.

8.3 Derivatives

Both XSLP and Knitro require the availability of derivative information for the constraints and objective function in order to solve a problem. In the Xpress NonLinear framework, several advanced approaches to the production of both first and second order derivatives (the Jacobian and Hessian matrices) are available, and which approach is used can be controlled by the user.

8.3.1 Finite Differences

The simplest such method is the use of finite differences, sometimes called numerical derivatives. This is a relatively coarse approximation, in which the function is evaluated in a small neighborhood of the

point in question. The standard argument from calculus indicates that an increasingly accurate approximation to the derivative of the function will be found as the size of the neighborhood decreases. This argument ignores the effects of floating point arithmetic, however, which can make it difficult to select values sufficiently small to give a good approximation to the function, and yet sufficiently large to avoid substantial numerical error.

The high performance implementation in XSLP makes use of subexpression caching to improve performance, but finite differences are inherently inefficient. They may however be necessary when the function itself is not known in closed form. When analytic approaches cannot be used, due to the use of expensive black box functions which do not provide derivatives (note that XSLP does allow user functions to provide their own derivatives), the cost of function evaluations may become a dominant factor in solve time. It is important to note that each second order numerical derivative costs twice as much as a first order numerical derivative, and this can make XSLP more attractive than Knitro for such problems.

8.3.2 Symbolic Differentiation

Xpress NonLinear will instead provide analytic derivatives where possible, which are both more accurate and more efficient. There are two major approaches to such calculations, and high quality implementations of both are available in this framework.

A symbolic differentiation engine calculates the derivative of an expression in closed form, using its formula representation. This is a very efficient way of recalculating individual entries of the Jacobian, and is the default approach to providing derivative information to XSLP in case the total size of all derivatives is expected to fit into memory.

8.3.3 Automatic Differentiation

An automatic differentiation engine in contrast can simultaneously compute multiple derivatives by repeated application of the chain rule. This is a very efficient means of calculating large numbers of Hessian entries, and is the default approach to providing derivative information to Knitro. It is also the default choice for SLP in case of large scale models.

8.4 Points of inflection

A point of inflection in a given variable occurs when the first and second order partial derivatives with respect to that variable become zero, but there exist nonzero derivatives of higher order. At such points, the approximations the iterative nonlinear methods create do not encapsulate enough information about the behavior of the function, and both first and second order methods may experience difficulties. For example, consider the following problem

$$\begin{array}{ll} \text{minimize} & x^3 \\ \text{subject to} & -1 \leq x \leq 1 \end{array}$$

for which the optimal solution is -1.

When the initial value of x is varied, XSLP and Knitro produce the solutions presented in Table 8.1 for this problem:

As a second order method, Knitro examines a local quadratic approximation to the function. Starting at both 0 and 1, this approximation will closely resemble the x^2 function, and so the solution will be attracted to zero. For XSLP, which is a first order method, the approximation at 0 will have a zero gradient. However, XSLP can detect this situation and will perform the analysis required to substitute an appropriate small nonzero (placeholder) value for the derivative during the first iterations. As can be seen, this allows XSLP find an optimal solution in all three cases.

Starting point:	-1	0	1
<hr/>			
Knitro :	-1	0	7.34639e-011
SLP :	-1	-1	-1

Figure 8.1: Effect of an inflection point on solution values.

This is only one example of the behaviour of these solvers without further tuning. The long steps which XSLP often takes can be both beneficial and harmful in different contexts. For example, if the function to be optimized includes many local minima, it is possible to see the opposite pattern for XSLP and Knitro. Consider

$$\begin{array}{ll} \text{minimize} & x \sin(100x^2) \\ \text{subject to} & -1 \leq x \leq 1 \end{array}$$

which has many local minima. For this problem, the results obtained are presented in Table 8.2:

Starting point:	-1	0	1
<hr/>			
Knitro :	-0.978883	0	-0.720008
SLP :	0.506366	0.506366	0.506366

Figure 8.2: Local solutions for a function with several local optima

In this case the same long steps made by XSLP lead to it finding the an identical, but unfortunate, local optimum no matter which initial point it begins from.

8.5 Trust regions

In a second order method like Knitro, there is a well-defined merit function which can be used to compare solutions, and which provides a measure of the progress being made by the algorithm. This is a significant advantage over first order methods, in which there is generally no such function.

Despite their speed and resilience to points of inflection, first order methods can also experience difficulties at points in which the current approximation is not well posed. Consider

$$\begin{array}{ll} \text{minimize} & x^2 \\ \text{subject to} & x \text{ free} \end{array}$$

at $x = 1$. A naive linearization is simply

$$\begin{array}{ll} \text{minimize} & 2x \\ \text{subject to} & x \text{ free} \end{array}$$

which is unbounded. To address such situations, XSLP will introduce trust regions to model the neighborhood in which the current approximation is believed to be applicable. When coupled with the use of derivative placeholders described in the previous section, this can lead XSLP to initially make large moves from its starting position.

CHAPTER 9

Handling Infeasibilities

By default, Xpress-SLP will include *penalty error vectors* in the augmented SLP structure. This feature adds explicit positive and negative slack vectors to all constraints (or, optionally, just to equality constraints) which include nonlinear coefficients. In many cases, this is itself enough to retain feasibility. There is also an opportunity to add penalty error vectors to all constraints, but this is not normally required.

During cascading (see next section), Xpress-SLP will ensure that the value of a cascaded variable is never set outside its lower and upper bounds (if these have been specified).

9.1 Infeasibility Analysis in the Xpress Optimizer

For problems which can be solved using the Xpress Optimizer, that is LP, convex QP and QCQP and their MIP counterparts, there is normally no difficulty with establishing feasibility. This is because for these convex problem classes, Xpress can produce globally optimal solutions, and any problem declared infeasible is globally infeasible. The concept of local infeasibility is primarily of use in the case of nonlinear problems, and in particular non-convex, nonlinear problems.

When the Xpress Optimizer declares a problem to be infeasible, the tools provided with the Xpress Optimizer console can be used to analyse the infeasibility, and hence to subsequently alter the model to overcome it. One important step in this respect is the ability to retrieve an irreducible infeasible set (IIS) (using the `iis` command). An IIS is a statement of a particular conflict in the model between a set of constraints and bounds, which make the problem certainly infeasible. An IIS is minimal in the sense that if any constraint or bound of the IIS were removed from the subproblem represented by the IIS, the resulting (relaxed) subproblem would be feasible. The Xpress Optimizer also contains a tool to identify the minimum weighted violations of constraints or bounds that would make the problem feasible (called `repairinfeas`).

Both `iis` and `repairinfeas` can be applied to any LP, convex QP, or convex QCQP problem, as well as to their mixed integer counterparts. Please refer to the Xpress Optimizer and Mosel reference manuals for more information.

9.2 Managing Infeasibility with Xpress Knitro

Xpress Knitro has three major controls which govern feasibility.

<code>XKTR_PARAM_FEASTOL</code>	This is the relative feasibility tolerance applied to a problem.
<code>XKTR_PARAM_FEASTOLABS</code>	This is the corresponding absolute feasibility tolerance.
<code>XKTR_PARAM_INFEASTOL</code>	This is the tolerance for declaring a problem infeasible.

The feasibility emphasis control, `XKTR_PARAM_BAR_FEASIBLE`, can be set for models on which Knitro has encountered difficulties in finding a feasible solution. If it is set to `get` or `get_stay`, particular

emphasis will be placed upon obtaining feasibility, rather than balancing progress toward feasibility and optimality as is the default.

If one of the built-in interior point methods is used, as determined by `XKTR_PARAM_ALGORITHM`, the feasibility emphasis control can force the iterates to strictly satisfy inequalities. It does not, however, require Knitro to satisfy all equality constraints at intermediate iterates.

The migration between a pure search for feasibility, and a balanced approach to feasibility and optimality, may be further fine tuned by using the `XKTR_PARAM_BAR_SWITCHRULE` control. Should a model still fail to converge to a feasible solution, the `XKTR_PARAM_BAR_PENCONS` control may be used to instruct Knitro to introduce penalty breakers of its own. This option has similar behaviour to the corresponding option in XSLP.

9.3 Managing Infeasibility with Xpress-SLP

There are two sources of infeasibility when XSLP is used

1. Infeasibility introduced by the error of the approximation, most noticeable when significant steps are made in the linearization.
2. Infeasibility introduced by the activation of penalty breakers, where it was not otherwise possible to make a meaningful step in the linearization.

The infeasibility induced by the former diminishes as the solution converges, provided mild assumptions regarding the continuity of the functions describing the model are satisfied. The focus of any analysis of infeasibility in XSLP must therefore most often be on the penalty breakers (also called error vectors).

For some problems, Xpress-SLP may terminate with a solution which is not sufficiently feasible for use in a desired application. The first controls to use to try to resolve such an issue are

<code>XSLP_ECFTOL_A</code>	The absolute linearization feasibility tolerance is compared for each constraint in the original, nonlinear problem to its violation by the current solution.
<code>XSLP_ECFTOL_R</code>	The relative linearization feasibility tolerance is compared for each constraint in the original, nonlinear problem to its violation by the current solution, relative to the maximum absolute value of the positive and negative contributions to the constraint.

9.4 Penalty Infeasibility Breakers in XSLP

Convergence will automatically address any errors introduced by movement within the linearization. When only small movements occur in the solution, then for differentiable functions the drift resulting from motion on the linearization is also limited.

However, it is not always possible to stay within the linearization and still make an improving step. XSLP is often able to resolve such situations automatically by the introduction of penalty infeasibility breakers. These allow the solver to violate the linearized constraints by a small amount. Such variables are associated with large cost penalties in the linearized problems, which prevents the solution process from straying too far from the approximated feasible region.

Note that if penalty breakers are required, the solution process may be very sensitive to the choice of cost penalties placed on the breakers. In most cases, XSLP's constraint analysis will automatically identify appropriate penalties as needed for each row, but for some problems additional tuning might be required.

Xpress-SLP will attempt to force all penalty breakers to zero in the limit by associating a substantial cost with them in the objective function. Such costs will be increased repeatedly should the penalty breaker

remain non-zero over a period of time. The current penalty cost for all such variables is available as `XSLP_CURRENTERRORCOST`. The control `XSLP_ERRORCOST` determines the initial value for this cost, while the `XSLP_ERRORCOSTFACTOR` controls the factor by which it increases if active error vectors remain. The maximum value of the penalty is determined by the control `XSLP_ERRORMAXCOST`. If the maximum error cost is reached, it is unlikely that XSLP will converge. It is possible in this situation to terminate the solve, by setting bit 11 of `XSLP_ALGORITHM`.

Some problems may be sensitive to the initial value of `XSLP_ERRORCOST`. If this value is too small relative to the original objective in the model, feasibility will not be sufficiently strongly encouraged during the solution process. This can cause SLP to explore highly infeasible solutions in the early stages, since the original objective will dominate any consideration of feasibility. It is even possible in this case for unboundedness of the linearizations to occur, although SLP is capable of automatic recovery from such a situation.

When the initial penalty cost is too high, the penalty term will dominate the objective. This in turn will may lead to initially low quality solutions being explored, with the attendant possibility of numerical errors accumulating. The control `XSLP_OBJTOPENALTYCOST` guides the process which selects an automatic value for `XSLP_ERRORCOST`, but determining such a value analytically can be difficult. For some difficult problems, there may be significant benefits to selecting the value directly.

Often for infeasible problems, the contribution of the individual constraints to the overall infeasibility is non-uniform. XSLP can automatically associate a weight with each row based upon the magnitude of the terms in the constraint. It is both possible to refine these weights, or alternatively to allow XSLP update them dynamically. The latter case is called escalation, and is controlled by bit 8 of `XSLP_ALGORITHM`.

Devising appropriate weights manually can be difficult, and in most cases it is preferable to leave the identification of these values to Xpress-SLP. However if it is necessary to do, the output of XSLP may provide hints as to appropriate values if detailed logging is enabled. This can be turned on with `XSLP_LOG`. The most important points in such output are the active error vectors at each iteration, where the most attractive constraints to modify are those which occur regularly in the log in association with non-zero error vectors.

CHAPTER 10

Cascading

Cascading is the process of recalculating the values of SLP variables to be more consistent with each other. The procedure involves sequencing the designated variables in order of dependence and then, starting from the current solution values, successively recalculating values for the variables, and modifying the stored solution values as required. Normal cascading is only possible if a *determining row* can be identified for each variable to be recalculated. A determining row is an equality constraint which uniquely determines the value of a variable in terms of other variables whose values are already known. Any variable for which there is no determining row will retain its original solution value. Defining a determining row for a column automatically makes the column into an SLP variable.

In extended MPS format, the SLPDATA record type "DR" is used to provide information about determining rows.

In the Xpress NonLinear function library, function `XSLPsetdetrow` allows the definition of a determining row for a column.

The cascading procedure is as follows:

- Produce an order of evaluation to ensure that variables are cascaded after any variables on which they are dependent.
- After each SLP iteration, evaluate the columns in order, updating coefficients only as required. If a determining row cannot calculate a new value for the SLP variable (for example, because the coefficient of the variable evaluates to zero), then the current value may be left unchanged, or (optionally) the previous value can be used instead.
- If a feedback loop is detected (that is, a determining row for a variable is dependent indirectly on the value of the variable), the evaluation sequence is carried out in the order in which the variables are weighted, or the order in which they are encountered if there is no explicit weighting.
- Check the step bounds, individual bounds and cascaded values for consistency. Adjust the cascaded result to ensure it remains within any explicit or implied bounds.

Normally, the solution value of a variable is exactly equal to its assumed value plus the solution value of its delta. Occasionally, this calculation is not exact (it may vary by up to the LP feasibility tolerance) and the difference may cause problems with the SLP solution path. This is most likely to occur in a quadratic problem when the quadratic part of the objective function contains SLP variables. Xpress NonLinear can re-calculate the value of an SLP variable to be equal to its assumed value plus its delta, rather than using the solution value itself.

`XSLP_CASCADE` is a bitmap which determines whether cascading takes place and whether the recalculation of solution values is extended from the use of determining rows to recalculation of the solution values for all SLP variables, based on the assumed value and the solution value of the delta.

In the following table, in the definitions under **Category**, *error* means the difference between the solution value and the assumed value plus the delta value. Bit settings in `XSLP_CASCADE` are used to determine which category of variable will have its value recalculated as follows:

Bit	Constant name	Category
0	XSLP_CASCADE_ALL	SLP variables with determining rows
1	XSLP_CASCADE_COEF_VAR	Variables appearing in coefficients where the error is greater than the feasibility tolerance
2	XSLP_CASCADE_ALL_COEF_VAR	Variables appearing in coefficients where the error is greater than 1.0E-14
3	XSLP_CASCADE_STRUCT_VAR	Variables not appearing in coefficients where the error is greater than the feasibility tolerance
4	XSLP_CASCADE_ALL_STRUCT_VAR	Variables not appearing in coefficients where the error is greater than 1.0E-14

In the presence of determining rows that include instantiated functions, SLP can attempt to group the corresponding variables together in the cascading order. This can be achieved by setting

Bit	Constant name	Effect
0	XSLP_CASCADE_SECONDARY_GROUPS	Create secondary order grouping DR rows with instantiated user functions together in the order

10.1 Determining rows and determining columns

Normally, Xpress-SLP automatically identifies if the constraint selected as determining row for a variable defines the value of the SLP variable which it determines or not. However, in certain situations, the value of a single other column determines if the determining row defines the variable or not; such a column is called the determining column for the variable.

This situation is typical when the determined and determining column form a bilinear term: $x * y + F(Z) = 0$ where y is the determined variable, Z is a set of other variables not including x or y , and F is an arbitrary function; in this case x is the determining column. These variable pairs are detected automatically. In case the absolute value of x is smaller than `XSLP_DRCOLTOL`, then variable y will not be cascaded, instead its value will be fixed and kept at its current value until the value of x becomes larger than the threshold.

Alternatively, the handling of variables for which a determining column has been identified can be customized by using a callback, see `XSLPsetcbdrcol`.

CHAPTER 11

Convergence criteria

11.1 Convergence criteria

In Xpress-SLP there are two levels of convergence criteria. On the higher level, convergence is driven by the target relative feasibility / validation control `XSLP_VALIDATIONTARGET_R`, and the target first order validation tolerance `XSLP_VALIDATIONTARGET_K`. These high level targets drive the traditional SLP convergence measures, of which there are three types for testing convergence:

- Strict convergence tests on variables
- Extended convergence tests on variables
- Convergence tests on the solution overall

11.2 Convergence overview

11.2.1 *Strict Convergence*

Three tolerances in XSLP are used to determine whether an individual variable has strictly converged, that is they describe the numerical behavior of convergence in the formal, mathematical sense.

<code>XSLP_CTOL</code>	The closure tolerance is compared against the movement of a variable relative to its initial step bound.
<code>XSLP_ATOL_A</code>	The absolute delta tolerance is compared against the absolute movement of a variable.
<code>XSLP_ATOL_R</code>	The relative delta tolerance is compared against the movement of a variable relative to its initial value.

11.2.2 *Extended Convergence*

There are six tolerances in XSLP used to determine whether an individual variable has converged according to the extended definition. These tests essentially measure the quality of the linearization, including the effect of changes to the nonlinear terms that contribute to a variable in the linearization. In order to be deemed to have converged in the extended sense, all terms in which it appears must satisfy at least one of the following:

<code>XSLP_MTOL_A</code>	The absolute matrix tolerance is compared against the approximation error relative only to the absolute value of the variable.
<code>XSLP_MTOL_R</code>	The relative matrix tolerance is compared against the approximation error relative to the size of the nonlinear term before any step is taken.

XSLP_ITOL_A	The absolute impact tolerance is compared against the approximation error of the nonlinear term.
XSLP_ITOL_R	The relative impact tolerance is compared against the approximation error relative to the positive and negative contributions to each constraint.
XSLP_STOL_A	The absolute slack impact tolerance is compared against the approximation error, but only for non-binding constraints, which is to say those for which the marginal value is small (as defined by XSLP_MVTOL).
XSLP_STOL_R	The relative slack impact tolerance is compared against the approximation error relative to the term's contribution to its constraints, but only for non-binding constraints, which is to say those for which the marginal value is small (as defined by XSLP_MVTOL).

11.2.3 Stopping Criterion

The stopping criterion requires that all variables in the problem have converged in one of the three senses. Detailed information regarding the conditions under which XSLP has terminated can be obtained from the XSLP_STATUS solver attribute. Note that a solution is deemed to have fully converged if all variables have converged in the strict sense. If all variables have converged either in the strict or extended sense, and there are no active step bounds, then the solution is called a practical solution. In contrast, the solution may be called converged if it is feasible and the objective is no longer improving.

The following four control sets can be applied by XSLP to determine whether the objective is stationary, depending on the convergence control parameter XSLP_CONVERGENCEOPS:

VTOL	This is the baseline static objective function tolerance, which is compared against the change in the objective over a given number of iterations, relative to the average objective value. Satisfaction of VTOL does not imply convergence of the variables.
XSLP_VCOUNT	This is the number of iterations over which to apply this measure of static objective convergence.
XSLP_VLIMIT	The static objective function test is applied only after at least XSLP_VLIMIT + XSLP_SBSTART XSLP iterations have taken place.
XSLP_VTOL_A	This is the absolute tolerance which is compared to the range of the objective over the last XSLP_VLIMIT iterations.
XSLP_VTOL_R	This is the tolerance used for a scaled version of the absolute test which considers the average size of the absolute value of the objective over the previous XSLP_VLIMIT iterations.
OTOL	This static objective function tolerance is applied when there are no unconverged variables in active constraints, although some variables with active step bounds might remain. It is compared to the change in the objective over a given number of iterations, relative to the average objective value.
XSLP_OCOUNT	This is the number of iterations over which to apply this measure of static objective convergence.
XSLP_OTOL_A	This is the absolute tolerance which is compared to the range of the objective over the last XSLP_OCOUNT iterations.
XSLP_OTOL_R	This is used for a scaled version of the absolute test which considers the average size of the absolute value of the objective over the previous XSLP_OCOUNT iterations.

XTOL	This static objective function tolerance is applied when a practical solution has been found. It is compared against the change in the objective over a given number of iterations, relative to the average objective value.
XSLP_XCOUNT	This is the number of iterations over which to apply this measure of static objective convergence.
XSLP_XLIMIT	This is the maximum number of iterations which can have occurred for this static objective function test to be applied. Once this number is exceeded, the solution is deemed to have converged if all the variables have converged by the strict or extended criteria.
XSLP_XTOL_A	This is the absolute tolerance which is compared to the range of the objective function over the last XSLP_XLIMIT iterations.
XSLP_XTOL_R	This is used for a scaled version of the absolute test which considers the average size of the absolute value of the objective over the last XSLP_XLIMIT iterations.
WTOL	The extended convergence continuation tolerance is applied when a practical solution has been found. It is compared to the change in the objective during the previous iteration.
XSLP_WCOUNT	This is number of iterations over which to calculate this measure of static objective convergence in the relative version of the test.
XSLP_WTOL_A	This is the absolute tolerance which is compared to the change in the objective in the previous iteration.
XSLP_WTOL_R	This is used for a scaled version of the test which considers the average size of the absolute value of the objective over the last XSLP_WCOUNT iterations.

11.2.4 Step Bounding

Step bounding in XSLP can be activated in two cases. It may be enabled adaptively in response to variable oscillation, or it may be enabled by after XSLP_SBSTART iterations, by setting XSLP_ALGORITHM appropriately. Two major controls define the behaviour of step bounds:

XSLP_SBSTART	This defines the number of iterations which must occur before XSLP may apply non-essential step bounding. When a linearization is unbounded, XSLP will introduce step bounding regardless of the value of this control.
XSLP_DEFAULTSTEPBOUND	This is the initial size of the step bounds introduced. Depending upon the value of XSLP_ALGORITHM, XSLP may use the iterations before XSLP_SBSTART to refine this initial value on a per variable basis.

11.3 Convergence: technical details

In the following sections we shall use the subscript 0 to refer to values used to build the linear approximation (the *assumed* value) and the subscript 1 to refer to values in the solution to the linear approximation (the *actual* value). We shall also use δ to indicate the change between the assumed and the actual values, so that for example:

$$\delta X = X_1 - X_0.$$

The tests are described in detail later in this section. Tests are first carried out on each variable in turn, according to the following sequence:

Strict convergence criteria:

1. **Closure tolerance (CTOL).**
This tests δX against the initial step bound of X .
2. **Delta tolerance (ATOL)**
This tests δX against X_0 .

If the strict convergence tests fail for a variable, it is tested against the extended convergence criteria:

3. **Matrix tolerance (MTOL)**
This tests whether the effect of a matrix coefficient is adequately approximated by the linearization. It tests the error against the magnitude of the effect.
4. **Impact tolerance (ITOL)**
This tests whether the effect of a matrix coefficient is adequately approximated by the linearization. It tests the error against the magnitude of the contributions to the constraint.
5. **Slack impact tolerance (STOL)**
This tests whether the effect of a matrix coefficient is adequately approximated by the linearization and is applied only if the constraint has a negligible marginal value (that is, it is regarded as "not constraining"). The test is the same as for the impact tolerance, but the tolerance values may be different.

The three extended convergence tests are applied simultaneously to all coefficients involving the variable, and each coefficient must pass at least one of the tests if the variable is to be regarded as converged. If any coefficient fails the test, the variable has not converged.

Regardless of whether the variable has passed the system convergence tests or not, if a convergence callback function has been set using `XSLPsetcbitervar` then it is called to allow the user to determine the convergence status of the variable.

6. **User convergence test**
This test is entirely in the hands of the user and can return one of three conditions: the variable has converged on user criteria; the variable has not converged; or the convergence status of the variable is unchanged from that determined by the system.

Once the tests have been completed for all the variables, there are several possibilities for the convergence status of the solution:

- (a) All variables have converged on strict criteria or user criteria.
- (b) All variables have converged, some on extended criteria, and there are no active step bounds (that is, there is no delta vector which is at its bound and has a significant reduced cost).
- (c) All variables have converged, some on extended criteria, and there are active step bounds (that is, there is at least one delta vector which is at its bound and has a significant reduced cost).
- (d) Some variables have not converged, but these have non-constant coefficients only in constraints which are not active (that is, the constraints do not have a significant marginal value);
- (e) Some variables have not converged, and at least one has a non-constant coefficient in an active constraint (that is, the constraint has a significant marginal value);

If (a) is true, then the solution has converged on *strict convergence criteria*.

If (b) is true, then the solution has converged on *extended convergence criteria*.

If (c) is true, then the solution is a *practical* solution. That is, the solution is an optimal solution to the linearization and, within the defined tolerances, it is a solution to the original nonlinear problem. It is

possible to accept this as the solution to the nonlinear problem, or to continue optimizing to see if a better solution can be obtained.

If (d) or (e) is true, then the solution has not converged. Nevertheless, there are tests which can be applied to establish whether the solution can be regarded as converged, or at least whether there is benefit in continuing with more iterations.

The first convergence test on the solution simply tests the variation in the value of the objective function over a number of SLP iterations:

7. Objective function convergence test 1 (VTOL)

This test measures the range of the objective function (the difference between the maximum and minimum values) over a number of SLP iterations, and compares this against the magnitude of the average objective function value. If the range is small, then the solution is deemed to have converged.

Notice that this test says nothing about the convergence of the variables. Indeed, it is almost certain that the solution is not in any sense a practical solution to the original nonlinear problem. However, experience with a particular type of problem may show that the objective function does settle into a narrow range quickly, and is a good indicator of the ultimate *value* obtained. This test can therefore be used in circumstances where only an estimate of the solution value is required, not how it is made up. One example of this is where a set of schedules is being evaluated. If a quick estimate of the value of each schedule can be obtained, then only the most profitable or economical ones need be examined further.

If the convergence status of the variables is as in (d) above, then it may be that the solution is practical and can be regarded as converged:

8. Objective function convergence test 2 (XTOL)

If there are no unconverged values in active constraints, then the inaccuracies in the linearization (at least for small errors) are not important. If a constraint is not active, then deleting the constraint does not change the feasibility or optimality of the solution. The convergence test measures the range of the objective function (the difference between the maximum and minimum values) over a number of SLP iterations, and compares this against the magnitude of the average objective function value. If the range is small, then the solution is deemed to have converged.

The difference between this test and the previous one is the requirement for the convergence status of the variables to be (d).

Unless test 7 (VTOL) is being applied, if the convergence status of the variables is (e) then the solution has not converged and another SLP iteration will be carried out.

If the convergence status is (c), then the solution is practical. Because there are active step bounds in the solution, a "better" solution would be obtained to the linearization if the step bounds were relaxed. However, the linearization becomes less accurate the larger the step bounds become, so it might not be the case that a better solution would also be achieved for the nonlinear problem. There are two convergence tests which can be applied to decide whether it is worth continuing with more SLP iterations in the hope of improving the solution:

9. Objective function convergence test 3 (OTOL)

If all variables have converged (even if some are converged on extended criteria only, and some of those have active step bounds), the solution is a practical one. If the objective function has not changed significantly over the last few iterations, then it is reasonable to suppose that the solution will not be significantly improved by continuing with more SLP iterations. The convergence test measures the range of the objective function (the difference between the maximum and minimum values) over a number of SLP iterations, and compares this against the magnitude of the average objective function value. If the range is small, then the solution is deemed to have converged.

10. Extended convergence continuation test (WTOL)

Once a solution satisfying (c) has been found, we have a practical solution against which to compare solution values from later SLP iterations. As long as there has been a significant improvement in the objective function, then it is worth continuing. If the objective function over the last few iterations has failed to improve over the practical solution, then the practical solution is restored and the solution is deemed to have converged.

The difference between tests 9 and 10 is that 9 (OTOL) tests for the objective function being stable, whereas 10 (WTOL) tests whether it is actually improving. In either case, if the solution is deemed to have converged, then it has converged to a practical solution.

11.3.1 Closure tolerance (CTOL)

If an initial step bound is provided for a variable, then the closure test measures the significance of the magnitude of the delta compared to the magnitude of the initial step bound. More precisely:

Closure test:

$$ABS(\delta X) \leq B * XSLP_CTOL$$

where B is the initial step bound for X. If no initial step bound is given for a particular variable, the closure test is not applied to that variable, even if automatic step bounds are applied to it during the solution process.

If a variable passes the closure test, then it is deemed to have converged.

11.3.2 Delta tolerance (ATOL)

The simplest tests for convergence measure whether the actual value of a variable in the solution is significantly different from the assumed value used to build the linear approximation.

The absolute test measures the significance of the magnitude of the delta; the relative test measures the significance of the magnitude of the delta compared to the magnitude of the assumed value. More precisely:

Absolute delta test:

$$ABS(\delta X) \leq XSLP_ATOL_A$$

Relative delta test:

$$ABS(\delta X) \leq X_0 * XSLP_ATOL_R$$

If a variable passes the absolute or relative delta tests, then it is deemed to have converged.

11.3.3 Matrix tolerance (MTOL)

The matrix tests for convergence measure the linearization error in the effect of a coefficient. The *effect* of a coefficient is its value multiplied by the activity of the column in which it appears.

$$E = V * C$$

where V is the activity of the matrix column in which the coefficient appears, and C is the value of the coefficient. The linearization approximates the effect of the coefficient as

$$E = V * C_0 + \delta X * C'_0$$

where V is as before, C_0 is the value of the coefficient C calculated using the assumed values for the variables and C'_0 is the value of $\frac{\partial C}{\partial X}$ calculated using the assumed values for the variables.

The error in the effect of the coefficient is given by

$$\delta E = V_1 * C_1 - (V_1 * C_0 + \delta X * C'_0)$$

Absolute matrix test:

$$ABS(\delta E) \leq XSLP_MTOL_A$$

Relative matrix test:

$$ABS(\delta E) \leq V_0 * X_0 * XSLP_MTOL_R$$

If all the coefficients which involve a given variable pass the absolute or relative matrix tests, then the variable is deemed to have converged.

11.3.4 Impact tolerance (ITOL)

The impact tests for convergence also measure the linearization error in the effect of a coefficient. The effect of a coefficient was described in the previous section. Whereas the matrix test compares the error against the magnitude of the coefficient itself, the impact test compares the error against a measure of the magnitude of the constraint in which it appears. All the elements of the constraint are examined: for each, the contribution to the constraint is evaluated as the element multiplied by the activity of the vector in which it appears; it is then included in a *total positive contribution* or *total negative contribution* depending on the sign of the contribution. If the predicted effect of the coefficient is positive, it is tested against the total positive contribution; if the effect of the coefficient is negative, it is tested against the total negative contribution.

As in the matrix tests, the predicted effect of the coefficient is

$$V * C_0 + \delta X * C'_0$$

and the error is

$$\delta E = V_1 * C_1 - (V_1 * C_0 + \delta X * C'_0)$$

Absolute impact test:

$$ABS(\delta E) \leq XSLP_ITOL_A$$

Relative impact test:

$$ABS(\delta E) \leq T_0 * XSLP_ITOL_R$$

where

$$T_0 = ABS\left(\sum_{v \in V} v_0 * c_0\right)$$

c is the value of the constraint coefficient in the vector v ; V is the set of vectors such that $v_0 * c_0 > 0$ if E is positive, or the set of vectors such that $v_0 * c_0 < 0$ if E is negative.

If a coefficient passes the matrix test, then it is deemed to have passed the impact test as well. If all the coefficients which involve a given variable pass the absolute or relative impact tests, then the variable is deemed to have converged.

11.3.5 Slack impact tolerance (STOL)

This test is identical in form to the impact test described in the previous section, but is applied only to constraints whose marginal value is less than $XSLP_MVTOL$. This allows a weaker test to be applied where the constraint is not, or is almost not, binding.

Absolute slack impact test:

$$ABS(\delta E) \leq XSLP_STOL_A$$

Relative slack impact test:

$$ABS(\delta E) \leq T_0 * XSLP_STOL_R$$

where the items in the expressions are as described in the previous section, and the tests are applied only when

$$ABS(\pi_i) < XSLP_MVTOL$$

where π_i is the marginal value of the constraint.

If all the coefficients which involve a given variable pass the absolute or relative matrix, impact or slack impact tests, then the variable is deemed to have converged.

11.3.6 Fixed variables due to determining columns smaller than threshold (FX)

Variables having a determining column, that are temporarily fixed due to the absolute value of the determining column being smaller than the threshold `XSLP_DRCOLTOL` are regarded as converged.

11.3.7 User-defined convergence

Regardless of what the Xpress-SLP convergence tests have said about the status of an individual variable, it is possible for the user to set the convergence status for a variable by using a function defined through the `XSLPsetcbitervar` callback registration procedure. The callback function returns an integer result *S* which is interpreted as follows:

$S < 0$	mark variable as unconverged
$S = 0$	leave convergence status of variable unchanged
$S \geq 11$	mark variable as converged with status <i>S</i>

Values of *S* in the range 1 to 10 are interpreted as meaning convergence on the standard system-defined criteria.

If a variable is marked by the user as converged, it is treated as if it has converged on strict criteria.

11.3.8 Static objective function (1) tolerance (VTOL)

This test does not measure convergence of individual variables, and in fact does not in any way imply that the solution has converged. However, it is sometimes useful to be able to terminate an optimization once the objective function appears to have stabilized. One example is where a set of possible schedules are being evaluated and initially only a good estimate of the likely objective function value is required, to eliminate the worst candidates.

The variation in the objective function is defined as

$$\delta Obj = MAX_{Iter}(Obj) - MIN_{Iter}(Obj)$$

where *Iter* is the `XSLP_VCOUNT` most recent SLP iterations and *Obj* is the corresponding objective function value.

Absolute static objective function (3) test:

$$ABS(\delta Obj) \leq XSLP_VTOL_A$$

Relative static objective function (3) test:

$$ABS(\delta Obj) \leq AVG_{Iter}(Obj) * XSLP_VTOL_R$$

The static objective function (3) test is applied only after at least `XSLP_VLIMIT` + `XSLP_SBSTART` SLP iterations have taken place. Where step bounding is being applied, this ensures that the test is not applied until after step bounding has been introduced.

If the objective function passes the relative or absolute static objective function (3) test then the solution will be deemed to have converged.

11.3.9 Static objective function (2) tolerance (OTOL)

This test does not measure convergence of individual variables. Instead, it measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables interacting with active constraints (those that have a marginal value of at least `XSLP_MVTOL`) have converged. The rationale is that if the remaining unconverged variables are not involved in active constraints and if the objective function is not changing significantly between iterations, then the solution is more-or-less practical.

The variation in the objective function is defined as

$$\delta Obj = MAX_{Iter}(Obj) - MIN_{Iter}(Obj)$$

where `Iter` is the `XSLP_OCOUNT` most recent SLP iterations and `Obj` is the corresponding objective function value.

Absolute static objective function (2) test:

$$ABS(\delta Obj) \leq XSLP_OTOL_A$$

Relative static objective function (2) test:

$$ABS(\delta Obj) \leq AVG_{Iter}(Obj) * XSLP_OTOL_R$$

If the objective function passes the relative or absolute static objective function (2) test then the solution is deemed to have converged.

11.3.10 Static objective function (3) tolerance (XTOL)

It may happen that all the variables have converged, but some have converged on extended criteria (MTOL, ITOL or STOL) and at least one of these is at its step bound. It is therefore possible that an improved result could be obtained by taking another SLP iteration. However, if the objective function has already been stable for several SLP iterations, then there is less likelihood of an improved result, and the converged solution can be accepted.

The static objective function (1) test measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables have converged, but some have converged on extended criteria (MTOL, ITOL or STOL) and at least one of these is at its step bound. Because all the variables have converged, the solution is already converged but the fact that some variables are at their step bound limit suggests that the objective function could be improved by going further.

The variation in the objective function is defined as

$$\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$$

where Iter is the XSLP_XCOUNT most recent SLP iterations and Obj is the corresponding objective function value.

Absolute static objective function (1) test:

$$\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_XTOL_A}$$

Relative static objective function (1) test:

$$\text{ABS}(\delta\text{Obj}) \leq \text{AVG}_{\text{Iter}}(\text{Obj}) * \text{XSLP_XTOL_R}$$

The static objective function (1) test is applied only until XSLP_XLIMIT SLP iterations have taken place. After that, if all the variables have converged on strict or extended criteria, the solution is deemed to have converged.

If the objective function passes the relative or absolute static objective function (1) test then the solution is deemed to have converged.

11.3.11 Extended convergence continuation tolerance (WTOL)

This test is applied after a converged solution has been found where at least one variable has converged on extended criteria and is at its step bound limit. As described under XTOL above, it is possible that by continuing with additional SLP iterations, the objective function might improve. The extended convergence continuation test measures whether any improvement is being achieved. If not, then the last converged solution will be restored and the optimization will stop.

For a maximization problem, the improvement in the objective function at the current iteration compared to the objective function at the last converged solution is given by:

$$\delta\text{Obj} = \text{Obj} - \text{ConvergedObj}$$

(for a minimization problem, the sign is reversed).

Absolute extended convergence continuation test:

$$\delta\text{Obj} > \text{XSLP_WTOL_A}$$

Relative extended convergence continuation test:

$$\delta\text{Obj} > \text{ABS}(\text{ConvergedObj}) * \text{XSLP_WTOL_R}$$

A solution is deemed to have a significantly better objective function value than the converged solution if δObj passes the relative *and* absolute extended convergence continuation tests.

When a solution is found which converges on extended criteria and with active step bounds, the solution is saved and SLP optimization continues until one of the following:

- a new solution is found which converges on some other criterion, in which case the SLP optimization stops with this new solution
- a new solution is found which converges on extended criteria and with active step bounds, and which has a significantly better objective function, in which case this is taken as the new saved solution
- none of the XSLP_WCOUNT most recent SLP iterations has a significantly better objective function than the saved solution, in which case the saved solution is restored and the SLP optimization stops

CHAPTER 12

Xpress-SLP Structures

12.1 SLP Matrix Structures

Xpress-SLP augments the original matrix to include additional rows and columns to model some or all of the variables involved in nonlinear relationships, together with first-order derivatives.

The amount and type of augmentation is determined by the bit map control variable
XSLP_AUGMENTATION:

Bit 0	Minimal augmentation. All SLP variables appearing in coefficients or matrix entries are provided with a corresponding update row and delta vector.
Bit 1	Even-handed augmentation. All nonlinear expressions are converted into terms. All SLP variables are provided with a corresponding update row and delta vector.
Bit 2	Create penalty error vectors (+ and -) for each equality row of the original problem containing a nonlinear coefficient or term. This can also be implied by the setting of bit 3.
Bit 3	Create penalty error vectors (+ and/or - as required) for each row of the original problem containing a nonlinear coefficient or term. Setting bit 3 to 1 implies the setting of bit 2 to 1 even if it is not explicitly carried out.
Bit 4	Create additional penalty delta vectors to allow the solution to exceed the step bounds at a suitable penalty.
Bit 8	Implement step bounds as constraint rows.
Bit 9	Create error vectors (+ and/or - as required) for each constraining row of the original problem.

If Bits 0-1 are not set, then Xpress-SLP will use standard augmentation: all SLP variables (appearing in coefficients or matrix entries, or variables with non constant coefficients) are provided with a corresponding update row and delta vector.

To avoid too many levels of super- and sub- scripting, we shall use X , Y and Z as variables, $F()$ as a function, and R as the row name. In the matrix structure, column and row names are shown *in italics*.

X_0 is the current estimate ("assumed value") of X . $F'_x(\dots)$ is the first derivative of F with respect to X .

12.1.1 Augmentation of a nonlinear coefficient

Original matrix structure

	X
R	$F(Y, Z)$

Matrix structure: minimal augmentation (XSLP_AUGMENTATION=1)

	X	Y	Z	dY	dZ	
R	$F(Y_0, Z_0)$			$X_0 * F'_y(Y_0, Z_0)$	$X_0 * F'_z(Y_0, Z_0)$	
uY		1		-1		$= Y_0$
uZ			1		-1	$= Z_0$

The original nonlinear coefficient (X, R) is replaced by its evaluation using the assumed values of the independent variables.

Two vectors and one equality constraint for each independent variable in the coefficient are created if they do not already exist.

The new vectors are:

- The SLP variable (e.g. Y)
- The SLP delta variable (e.g. dY)

The new constraint is the SLP update row (e.g. uY) and is always an equality. The only entries in the update row are the +1 and -1 for the SLP variable and delta variable respectively. The right hand side is the assumed value for the SLP variable.

The entry in the original nonlinear constraint row for each independent variable is the first-order partial derivative of the implied term $X * F(Y, Z)$, evaluated at the assumed values.

The delta variables are bounded by the current values of the corresponding step bounds.

Matrix structure: standard augmentation (XSLP_AUGMENTATION=0)

	X	Y	Z	dX	dY	dZ	
R	$F(Y_0, Z_0)$				$X_0 * F'_y(Y_0, Z_0)$	$X_0 * F'_z(Y_0, Z_0)$	
uX	1			-1			$= X_0$
uY		1			-1		$= Y_0$
uZ			1			-1	$= Z_0$

The original nonlinear coefficient (X, R) is replaced by its evaluation using the assumed values of the independent variables.

Two vectors and one equality constraint for each independent variable in the coefficient are created if they do not already exist.

The new vectors are:

- The SLP variable (e.g. Y)
- The SLP delta variable (e.g. dY)

The new constraint is the SLP update row (e.g. uY) and is always an equality. The only entries in the update row are the +1 and -1 for the SLP variable and delta variable respectively. The right hand side is the assumed value for the SLP variable.

The entry in the original nonlinear constraint row for each independent variable is the first-order partial derivative of the implied term $X * F(Y, Z)$, evaluated at the assumed values.

The delta variables are bounded by the current values of the corresponding step bounds.

One new vector and one new equality constraint are created for the variable containing the nonlinear coefficient.

The new vector is:

- The SLP delta variable (e.g. dX)

The new constraint is the SLP update row (e.g. uX) and is always an equality. The only entries in the update row are the +1 and -1 for the original variable and delta variable respectively. The right hand side is the assumed value for the original variable.

The delta variable is bounded by the current values of the corresponding step bounds.

Matrix structure: even-handed augmentation (XSLP_AUGMENTATION=2)

	=	X	Y	Z	dX	dY	dZ	
R	$X_0 * F(Y_0, Z_0)$				$F(Y_0, Z_0)$	$X_0 * F'_Y(Y_0, Z_0)$	$X_0 * F'_Z(Y_0, Z_0)$	
uX		1			-1			= X_0
uY			1			-1		= Y_0
uZ				1			-1	= Z_0

The coefficient is treated as if it was the term $X * F(Y, Z)$ and is expanded in the same way as a *nonlinear term*.

12.1.2 Augmentation of a nonlinear term

Original matrix structure

	=
R	$F(X, Y, Z)$

The column name = is a reserved name for a column which has a fixed activity of 1.0 and can conveniently be used to hold nonlinear terms, particularly those which cannot be expressed as coefficients of variables.

Matrix structure: all augmentations

	=	X	Y	Z	dX	dY	dZ	
R	$F(X_0, Y_0, Z_0)$				$F'_X(X_0, Y_0, Z_0)$	$F'_Y(X_0, Y_0, Z_0)$	$F'_Z(X_0, Y_0, Z_0)$	
uX		1			-1			= X_0
uY			1			-1		= Y_0
uZ				1			-1	= Z_0

The original nonlinear coefficient (=,R) is replaced by its evaluation using the assumed values of the independent variables.

Two vectors and one equality constraint for each independent variable in the coefficient are created if they do not already exist.

The new vectors are:

- The SLP variable (e.g. Y)
- The SLP delta variable (e.g. dY)

The new constraint is the SLP update row (e.g. uY) and is always an equality. The only entries in the update row are the +1 and -1 for the SLP variable and delta variable respectively. The right hand side is the assumed value for the SLP variable.

The entry in the original nonlinear constraint row for each independent variable is the first-order partial derivative of the term $F(X, Y, Z)$, evaluated at the assumed values.

The delta variables are bounded by the current values of the corresponding step bounds.

One new vector and one new equality constraint are created for the variable containing the nonlinear coefficient.

The new vector is:

- The SLP delta variable (e.g. dX)

The new constraint is the SLP update row (e.g. uX) and is always an equality. The only entries in the update row are the +1 and -1 for the original variable and delta variable respectively. The right hand side is the assumed value for the original variable.

The delta variable is bounded by the current values of the corresponding step bounds.

Note that if $F(X,Y,Z) = X \cdot F(Y,Z)$ then this translation is exactly equivalent to that for the nonlinear coefficient described earlier.

12.1.3 Augmentation of a user-defined SLP variable

Typically, this will arise when a variable represents the result of a nonlinear function, and is required to converge, or to be constrained by step-bounding to force convergence. In essence, it would arise from a relationship of the form

$$X = F(Y, Z)$$

Original matrix structure

	=	X
R	$F(Y, Z)$	-1

Matrix structure: all augmentations

	=	X	Y	Z	dX	dY	dZ	
R	$F(Y_0, Z_0)$	-1				$F'_Y(Y_0, Z_0)$	$F'_Z(Y_0, Z_0)$	
uX		1			-1			= X_0
uY			1			-1		= Y_0
uZ				1			-1	= Z_0

The Y, Z structures are identical to those which would result from a nonlinear term or coefficient. The X, dX and uX structures effectively define dX as the deviation of X from X_0 which can be controlled with step bounds.

The augmented and even-handed structures include more delta vectors, and so allow for more measurement and control of convergence.

Type of structure	Minimal	Standard	Even-handed
Type of variable			
Variables in nonlinear coefficients	Y	Y	Y
Variables with nonlinear coefficients	N	Y	Y
User-defined SLP variable	Y	Y	Y
Nonlinear term	Y	Y	Y

Y SLP variable has a delta vector which can be measured and/or controlled for convergence.

N SLP variable does not have a delta and cannot be measured and/or controlled for convergence.

There is no mathematical difference between the augmented and even-handed structures.

The even-handed structure is more elegant because it treats all variables in an identical way. However, the original coefficients are lost, because their effect is transferred to the "=" column as a term and so it is not possible to look up the coefficient value in the matrix after the SLP solution process has finished (whether because it has converged or because it has terminated for some other reason). The values of the SLP variables are still accessible in the usual way.

Some of the extended convergence criteria will be less effective because the effects of the individual coefficients may be amalgamated into one term (so, for example, the total positive and negative contributions to a constraint are no longer available).

12.1.4 SLP penalty error vectors

Bits 2, 3 and 9 of control variable `XSLP_AUGMENTATION` determine whether SLP penalty error vectors are added to constraints. Bit 9 applies penalty error vectors to all constraints; bits 2 and 3 apply them only to constraints containing nonlinear terms. When bit 2 or bit 3 is set, two penalty error vectors are added to each such equality constraint; when bit 3 is set, one penalty error vector is also added to each such inequality constraint. The general form is as follows:

Original matrix structure

$$\begin{array}{c} \text{=} \\ \hline R \quad F(Y, Z) \end{array}$$

Matrix structure with error vectors

	<i>X</i>	<i>R+</i>	<i>R-</i>
<i>R</i>	$F(Y, Z)$	+1	-1
<i>P_ERROR</i>		+Weight	+Weight

For equality rows, two penalty error vectors are added. These have penalty weights in the penalty error row `P_ERROR`, whose total is transferred to the objective with a cost of `XSLP_CURRENTERRORCOST`. For inequality rows, only one penalty error vector is added — the one corresponding to the slack is omitted. If any error vectors are used in a solution, the transfer cost from the cost penalty error row will be increased by a factor of `XSLP_ERRORCOSTFACTOR` up to a maximum of `XSLP_ERRORMAXCOST`.

Error vectors are ignored when calculating cascaded values.

The presence of error vectors at a non-zero level in an SLP solution normally indicates that the solution is not self-consistent and is therefore not a solution to the nonlinear problem.

Control variable `XSLP_ERRORTOL_A` is a tolerance on error vectors. Any error vector with a value less than `XSLP_ERRORTOL_A` will be regarded as having a value of zero.

Bit 9 controls whether error vectors are added to all constraints. If bit 9 is set, then error vectors are added in the same way as for the setting of bit 3, but to all constraints regardless of whether or not they have nonlinear coefficients.

12.2 Xpress-SLP Matrix Name Generation

Xpress-SLP adds rows and columns to the nonlinear problem in order to create a linear approximation. The new rows and columns are given names derived from the row or column to which they are related as follows:

Row or column type	Control parameter containing format	Default format
Update row	XSLP_UPDATEFORMAT	pU_r
Delta vector	XSLP_DELTAFORMAT	pD_c
Penalty delta (below step bound)	XSLP_MINUSDELTAFORMAT	pD-c
Penalty delta (above step bound)	XSLP_PLUSDELTAFORMAT	pD+c
Penalty error (below RHS)	XSLP_MINUSERERRORFORMAT	pE-r
Penalty error (above RHS)	XSLP_PLUSERERRORFORMAT	pE+r
Row for total of all penalty vectors (error or delta)	XSLP_PENALTYROWFORMAT	pPR_x
Column for standard penalty cost (error or delta)	XSLP_PENALTYCOLFORMAT	pPC_x
LO step bound formulated as a row	XSLP_SBLOROWFORMAT	pSB-c
UP step bound formulated as a row	XSLP_SBUPROWFORMAT	pSB+c

In the default formats:

p	a unique prefix (one or more characters not used as the beginning of any name in the problem).
r	the original row name.
c	the original column name.
x	The penalty row and column vectors are suffixed with "ERR" or "DELT" (for error and delta respectively).

Other characters appear "as is".

The format of one of these generated names can be changed by setting the corresponding control parameter to a formatting string using standard "C"-style conventions. In these cases, the unique prefix is not available and the only obvious choices, apart from constant names, use "%s" to include the original name — for example:

U_%s would create names like U_abcdefghi

U_%-8s would create names like U_abcdefgh (always truncated to 8 characters).

You can use a part of the name by using the XSLP_*OFFSET control parameters (such as XSLP_UPDATEOFFSET) which will offset the start of the original name by the number of characters indicated (so, setting XSLP_UPDATEOFFSET to 1 would produce the name U_bdefghi).

12.3 Xpress-SLP Statistics

When a matrix is read in using XSLPreadprob, statistics on the model are produced. They should be interpreted as described in the numbered footnotes:

Reading Problem xxx	(1)
Problem Statistics	
1920 (0 spare) rows	(2)
899 (0 spare) structural columns	(3)
6683 (3000 spare) non-zero elements	(4)
MIP Entity Statistics	
0 entities 0 sets 0 set members	(5)
Xpress-SLP Statistics:	

3632 coefficients	(6)
14 extended variable arrays	(7)
1 user functions	(8)
1011 SLP variables	(9)

Notes:

1. Standard output from XPRSreadprob reading the linear part of the problem
2. Number of rows declared in the ROWS section
3. Number of columns with at least one constant coefficient
4. Number of constant elements
5. Integer and SOS statistics if appropriate
6. Number of non-constant coefficients
7. Number of XVs defined
8. Number of user functions defined
9. Number of variables identified as SLP variables (interacting with a non-linear coefficient)

When the original problem is augmented prior to optimization, the following statistics are produced:

Xpress-SLP Augmentation Statistics:	
Columns:	
754 implicit SLP variables	(10)
1010 delta vectors	(11)
2138 penalty error vectors (1177 positive, 961 negative)	(12)
Rows:	
1370 nonlinear constraints	(13)
1010 update rows	(14)
1 penalty error rows	(15)
Coefficients:	
11862 non-constant coefficients	(16)

Notes:

10. SLP variables appearing only in coefficients and having no constant elements
11. Number of delta vectors created
12. Numbers of penalty error vectors
13. Number of constraints containing nonlinear terms
14. Number of update rows (equals number of delta vectors)
15. Number of rows totaling penalty vectors (error or delta)
16. Number of non-constant coefficients in the linear augmented matrix

If the matrix is read in using the XPRSloadxxx and XSLPloadxxx functions then these statistics may not be produced. However, most of the values are accessible through Xpress NonLinear integer attributes using the XSLPgetintattrib function.

12.4 SLP Variable History

Xpress-SLP maintains a history value for each SLP variable. This value indicates the direction in which the variable last moved and the number of consecutive times it moved in the same direction. All variables start with a history value of zero.

Current History	Change in activity of variable	New History
0	>0	1
0	<0	-1
>0	>0	No change unless delta vector is at its bound. If it is, then new value is Current History + 1
>0	<0	-1
<0	<0	No change unless delta vector is at its bound. If it is, then new value is Current History - 1
<0	>0	1
anything	0	No change

Tests of variable movement are based on comparison with absolute and relative (and, if set, closure) tolerances. Any movement within tolerance is regarded as zero.

If the new absolute value of History exceeds the setting of `XSLP_SAMECOUNT`, then the step bound is reset to a larger value (determined by `XSLP_EXPAND`) and History is reset as if it had been zero.

If History and the change in activity are of opposite signs, then the step bound is reset to a smaller value (determined by `XSLP_SHRINK`) and History is reset as if it had been zero.

With the default settings, History will normally be in the range -1 to -3 or +1 to +3.

CHAPTER 13

Xpress NonLinear Formulae

Xpress NonLinear can handle formulae described in three different ways:

Character strings The formula is written exactly as it would appear in, for example, the Extended MPS format used for text file input.

Internal unparsed format The tokens within the formula are replaced by a {tokentype, tokenvalue} pair. The list of types and values is in the table below.

Internal parsed format The tokens are converted as in the unparsed format, but the order is changed so that the resulting array forms a reverse-Polish execution stack for direct evaluation by the system.

13.1 Parsed and unparsed formulae

All formulae input into Xpress NonLinear are parsed into a reverse-Polish execution stack. Tokens are identified by their type and a value. The table below shows the values used in interface functions.

All formulae are provided in the interface functions as two parallel arrays:

- an integer array of token types;
- a double array of token values.

The last token type in the array should be an end-of-formula token (XSLP_EOF, which evaluates to zero).

If the value required is an integer, it should still be provided in the array of token values as a double precision value.

Even if a token type requires no token value, it is best practice to initialize such values as zeros. In particular, any other provided values will generally not be preserved when queried later.

Type	Description	Value
XSLP_COL	column	index of matrix column.
XSLP_CON	constant	(double) value.
XSLP_DEL	delimiter	XSLP_COMMA (1) = comma (",") XSLP_COLON (2) = colon (":")
XSLP_EOF	end of formula	not required: use zero
XSLP_FUN	user function	index of function
XSLP_IFUN	internal function	index of function
XSLP_LB	left bracket	not required: use zero
XSLP_OP	operator	XSLP_UMINUS (1) = unary minus ("-") XSLP_EXPONENT (2) = exponent ("**" or "^") XSLP_MULTIPLY (3) = multiplication ("*") XSLP_DIVIDE (4) = division ("/") XSLP_PLUS (5) = addition ("+") XSLP_MINUS (6) = subtraction ("-")
XSLP_RB	right bracket	not required: use zero

When a formula is passed to Xpress NonLinear in "internal unparsed format" — that is, with the formula already converted into tokens — the full range of token types is permitted.

When a formula is passed to Xpress NonLinear in "parsed format" — that is, in reverse Polish — the following rules apply:

XSLP_DEL	comma is optional.
XSLP_FUN	implies a following left-bracket, which is not included explicitly.
XSLP_IFUN	implies a following left-bracket, which is not included explicitly.
XSLP_LB	never used.
XSLP_RB	only used to terminate the list of arguments to a function.

Brackets are not used in the reverse Polish representation of the formula: the order of evaluation is determined by the order of the items on the stack. Functions which need the brackets — for example `XPRSslpgetcoefstr` — fill in brackets as required to achieve the correct evaluation order. The result may not match the formula as originally provided.

13.2 Example of an arithmetic formula

$$x^2 + 4y(z - 3)$$

Written as an unparsed formula, each token is directly transcribed as follows:

Type	Value
XSLP_COL	index of x
XSLP_OP	XSLP_EXPONENT
XSLP_CON	2
XSLP_OP	XSLP_PLUS
XSLP_CON	4
XSLP_OP	XSLP_MULTIPLY
XSLP_COL	index of y
XSLP_OP	XSLP_MULTIPLY
XSLP_LB	0
XSLP_COL	index of z
XSLP_OP	XSLP_MINUS
XSLP_CON	3
XSLP_RB	0
XSLP_EOF	0

Written as a parsed formula (in reverse Polish), an evaluation order is established first, for example:

`x 2 ^ 4 y * z 3 - * +`

and this is then transcribed as follows:

Type	Value
XSLP_COL	index of x
XSLP_CON	2
XSLP_OP	XSLP_EXPONENT
XSLP_CON	4
XSLP_COL	index of y
XSLP_OP	XSLP_MULTIPLY
XSLP_COL	index of z
XSLP_CON	3
XSLP_OP	XSLP_MINUS
XSLP_OP	XSLP_MULTIPLY
XSLP_OP	XSLP_PLUS
XSLP_EOF	0

Notice that the brackets used to establish the order of evaluation in the unparsed formula are not required in the parsed form.

13.3 Example of a formula involving a simple function

$y * \text{MyFunc}(z, 3)$

Written as an unparsed formula, each token is directly transcribed as follows:

Type	Value
XSLP_COL	index of y
XSLP_OP	XSLP_MULTIPLY
XSLP_FUN	index of MyFunc
XSLP_LB	0
XSLP_COL	index of z
XSLP_DEL	XSLP_COMMA
XSLP_CON	3
XSLP_RB	0
XSLP_EOF	0

Written as a parsed formula (in reverse Polish), an evaluation order is established first, for example:

$y \) \ 3 \ , \ z \ \text{MyFunc}(\ *$

and this is then transcribed as follows:

Type	Value
XSLP_COL	index of y
XSLP_RB	0
XSLP_CON	3
XSLP_DEL	XSLP_COMMA
XSLP_COL	index of z
XSLP_FUN	index of MyFunc
XSLP_OP	XSLP_MULTIPLY
XSLP_EOF	0

Notice that the function arguments are in reverse order, and that a right bracket is used as a delimiter to indicate the end of the argument list. The left bracket indicating the start of the argument list is implied by the XSLP_FUN token.

CHAPTER 14

User Functions

14.1 Callbacks and user functions

Callbacks and user functions both provide mechanisms for connecting user-written functions to Xpress NonLinear. However, they have different capabilities and are not interchangeable.

A *callback* is called at a specific point in the SLP optimization process (for example, at the start of each SLP iteration). It has full access to all the problem data and can, in principle, change the values of any items — although not all such changes will necessarily be acted upon immediately or at all.

A *user function* is essentially the same as any other mathematical function, used in a formula to calculate the current value of a coefficient. The function is called when a new value is needed; for efficiency, user functions are not usually called if the value is already known (for example, when the function arguments are the same as on the previous call). Therefore, there is no guarantee that a user function will be called at any specific point in the optimization procedure or at all.

Although a user function is normally free-standing and needs no access to problem or other data apart from that which it receives through its argument list, there are facilities to allow it to access the problem and its data if required. The following limitations should be observed:

1. The function should not make use of any variable data which is not in its list of arguments;
2. The function should not change any of the problem data.

The reasons for these restrictions are as follows:

1. Xpress NonLinear determines which variables are linked to a formula by examining the list of variables and arguments to functions in the formula. If a function were to access and use the value of a variable not in this list, then incorrect relationships would be established, and incorrect or incomplete derivatives would be calculated. The predicted and actual values of the coefficient would then always be open to doubt.
2. Xpress NonLinear generally allows problem data to be changed between function calls, and also by callbacks called from within an Xpress NonLinear function. However, user functions are called at various points during the optimization and no checks are generally made to see if any problem data has changed. The effects of any such changes will therefore at best be unpredictable.

For a description of how to access the problem data from within a user function, see the section on "More complicated user functions" later in this chapter.

14.2 User function interface

In its simplest form, a user function is exactly the same as any other mathematical function: it takes a set of arguments (constants or values of variables) and returns a value as its result. In this form, which is

the usual implementation, the function needs no information apart from the values of its arguments. It is possible to create more complicated functions which do use external data in some form.

Xpress NonLinear distinguishes six different types of user functions.

- A user function is called a `map`, if it takes and returns a single value.
- A user function is called a `mapvec`, if it takes an array of inputs, and returns a single evaluation value.
- A user function is called a `multimap`, if it takes an array of inputs, and also returns an array of evaluation values.
- A `mapdelta` user function is an extended version of a `map` that also returns its own partial derivatives when requested
- A `mapvecdelta` user function is an extended version of a `mapvec` that also returns its own partial derivatives when requested
- A `multimapdelta` user function is an extended version of a `multimap` that also returns its own partial derivatives when requested

14.3 User Function declaration in native languages

This section describes how to declare a user function in C. The general shape of the declaration is shown. Not all the possible arguments will necessarily be used by any particular function, and the actual arguments required will depend on the way the function is declared to Xpress NonLinear.

14.3.1 User function declaration in C

If the function is placed in a library, `XSLPimportlibfunc` may be used to retrieve a pointer to be passed to `XSLPadduserfunction`.

A user function can be included in the executable program which calls Xpress NonLinear.

A `multimapdelta` function's `deltas` is an array with the same number of items as `inputvalues`. It is used as an indication of which derivatives (if any) are required on a particular function call. If `deltas[i]` is zero then a derivative for input variable `i` is not required and must not be returned. If `deltas[i]` is nonzero then a derivative for input variable `i` is required and must be returned.

When no derivatives are calculated, the array of return values simply contains the results.

When derivatives are calculated, the array contains the values and the derivatives as follows (DV_i is the i^{th} variable for which derivatives are required, which may not be the same as the i^{th} input value):

```
Result1
Derivative of Result1 w.r.t. DV1
Derivative of Result1 w.r.t. DV2
...
Derivative of Result1 w.r.t. DVn
Result2
Derivative of Result2 w.r.t. DV1
Derivative of Result2 w.r.t. DV2
...
Derivative of Result2 w.r.t. DVn
...
Derivative of Resultm w.r.t. DVn
```

It is therefore important to check whether derivatives are required and, if so, how many.

The return value of the user functions that return an int (as opposed to the evaluation value) is a status code indicating whether the function has completed normally. Possible values are:

- 0 No errors: the function has completed normally.
- 1 The function has encountered an error. This will terminate the optimization.
- 1 The calling function must estimate the function value from the last set of values calculated. This will cause an error if no values are available.

14.4 Programming Techniques for User Functions

This section is principally concerned with the programming of large or complicated user functions, perhaps taking a potentially large number of input values and calculating a large number of results. However, some of the issues raised are also applicable to simpler functions.

The first part describes in more detail some of the possible arguments to the function. The remainder of the section looks at function instances, function objects and direct calls to user functions.

14.4.1 Deltas

The `Deltas` array has the same dimension as `InputValues` and is used to indicate which of the input variables should be used to calculate derivatives. If `Deltas[i]` is zero, then no derivative should be returned for input variable *i*. If `Deltas[i]` is nonzero, then a derivative is required for input variable *i*. The value of `Deltas[i]` can be used as a suggested perturbation for numerical differentiation (a negative sign indicates that if a one-sided derivative is calculated, then a backward one is preferred). If derivatives are calculated analytically, or without requiring a specific perturbation, then `Deltas` can be interpreted simply as an array of flags indicating which derivatives are required.

14.4.2 Return values and ReturnArray

The `ReturnArray` array is provided for those user functions which return more than one value, either because they do calculate more than one result, or because they also calculate derivatives. The function must either return the address of an array which holds the values, or pass the values to the calling program through the `ReturnArray` array.

The total number of values returned depends on whether derivatives are being calculated. The total number of values (and hence the minimum size of the array) is $nRet * (nDeriv + 1)$. Xpress NonLinear guarantees that `ReturnArray` will be large enough to hold the total number of values requested.

A function which calculates and returns a single value can use the `ReturnArray` array provided that the declarations of the function in Xpress NonLinear and in the native language both include the appropriate argument definition.

functions which use the `ReturnArray` array must also return a status code as their return value. Zero is the normal return value. A nonzero value is an error code which will cause any formula evaluation to stop and will normally interrupt any optimization or other procedure.

14.4.3 Returning Derivatives

A multi-valued function which does not calculate its own derivatives will return its results as a one-dimensional array.

As already described, when derivatives are calculated as well, the order is changed, so that the required derivatives follow the value for each result. That is, the order becomes:

$A, \frac{\partial A}{\partial X_1}, \frac{\partial A}{\partial X_2}, \dots, \frac{\partial A}{\partial X_n}, B, \frac{\partial B}{\partial X_1}, \frac{\partial B}{\partial X_2}, \dots, \frac{\partial B}{\partial X_n}, \dots, \frac{\partial Z}{\partial X_n}$
 where A, B, Z are the return values, and X_1, X_2, X_n , are the input (independent) variables (in order) for which derivatives have been requested.

Not all calls to a user function necessarily require derivatives to be calculated. Check `FunctionInfo` for the number of derivatives required (it will be zero if only a value calculation is needed), and `Deltas` for the indications as to which independent variables are required to produce derivatives. Xpress NonLinear will not ask for, nor will it expect to receive, derivatives for function arguments which are actually constant in a particular problem. A function which provides uncalled-for derivatives will cause errors in subsequent calculations and may cause other unexpected side-effects if it stores values outside the expected boundaries of the return array.

14.4.4 Function Instances

Xpress NonLinear defines an *instance* of a user function to be a unique combination of function and arguments. For functions which return an array of values, the specific return argument is ignored when determining instances. Thus, given the following formulae:

$$\begin{aligned} &f(x) + f(y) + g(x, y : 1) \\ &f(y) * f(x) * g(x, y : 2) \\ &f(z) \end{aligned}$$

the following instances are created:

$f(x)$
 $f(y)$
 $f(z)$
 $g(x, y)$

(A function reference of the form $g(x, y : n)$ means that g is a multi-valued function of x and y , and we want the n^{th} return value.)

Xpress NonLinear regards as *complicated* any user function which returns more than one value, which uses input or return names, or which calculates its own derivatives. All complicated functions give rise to function instances, so that each function is called only once for each distinct combination of arguments.

Functions which are not regarded as complicated are normally called each time a value is required.

Note that conditional re-evaluation of the function is only possible if it generates function instances.

Using function instances can improve the performance of a problem, because the function is called only once for each combination of arguments, and is not re-evaluated if the values have not changed significantly. If the function is computationally intensive, the improvement can be significant.

Xpress NonLinear normally expects to obtain a set of partial derivatives from a user function at a particular base-point and then to use them as required, depending on the evaluation settings for the various functions. If for any reason this is not appropriate, then the integer control parameter `XSLP_EVALUATE` can be set to 1, which will force re-evaluation every time.

A function instance is not re-evaluated if all of its arguments are unchanged.

A simple function which does not have a function instance is evaluated every time.

If `XSLP_EVALUATE` is not set, then it is still possible to by-pass the re-evaluation of a function if the values have not changed significantly since the last evaluation. If the input values to a function have all converged to within their strict convergence tolerance (`CTOL`, `ATOL_A`, `ATOL_R`), and bit 4 of `XSLP_FUNCEVAL` is set to 1, then the existing values and derivatives will continue to be used. At the option of the user, an individual function, or all functions, can be re-evaluated in this way or at each SLP iteration. If a function is not re-evaluated, then all the required values will be calculated from the base point and the partial derivatives; the input and return values used in making the original function calculation are unchanged.

Bits 3-5 of integer control parameter `XSLP_FUNCEVAL` determine the nature of function evaluations. The meaning of each bit is as follows:

- Bit 3** evaluate functions whenever independent variables change.
- Bit 4** evaluate functions when independent variables change outside tolerances.
- Bit 5** apply evaluation mode to all functions.

If bits 3-4 are zero, then the settings for the individual functions are used.

If bit 5 is zero, then the settings in bits 3-4 apply only to functions which do not have their own specific evaluation modes set.

Examples:

Bits 3-5 = 1 (set bit 3) Evaluate functions whenever their input arguments (independent variables) change, unless the functions already have their own evaluation options set.

Bits 3-5 = 5 (set bits 3 and 5) Evaluate all functions whenever their input arguments (independent variables) change.

Bits 3-5 = 6 (set bits 4 and 5) Evaluate functions whenever input arguments (independent variables) change outside tolerance. Use existing calculation to estimate values otherwise.

Bits 6-8 of integer control parameter XSLP_FUNCEVAL determine the nature of derivative calculations. The meaning of each bit is as follows:

- Bit 6** tangential derivatives.
- Bit 7** forward derivatives.
- Bit 8** apply evaluation mode to all functions.

If bits 6-7 are zero, then the settings for the individual functions are used.

If bit 8 is zero, then the settings in bits 6-7 apply only to functions which do not have their own specific derivative calculation modes set.

Examples:

Bits 6-8 = 1 (set bit 6) Use tangential derivatives for all functions which do not already have their own derivative options set.

Bits 6-8 = 5 (set bits 6 and 8) Use tangential derivatives for all functions.

Bits 6-8 = 6 (set bits 7 and 8) Use forward derivatives for all functions.

The following constants are provided for setting these bits:

Setting bit 3 XSLP_RECASC
 Setting bit 4 XSLP_TOLCALC
 Setting bit 5 XSLP_ALLCALCS
 Setting bit 6 XSLP_2DERIVATIVE
 Setting bit 7 XSLP_1DERIVATIVE
 Setting bit 8 XSLP_ALLDERIVATIVES

When analytical derivatives are used, for user functions not returning their own derivatives, SLP will calculate approximated derivatives using finite differences for instantiated functions and use these values when deriving analytical derivatives.

Management of zero placeholder entries

15.1 The augmented matrix structure

During the augmentation process, Xpress-SLP builds additional matrix structure to represent the linear approximation of the nonlinear constraints within the problem (see Xpress-SLP Structures). In effect, it adds a generic structure which approximates the effect of changes to variables in nonlinear expressions, over and above that which would apply if the variables were simply replaced by their current values.

As a very simple example, consider the nonlinear constraint (R1, say)

$$X * Y \leq 10$$

The variables X and Y are replaced by $X_0 + \delta X$ and $Y_0 + \delta Y$ respectively, where X_0 and Y_0 are the values of X and Y at which the approximation will be made.

The original constraint is therefore

$$(X_0 + \delta X) * (Y_0 + \delta Y) \leq 10$$

Expanding this into individual terms, we have

$$X_0 * Y_0 + X_0 * \delta Y + Y_0 * \delta X + \delta X * \delta Y \leq 10$$

The first term is constant, the next two terms are linear in δY and δX respectively, and the last term is nonlinear.

The augmented structure deletes the nonlinear term, so that the remaining structure is a linear approximation to the original constraint. The justification for doing this is that if δX or δY (or both) are small, then the error involved in ignoring the term is also small.

The resulting matrix structure has entries of Y_0 in the delta variable δX and X_0 in the delta variable δY . The constant entry $X_0 * Y_0$ is placed in the special "equals" column which has a fixed activity of 1. All these entries are updated at each SLP iteration as the solution process proceeds and the problem is linearized at a new point. The positions of these entries – (R1, δX), (R1, δY) and (R1, =) – are known as *placeholders*.

15.2 Derivatives and zero derivatives

At each SLP iteration, the values of the placeholders are re-calculated. In the example in the previous section, the values X_0 in the delta variable δY and Y_0 in the delta variable δX were effectively determined by analytic methods – that is, we differentiated the original formula to determine what values would be required in the placeholders.

In general, analytic differentiation may not be possible: the formula may contain functions which cannot be differentiated (because, for example, they are not smooth or not continuous), or for which the analytic derivatives are not known (because, for example, they are functions providing values from "black boxes" such as databases or simulators). In such cases, Xpress-SLP approximates the differentiation process by numerical methods. The example in the previous section would have approximate derivatives calculated as follows:

The current value of X (X_0) is perturbed by a small amount (dX), and the value of the formula is recalculated in each case.

$$f_d = (X_0 - dX) * Y_0$$

$$f_u = (X_0 + dX) * Y_0$$

$$\text{derivative} = (f_u - f_d) / (2 * dX)$$

In this particular example, the value obtained by numerical methods is the same as the analytic derivative. For more complex functions, there may be a slight difference, depending on the magnitude of dX .

This derivative represents the effect on the constraint of a change in the value of X . Obviously, if Y changes as well, then the combined effect will not be fully represented although, in general, it will be directionally correct.

The problem comes when Y_0 is zero. In such a case, the derivative is calculated as zero, meaning that changing X has no effect on the value of the formula. This can impact in one of two ways: either the value of X never changes because there is no incentive to do so, or it changes by unreasonably large amounts because there is no effect from doing so. If X and Y are linked in some other way, so that Y becomes nonzero when X changes, the approximation using zero as the derivative can cause the optimization process to behave badly.

Xpress-SLP tries to avoid the problem of zero derivatives by using small nonzero values for variables which are in fact zero. In most cases this gives a small nonzero value for the derivative, and hence for the placeholder entry. The model then contains some effect for the change in a variable, even if instantaneously the effect is zero.

The same principle is applied to analytic derivatives, so that the values obtained by either method are broadly similar.

15.3 Placeholder management

The default action of Xpress-SLP is to retain all the calculated values for all the placeholder entries. This includes values which would be zero without the special handling described in the previous section. We will call such values "zero placeholders".

Although retaining all the values gives the best chance of finding a good optimum, the presence of a large dense area of small values often gives rise to considerable numerical instability which adversely affects the optimization process. Xpress-SLP therefore offers a way of deleting small values which is less likely to affect the final outcome whilst improving numerical stability.

Most of the candidate placeholders are in the delta variables (represented by the δX and δY variables above). Various criteria can be selected for deletion of zero placeholder entries without affecting the validity of the basis (and so making the next SLP iteration more costly in time and stability). The criteria are selected using the control parameter `XSLP_ZEROCRITERION` as follows:

- **Bit 0 (=1)** Remove placeholders in nonbasic SLP variables
This criterion applies to placeholders which are in the SLP variable (not the delta). Any value can be deleted from a nonbasic variable without upsetting the basis, so all eligible zero placeholders can be deleted.
- **Bit 1 (=2)** Remove placeholders in nonbasic delta variables
Any value can be deleted from a nonbasic variable without upsetting the basis, so all eligible zero placeholders can be deleted.
- **Bit 2 (=4)** Remove placeholders in a basic SLP variable if its update row is nonbasic
If the update row is nonbasic, then generally the basic SLP variable can be pivoted in the update row, so the basis is still valid if other entries are deleted. The entry in the update row is always 1.0 and will never be deleted.

- **Bit 3 (=8)** Remove placeholders in a basic delta variable if its update row is nonbasic and the corresponding SLP variable is nonbasic
If the delta is basic and the corresponding SLP variable is nonbasic, then the delta will pivot in the update row (the delta and the SLP variable are the only two variables in the update row), so the basis is still valid if other entries are deleted. The entry in the update row is always -1.0 and will never be deleted.
- **Bit 4 (=16)** Remove placeholders in a basic delta variable if the determining row for the corresponding SLP variable is nonbasic
If the delta variable is basic and the determining row for the corresponding SLP variable is nonbasic then it is generally possible (although not 100% guaranteed) to pivot the delta variable in the determining row. so the basis is still valid if other entries are deleted. The entry in the determining row is never deleted even if it is otherwise eligible.

The following constants are provided for setting these bits:

Setting bit 0 `XSLP_ZEROCRTIERION_NBSLPVAR`
 Setting bit 1 `XSLP_ZEROCRTIERION_NBDELTA`
 Setting bit 2 `XSLP_ZEROCRTIERION_SLPVARNBUPDATEROW`
 Setting bit 3 `XSLP_ZEROCRTIERION_DELTANBUPDATEROW`
 Setting bit 4 `XSLP_ZEROCRTIERION_DELTANBDRROW`

There are two additional control parameters used in this procedure:

- `XSLP_ZEROCRITERIONSTART`
This is the first SLP iteration at which zero placeholders will be examined for eligibility. Use of this parameter allows a balance to be made between optimality and numerical stability.
- `XSLP_ZEROCRITERIONCOUNT`
This is the number of consecutive SLP iterations that a placeholder is a zero placeholder before it is deleted. So, if in the earlier example `XSLP_ZEROCRITERIONCOUNT` is 2, the entry in the delta variable `dX` will be deleted only if `Y` was also zero on the previous SLP iteration.

Regardless of the basis status of a variable, its delta, update row and determining row, if a zero placeholder was deleted on the previous SLP iteration, it will always be deleted in the current SLP iteration (keeping a zero matrix entry at zero does not upset the basis).

If the optimization method is barrier, or the basis is not being used, then the bit settings of `XSLP_ZEROCRITERION` are not used as such: if `XSLP_ZEROCRITERION` is nonzero, all zero placeholders will be deleted subject to `XSLP_ZEROCRITERIONCOUNT` and `XSLP_ZEROCRITERIONSTART`.

Special Types of Problem

16.1 Nonlinear objectives

Xpress NonLinear works with nonlinear constraints. If a nonlinear objective is required (except for the special case of a quadratic objective — see below) then the objective should be provided using a constraint in the problem. For example, to optimize $f(x)$ where f is a nonlinear function and x is a set of one or more variables, create the constraint

$$f(x) - X = 0$$

where X is a new variable, and then optimize X .

In general, x should be made a free variable, so that the problem does not converge prematurely on the basis of an unchanging objective function. It is generally important that the objective is not artificially constrained (for example, by bounding x) because this can distort the solution process. Also, as such an objective transfer row is not a real constraint, no error vectors should be added (row can be enforced); feasibility should be provided by the transfer variable x being free.

16.2 Convex Quadratic Programming

Convex quadratic programming (QP) is a special case of nonlinear programming where the constraints are linear but the objective is quadratic (that is, it contains only terms which are constant, variables multiplied by a constant, or products of two variables multiplied by a constant) and convex (convexity is checked by the Xpress Optimizer). It is possible to solve convex quadratic problems using SLP, but it is not usually the best way. The reason is that the solution to a convex QP problem is typically not at a vertex. In SLP a non-vertex solution is achieved by applying step bounds to create additional constraints which surround the solution point, so that ultimately the solution has been obtained within suitable tolerances. Because of the nature of the problem, successive solutions will often swing from one step bound to the other; in such circumstances, the step bounds are reduced on each SLP iteration but it will still take a long time before convergence. In addition, unless the linear approximation is adequately constrained, it will be unbounded because the linear approximation will not recognize the change in direction of the relationship with the derivative as the variable passes through a stationary point. The easiest way to ensure that the linear problem is constrained is to provide realistic upper and lower bounds on all variables.

In Xpress NonLinear, convex quadratic problems can be solved using the quadratic optimizer within the Xpress optimizer package. For pure QP (or MIQP) problems, therefore, SLP is not required. However, the SLP algorithm can be used together with QP to solve problems with a quadratic objective and also nonlinear constraints. The constraints are handled using the normal SLP techniques; the objective is handled by the QP optimizer. If the objective is not convex (not semi-definite), the QP optimizer may not give a solution (with default settings, it will produce an error message); SLP will find a solution but — as always — it may be a local optimum.

If a QP problem is to be solved, then the quadratic component should be input in the normal way (using `QMATRIX` or `QUADOBJ` in MPS file format, or the library functions `XPRSloadqp` or `XPRSloadmiqp`). Xpress NonLinear will then automatically use the QP optimizer. If the problem is to be solved using the SLP routines throughout, then the objective should be provided via a constraint as described in the previous section.

This applies to quadratically constrained (QCQP and MIQCQP) problems as well.

For a description on when it's more beneficial to use the XPRS library to solve QP or QCQP problems, please see *Selecting the right algorithm for a nonlinear problem - when to use the XPRS library instead of XSLP*.

16.3 Mixed Integer Nonlinear Programming

Mixed Integer Non-Linear Programming (MINLP) is the application of mixed integer techniques to the solution of problems including non-linear relationships. Xpress NonLinear offers a set of components to implement MINLP using Mixed Integer Successive Linear Programming (MISLP).

16.3.1 Mixed Integer SLP

The mixed integer successive linear programming (MISLP) solver is a generalization of the traditional branch and bound procedure to nonlinear programming. The MIP engine is used to control the branch-and-bound algorithm, with each node being evaluated using SLP. MIP then compares the SLP solutions at each node to decide which node to explore next, and to decide when an integer feasible and ultimately optimal solution have been obtained.

MISLP, also known as SLP within MIP, offers nonlinear specific root heuristics controlled by control `XSLP_HEURSTRATEGY`.

Other generic heuristics are controlled by the respective XPRS heuristics controls.

The branch and bound tree exploration is executed in parallel. Use the XPRS control `MIPTHREADS` to limit the number of threads used.

Normally, the relaxed problem is solved first, using `XSLPnloptimize` with the `-1` flag to ignore the integer elements of the problem. It is also possible to call the `XSLPnloptimize` routine with the `-g` flag and allow it to do the initial SLP optimization as well. In either case, ensure that the control parameter `XSLP_OBJSENSE` is set to `+1` (minimization) or `-1` (maximization) before calling `XSLPnloptimize`.

The actual algorithm employed is controlled by a number of control parameters, as well as offering the possibility of direct user interaction through call-backs at key points in the solution process.

16.3.2 Heuristics for Mixed Integer SLP

For hard MINLP problems, or where a solution must quickly be generated, the root heuristics of MISLP can be executed as stand alone methods. These approaches can be used by changing the value of the control parameter `XSLP_MIPALGORITHM`.

there are two MISLP heuristics:

1. MIP within SLP. In this, each SLP iteration is optimized using MIP to obtain an integer optimal solution to the linear approximation of the original problem. SLP then compares this MIP solution to the MIP solution of the previous SLP iteration and determines convergence based on the differences between the successive MIP solutions.
2. SLP then MIP. In this, SLP is used to find a converged solution to the relaxed problem. The resulting

linearization is then fixed (i.e. the base point and the partial derivatives do not change) and MIP is run to find an integer optimum. SLP is then run again to find a converged solution to the original problem with these integer settings.

The approach described in (1) seems potentially dangerous, in that changes in the integer variables could have disproportionate effects on the solution and on the values of the SLP variables. There are also question-marks over the use of step-bounding to control convergence, particularly if any of the integer variables are also SLP variables.

The approach described in (2) has the big advantage that MIP is working on a linear problem and so can take advantage of all of the special attributes of such a problem. This means that the solution time is likely to be much faster than the alternatives. However, if the real problem is significantly non-linear, the integer solution to the initial SLP solution may not be a good integer solution to the original problem and so a false optimum may occur.

16.3.3 Fixing or relaxing the values of the SLP variables

The solution process may involve step-bounding to obtain the converged solution. Some MIP solution strategies may want to fix the values of some of the SLP variables before moving on to the MIP part of the process, or they may want to allow the child nodes more freedom than would be allowed by the final settings of the step bounds. Control parameters `XSLP_MIPALGORITHM`, `XSLP_MIPFIXSTEPBOUNDS` and `XSLP_MIPRELAXSTEPBOUNDS` can be used to free, or fix to zero, various categories of step bounds, thus effectively freeing the SLP variables or fixing them to their values in the initial solution.

At each node, step bounds may again be fixed to zero or relaxed or left in the same state as in the solution to the parent node.

`XSLP_MIPALGORITHM` uses bits 2-3 (for the root node) and 4-5 (for other nodes) to determine which step bounds are fixed to zero (thus fixing the values of the corresponding variables) or freed (thus allowing the variables to change, possibly beyond the point they were restricted to in the parent node). Set bit 2 (4) of `XSLP_MIPALGORITHM` to implement relaxation of defined categories of step bounds as determined by `XSLP_MIPRELAXSTEPBOUNDS` at the root node (at each node). Set bit 3 (5) of `XSLP_MIPALGORITHM` to implement fixing of defined categories of step bounds as determined by `XSLP_MIPFIXSTEPBOUNDS` at the root node (at each node).

Alternatively, specific actions on setting bounds can be carried out by the user callback defined by `XSLPsetcbprenode`.

The default setting of `XSLP_MIPALGORITHM` is 17 which relaxes step bounds at all nodes except the root node. The step bounds from the initial SLP optimization are retained for the root node.

`XSLP_MIPRELAXSTEPBOUNDS` and `XSLP_MIPFIXSTEPBOUNDS` are bitmaps which determine which categories of SLP variables are processed.

- | | |
|-------|--|
| Bit 1 | Process SLP variables which do not appear in coefficients but which do have coefficients (constant or variable) in the original problem. |
| Bit 2 | Process SLP variables which have coefficients (constant or variable) in the original problem. |
| Bit 3 | Process SLP variables which appear in coefficients but which do not have coefficients (constant or variable) in the original problem. |
| Bit 4 | Process SLP variables which appear in coefficients. |

In most cases, the default settings (`XSLP_MIPFIXSTEPBOUNDS=0`, `XSLP_MIPRELAXSTEPBOUNDS=15`) are appropriate.

16.3.4 Iterating at each node

Any number of SLP iterations can be carried out at each node. The maximum number is set by control parameter `XSLP_MIPITERLIMIT` and is activated by `XSLP_MIPALGORITHM`. The significant values for `XSLP_MIPITERLIMIT` are:

- 0 Perform an LP optimization with the current linearization. This means that, subject to the step bounds, the SLP variables can take on other values, but the coefficients are not updated.
- 1 As for 0, but the model is updated after each iteration, so that each node starts with a new linearization based on the solution of its parent.
- $n > 1$ Perform up to n SLP iterations, but stop when a termination criterion is satisfied. If no other criteria are set, the SLP will terminate on `XSLP_ITERLIMIT` or `XSLP_MIPITERLIMIT` iterations, or when the SLP converges.

After the last MIP node has been evaluated and the MIP procedure has terminated, the final solution can be re-optimized using SLP to obtain a converged solution. This is only necessary if the individual nodes are being terminated on a criterion other than SLP convergence.

16.3.5 Termination criteria at each node

Because the intention at each node is to get a reasonably good estimate for the SLP objective function rather than to obtain a fully converged solution (which is only required at the optimum), it may be possible to set looser but practical termination criteria. The following are provided:

Testing for movement of the objective function

This functions in a similar way to the extended convergence criteria for ordinary SLP convergence, but does not require the SLP variables to have converged in any way. The test is applied once step bounding has been applied (or `XSLP_SBSTART` SLP iterations have taken place if step bounding is not being used). The node will be terminated at the current iteration if the range of the objective function values over the last `XSLP_MIPOCOUNT` SLP iterations is within `XSLP_MIPOTOL_A` or within `XSLP_MIPOTOL_R * OBJ` where `OBJ` is the average value of the objective function over those iterations.

Related control parameters:

<code>XSLP_MIPOTOL_A</code>	Absolute tolerance
<code>XSLP_MIPOTOL_R</code>	Relative tolerance
<code>XSLP_MIPOCOUNT</code>	Number of SLP iterations over which the movement is measured

Testing the objective function against a cutoff

If the objective function is worse by a defined amount than the best integer solution obtained so far, then the SLP will be terminated (and the node will be cut off). The node will be cut off at the current SLP iteration if the objective function for the last `XSLP_MIPCUTOFFCOUNT` SLP iterations are all worse than the best obtained so far, and the difference is greater than `XSLP_MIPCUTOFF_A` and `XSLP_MIPCUTOFF_R * OBJ` where `OBJ` is the best integer solution obtained so far.

Related control parameters:

<code>XSLP_MIPCUTOFF_A</code>	Absolute amount by which the objective function is worse
<code>XSLP_MIPCUTOFF_R</code>	Relative amount by which the objective function is worse
<code>XSLP_MIPCUTOFFCOUNT</code>	Number of SLP iterations checked
<code>XSLP_MIPCUTOFFLIMIT</code>	Number of SLP iterations before which the cutoff takes effect

16.3.6 Callbacks

User callbacks are provided as follows:

```
XSLPsetcbintsol(XSLPprob Prob,
               int (*UserFunc)(XSLPprob myProb, void *myObject),
               void *Object);
```

UserFunc is called when an integer solution has been obtained. The return value is ignored.

```
XSLPsetcboptnode(XSLPprob Prob,
                 int (*UserFunc)(XSLPprob myProb, void *myObject, int *feas),
                 void *Object);
```

UserFunc is called when an optimal solution is obtained at a node.

If the feasibility flag *feas is set nonzero or if the function returns a nonzero value, then further processing of the node will be terminated (it is declared infeasible).

```
XSLPsetcbprenode(XSLPprob Prob,
                 int (*UserFunc)(XSLPprob myProb, void *myObject, int *feas),
                 void *Object);
```

UserFunc is called at the beginning of each node after the SLP problem has been set up but before any SLP iterations have taken place.

If the feasibility flag *feas is set nonzero or if the function returns a nonzero value, then the node will be declared infeasible and cut off. In particular, the SLP optimization at the node will not be performed.

```
XSLPsetcbslpnode(XSLPprob Prob,
                 int (*UserFunc)(XSLPprob myProb, void *myObject, int *feas),
                 void *Object);
```

UserFunc is called after each SLP iteration at each node, after the SLP iteration, and after the convergence and termination criteria have been tested.

If the feasibility flag *feas is set nonzero or if the function returns a nonzero value, then the node will be declared infeasible and cut off.

16.4 Integer and semi-continuous delta variables

Functions implementing piecewise linear expressions often lead to local stalling due to the partial derivatives not capturing the true nature of the behaviour of the function. Such functions are often implemented as user functions or expressions using the abs function. To provide Xpress with a better way of evaluating such expressions, it is possible to mark variables (typically the key dependencies of the expression) as having a semi-continuous delta variable with a minimum perturbation size associated, which means the value of any expression that involves this variable is expected to meaningfully change if the variable's value in the current solution is changed by at least the semi-continuous bound of the delta. If a minimum meaningful perturbation is not known, the variable's delta may be set up to being of type explore, when SLP will trial several values up to the provided maximum in case zero partials are detected. Using exploration deltas may significantly increase the number of times the formulas the variable is used in are evaluated.

It is important to note that the value with a semi-continuous delta will still be allowed to take any value and make arbitrary steps between iterations, the extra information of the delta variable is solely used as a means of better evaluating the effect of change per variable.

User functions that can only be evaluated at given values (e.g. lookup tables or simulations over integer input) may be modelled with variables with an integer delta variable. If a variable's delta variable is flagged as being integer, with a step value of 'delta', then assuming the variable has an initial value of

'x0', the possible values of the variable are ' $x0 + i * \text{delta}$ ' where 'i' is an integer number. If no initial value is provided, the lower bound (or zero if no lower bound) is used to start the possible values from.

Variables with a semi-continuous delta are not expected to make the problem harder, in fact, the extra information usually aids the solve noticeably.

A model with variables with integer deltas is considered to be hard. An integer delta is expected to be used to model the domain of user functions, and should not be used to otherwise model integrality of the original variable. Variables with an integer delta used in constraints tend to make the problem difficult to solve unless their use is balanced by the presence of infeasibility breaker variables (penalty slacks).

To change the type of a delta variable, use 'XSLPchgdelatetype' in the API and the 'setdelatetype' method in Mosel.

If variables with integer deltas are present in the problem, then SLP will run a number of heuristics as part of the solve, please refer to XSLP_GRIDHEURSELECT.

CHAPTER 17

Xpress NonLinear multistart

The feature is an additive feature that minimizes the development overhead and effort of implementing parallel multistart searches. The purpose of multistart is two-fold. Traditionally, multistart is a so called "globalization" feature. It is important to correctly understand what this technology offers, and what it does not. It offers a convenient and efficient way of exploring a larger feasible space building on top of existing local solver algorithms by the means of perturbing initial points and/or parameters or even the problem statement itself. Multistarts increase the chance that one of the found local optima is indeed the global one. However, multistarts do not guarantee that a global optimum is found and never provide a proof of global optimality. For solving MINLPs to proven global optimality, consider using [FICO Xpress Global](#).

Multistart can also be viewed as a left-alone feature. In a typical situation, versions of a model react favourably to a set of control settings, dependent on data. Multistart allows for a simple way of combining different control setting scenarios, increasing the robustness of the model.

The initial problem is defined as the baseline: as the model is normally loaded it without any multistart information, including problem description, callbacks and controls. A run or a job is defined as a problem instance that needs to be solved as part of multistart.

On completion, the initial problem is set up to match that of the winner, allowing examination of the winning strategy and solution using the normal means.

The initial prob object is not reused, all runs are made on a copy of the problem, allowing full customization from the callbacks, including changes to structure.

Callbacks are inherited by the multistart jobs from the initial problem and can be customized from the multistart callbacks. `XSLPInterrupt` has a global scope, and calling it terminates the multistart search.

Although not intended as the primary use, multistart allows the execution of all supported problem classes, so for example alternate MIP strategies can be used in parallel.

The multistart job pool is maintained and can be extended until the first call to `optimize / nloptimize` with `XSLP_MULTISTART` on. This allows for doing optimizations runs aimed at generating multistart jobs. The multistart pool is dynamic and new jobs can be added on the fly from the `jobstart` and `jobend` callbacks.

III. Reference

CHAPTER 18

Reference Documentation by Topic/Functionality

This section lists all functions, controls, and attributes of FICO XPress Nonlinear 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 nonlinear solver.

Every function, control or attribute in this section is displayed with all their related topic areas.

18.1 Bit-vector

Reference section for all bit-vector controls.

18.1.1 Bit-vector library functions

XSLPchgrowstatus, XPRSslpchgrowstatus	Change the status setting of a constraint [Bit-vector, Data Input, SLP]	p. 300
XSLPgetrowstatus, XPRSslpgetrowstatus	Retrieve the status setting of a constraint [Bit-vector, SLP, Solution]	p. 340

18.1.2 Bit-vector controls

XSLP_ALGORITHM, SLPALGORITHM	Bit map describing the SLP algorithm(s) to be used [Bit-vector, SLP]	p. 216
XSLP_ANALYZE, SLPANALYZE	Bit map activating additional options supporting model / solution path analysis [Bit-vector, Logging, SLP]	p. 218
XSLP_AUGMENTATION, SLPAUGMENTATION	Bit map describing the SLP augmentation method(s) to be used [Bit-vector, SLP]	p. 219
XSLP_CONTROL	Bit map describing which Xpress NonLinear functions also activate the corresponding Optimizer Library function [Bit-vector]	p. 225
XSLP_CONVERGENCEOPS, SLPCONVERGENCEOPS	Bit map describing which convergence tests should be carried out [Bit-vector, SLP, SLP-convergence]	p. 225
XSLP_FILTER, SLPFILTER	Bit map for controlling solution updates [Bit-vector, SLP, Solution]	p. 231

XSLP_MIPALGORITHM, SLP_MIPALGORITHM	Bitmap describing the MISLP algorithms to be used [Bit-vector, MISLP]	p. 238
XSLP_PRESOLVEOPS, NLPPRESOLVEOPS	Bitmap indicating the SLP presolve actions to be taken [Bit-vector, Presolve]	p. 247
XSLP_TRACEMASKOPS, SLPTRACEMASKOPS	Controls the information printed for XSLP_TRACEMASK. [Bit-vector, Logging, SLP]	p. 253
XSLP_ZEROCRITERION, SLPZEROCRITERION	Bitmap determining the behavior of the placeholder deletion procedure [Bit-vector, SLP]	p. 259

18.2 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.

18.2.1 Branching controls

XKTR_PARAM_MIP_BRANCHRULE, KNITRO_PARAM_MIP_BRANCHRULE	Specifies which branching rule to use for MIP branch and bound procedure. [Branching, Knitro-MINLP]	p. 467
XKTR_PARAM_MIP_GUB_BRANCH, KNITRO_PARAM_MIP_GUB_BRANCH	Specifies whether or not to branch on generalized upper bounds (GUBs). [Branching, Knitro-MINLP]	p. 468
XKTR_PARAM_MIP_PSEUDOINIT, KNITRO_PARAM_MIP_PSEUDOINIT	Specifies the method used to initialize pseudo-costs corresponding to variables that have not yet been branched on in the MIP method. [Branching, Knitro-MINLP]	p. 471
XKTR_PARAM_MIP_STRONG_CANDLIM, KNITRO_PARAM_MIP_STRONG_CANDLIM	Specifies the maximum number of candidates to explore for MIP strong branching. [Branching, Knitro-MINLP, Limits]	p. 472
XKTR_PARAM_MIP_STRONG_LEVEL, KNITRO_PARAM_MIP_STRONG_LEVEL	Specifies the maximum number of tree levels on which to perform MIP strong branching. [Branching, Knitro-MINLP, Limits]	p. 472
XKTR_PARAM_MIP_STRONG_MAXIT, KNITRO_PARAM_MIP_STRONG_MAXIT	Specifies the maximum number of iterations to allow for MIP strong branching solves. [Branching, Knitro-MINLP, Limits]	p. 472

18.3 Callback

Reference section for functions, controls, and attributes related to the use of callback functions. Callbacks enable the user to interact with the solver 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.

18.3.1 Callback library functions

<code>XSLPcopycallbacks</code>	Copy the user-defined callbacks from one SLP problem to another [Callback]	p. 307
<code>XSLPsetcbcascadeend</code> , <code>XPRSsetcbslpcascadeend</code>	Set a user callback to be called at the end of the cascading process, after the last variable has been cascaded [Callback, Cascading, SLP]	p. 379
<code>XSLPsetcbcascadestart</code> , <code>XPRSsetcbslpcascadestart</code>	Set a user callback to be called at the start of the cascading process, before any variables have been cascaded [Callback, Cascading, SLP]	p. 380
<code>XSLPsetcbcascadevar</code> , <code>XPRSsetcbslpcascadevar</code>	Set a user callback to be called after each column has been cascaded [Callback, Cascading, SLP]	p. 381
<code>XSLPsetcbcascadevarfail</code> , <code>XPRSsetcbslpcascadevarfail</code>	Set a user callback to be called after cascading a column was not successful [Callback, Cascading, SLP]	p. 382
<code>XSLPsetcbcoefevalerror</code> , <code>XPRSsetcbnlpcoeefevalerror</code>	Set a user callback to be called when an evaluation of a coefficient fails during the solve [Callback, Numerics]	p. 383
<code>XSLPsetcbconstruct</code> , <code>XPRSsetcbslpconstruct</code>	Set a user callback to be called during the Xpress-SLP augmentation process [Callback, SLP]	p. 384
<code>XSLPsetcbdestroy</code>	Set a user callback to be called when an SLP problem is about to be destroyed [Callback]	p. 386
<code>XSLPsetcbdrcol</code> , <code>XPRSsetcbslpdrcol</code>	Set a user callback used to override the update of variables with small determining column [Callback, Cascading, SLP]	p. 387
<code>XSLPsetcbintsol</code> , <code>XPRSsetcbslpintsol</code>	Set a user callback to be called during MISLP when an integer solution is obtained [Callback, MISLP]	p. 388
<code>XSLPsetcbiterend</code> , <code>XPRSsetcbslpiterend</code>	Set a user callback to be called at the end of each SLP iteration [Callback, SLP]	p. 389
<code>XSLPsetcbiterstart</code> , <code>XPRSsetcbslpiterstart</code>	Set a user callback to be called at the start of each SLP iteration [Callback, SLP]	p. 390
<code>XSLPsetcbitervar</code> , <code>XPRSsetcbslpitervar</code>	Set a user callback to be called after each column has been tested for convergence [Callback, SLP, SLP-convergence]	p. 391
<code>XSLPsetcbmessage</code>	Set a user callback to be called whenever Xpress NonLinear outputs a line of text according to <code>XSLP_ECHOXPRSMESSAGES</code> . [Callback, Logging]	p. 392

<code>XSLPsetcbmsjobend</code> , <code>XPRSsetcbmsjobend</code>	Set a user callback to be called every time a new multistart job finishes. [Callback, Multistart]	p. 394
<code>XSLPsetcbmsjobstart</code> , <code>XPRSsetcbmsjobstart</code>	Set a user callback to be called every time a new multistart job is created, and the pre-loaded settings are applied [Callback, Multistart]	p. 395
<code>XSLPsetcbmswinner</code> , <code>XPRSsetcbmswinner</code>	Set a user callback to be called every time a multistart winner has been declared [Callback, Multistart]	p. 396
<code>XSLPsetcboptnode</code>	Set a user callback to be called during MISLP when an optimal SLP solution is obtained at a node [Callback, MISLP]	p. 397
<code>XSLPsetcbprenode</code>	Set a user callback to be called during MISLP after the set-up of the SLP problem to be solved at a node, but before SLP optimization [Callback, MISLP]	p. 398
<code>XSLPsetcbpresolved</code>	Set a user callback to be called after the nonlinear presolver has been applied. [Callback, SLP]	p. 401
<code>XSLPsetcbpreupdatelinearization</code> , <code>XPRSsetcbslppreupdatelinearization</code>	Set a user callback to be called before the linearization is updated [Callback, SLP]	p. 399
<code>XSLPsetcbslpend</code>	Set a user callback to be called at the end of the SLP optimization [Callback, SLP]	p. 402
<code>XSLPsetcbslpnode</code>	Set a user callback to be called during MISLP after the SLP optimization at each node. [Callback, MISLP]	p. 403
<code>XSLPsetcbslpstart</code>	Set a user callback to be called at the start of the SLP optimization [Callback, SLP]	p. 404

18.4 Cascading

Reference section for functions, controls, and attributes for using cascading in SLP on problems with pooling structure to recompute implied values after the iteration.

18.4.1 Cascading library functions

<code>XSLPcascade</code> , <code>XPRSslpcascade</code>	Re-calculate consistent values for SLP variables based on the current values of the remaining variables. [Cascading, SLP, Solution Process]	p. 288
<code>XSLPsetcbcascadeend</code> , <code>XPRSsetcbslpcascadeend</code>	Set a user callback to be called at the end of the cascading process, after the last variable has been cascaded [Callback, Cascading, SLP]	p. 379

<code>XSLPsetcbascadestart</code> , <code>XPRSsetcbslpcascadestart</code>	Set a user callback to be called at the start of the cascading process, before any variables have been cascaded	[Callback, Cascading, SLP]	p. 380
<code>XSLPsetcbcascadevar</code> , <code>XPRSsetcbslpcascadevar</code>	Set a user callback to be called after each column has been cascaded	[Callback, Cascading, SLP]	p. 381
<code>XSLPsetcbcascadevarfail</code> , <code>XPRSsetcbslpcascadevarfail</code>	Set a user callback to be called after cascading a column was not successful	[Callback, Cascading, SLP]	p. 382
<code>XSLPsetcbdrcol</code> , <code>XPRSsetcbslpdrcol</code>	Set a user callback used to override the update of variables with small determining column	[Callback, Cascading, SLP]	p. 387
<code>XSLPsetdetrow</code> , <code>XPRSslpsetdetrow</code>	Set the determining row of a variable	[Cascading, Data Input, SLP]	p. 409

18.4.2 Cascading controls

<code>XSLP_CASCADE</code> , <code>SLPCASCADE</code>	Bit map describing the cascading to be used	[Cascading, SLP]	p. 223
<code>XSLP_CASCADENLIMIT</code> , <code>SLPCASCADENLIMIT</code>	Maximum number of iterations for cascading with non-linear determining rows	[Cascading, Limits, SLP]	p. 224
<code>XSLP_CASCADETOL_PA</code> , <code>SLPCASCADETOL_PA</code>	Absolute cascading print tolerance	[Cascading, Logging, SLP]	p. 177
<code>XSLP_CASCADETOL_PR</code> , <code>SLPCASCADETOL_PR</code>	Relative cascading print tolerance	[Cascading, Logging, SLP]	p. 177
<code>XSLP_DRCOLDJTOL</code> , <code>SLPDRCOLDJTOL</code>	Reduced cost tolerance on the delta variable when fixing due to the determining column being below <code>XSLP_DRCOLTOL</code> .	[Cascading, SLP, Tolerances]	p. 186
<code>XSLP_DRCOLTOL</code> , <code>SLPDRCOLTOL</code>	The minimum absolute magnitude of a determining column, for which the determined variable is still regarded as well defined	[Cascading, SLP, Tolerances]	p. 186
<code>XSLP_DRFIXRANGE</code> , <code>SLPDRFIXRANGE</code>	The range around the previous value where variables are fixed in cascading if the determining column is below <code>XSLP_DRCOLTOL</code> .	[Cascading, SLP]	p. 187

18.5 Controls and Attributes

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

18.5.1 Controls and Attributes library functions

<code>XSLPcopycontrols</code>	Copy the values of the control variables from one SLP problem to another [Controls and Attributes]	p. 308
<code>XSLPgetdblattrib</code>	Retrieve the value of a double precision problem attribute [Controls and Attributes]	p. 327
<code>XSLPgetdblcontrol</code>	Retrieve the value of a double precision problem control [Controls and Attributes]	p. 328
<code>XSLPgetintattrib</code>	Retrieve the value of an integer problem attribute [Controls and Attributes]	p. 335
<code>XSLPgetintcontrol</code>	Retrieve the value of an integer problem control [Controls and Attributes]	p. 336
<code>XSLPgetptrattrib</code>	Retrieve the value of a problem pointer attribute [Controls and Attributes]	p. 338
<code>XSLPgetstrattrib</code>	Retrieve the value of a string problem attribute [Controls and Attributes]	p. 343
<code>XSLPgetstrcontrol</code>	Retrieve the value of a string problem control [Controls and Attributes]	p. 344
<code>XSLPsetdblcontrol</code>	Set the value of a double precision problem control [Controls and Attributes]	p. 406
<code>XSLPsetdefaultcontrol</code>	Set the values of one SLP control to its default value [Controls and Attributes]	p. 407
<code>XSLPsetdefaults</code>	Set the values of all SLP controls to their default values [Controls and Attributes]	p. 408
<code>XSLPsetintcontrol</code>	Set the value of an integer problem control [Controls and Attributes]	p. 412
<code>XSLPsetparam</code>	Set the value of a control parameter by name [Controls and Attributes]	p. 414
<code>XSLPsetstrcontrol</code>	Set the value of a string problem control [Controls and Attributes]	p. 415

18.6 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.

18.6.1 Cuts controls

<code>XKTR_PARAM_MIP_KNAPSACK</code> , <code>KNITRO_PARAM_MIP_KNAPSACK</code>	Specifies rules for adding MIP knapsack cuts. [Cuts, Knitro-MINLP]	p. 469
<code>XSLP_CUTSTRATEGY</code> , <code>SLPCUTSTRATEGY</code>	Determines which cuts to apply in the MISLP search when the default SLP-in-MIP strategy is used. [Cuts, MISLP]	p. 226

18.7 Data Input

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

18.7.1 Data Input library functions

<code>XSLPcascadeorder</code> , <code>XPRSslpcascadeorder</code>	Establish a re-calculation sequence for SLP variables with determining rows.	
	[Data Input, SLP]	p. 289
<code>XSLPchgcascadenlimit</code> , <code>XPRSslpchgcascadenlimit</code>	Set a variable specific cascade iteration limit	
	[Data Input, SLP]	p. 290
<code>XSLPchgrowstatus</code> , <code>XPRSslpchgrowstatus</code>	Change the status setting of a constraint	
	[Bit-vector, Data Input, SLP]	p. 300
<code>XSLPchgrowwt</code> , <code>XPRSslpchgrowwt</code>	Set or change the initial penalty error weight for a row	
	[Data Input, SLP]	p. 301
<code>XSLPmsaddcustompreset</code> , <code>XPRSmsaddcustompreset</code>	A combined version of <code>XSLPmsaddjob</code> and <code>XSLPmsaddpreset</code> .	
	[Data Input, Multistart]	p. 362
<code>XSLPmsaddjob</code> , <code>XPRSmsaddjob</code>	Adds a multistart job to the multistart pool	
	[Data Input, Multistart]	p. 363
<code>XSLPmsaddpreset</code> , <code>XPRSmsaddpreset</code>	Loads a preset of jobs into the multistart job pool.	
	[Data Input, Multistart]	p. 364
<code>XSLPsetcurrentiv</code> , <code>XPRSnlpsetcurrentiv</code>	Transfer the current solution to initial values	
	[Data Input]	p. 405
<code>XSLPsetdetrow</code> , <code>XPRSslpsetdetrow</code>	Set the determining row of a variable	
	[Cascading, Data Input, SLP]	p. 409
<code>XSLPsetinitval</code> , <code>XPRSnlpsetinitval</code>	Set the initial value of a variable	
	[Data Input]	p. 411

18.7.2 Data Input controls

<code>XSLP_DEFAULTIV</code> , <code>NLPDEFAULTIV</code>	Default initial value for an SLP variable if none is explicitly given	
	[Data Input]	p. 182

18.8 Data Information

Reference section for functions, controls, and attributes related to querying information about auxiliary data.

18.8.1 Data Information library functions

<code>XSLPgetrowwt</code> , <code>XPRSslpgetrowwt</code>	Get the initial penalty error weight for a row	
	[Data Information, SLP]	p. 341

18.8.2 Data Information attributes

XSLP_VARIABLES, NLPVARIABLES	Number of SLP variables [Data Information, SLP]	p. 159
------------------------------	--	--------

18.9 Derivatives

Reference section for functions, controls, and attributes affecting the calculation of derivatives for the local solvers.

18.9.1 Derivatives controls

XKTR_PARAM_GRADOPT, KNITRO_PARAM_GRADOPT	Specifies how to compute the gradients of the objective and constraint functions. [Derivatives, Knitro]	p. 465
XKTR_PARAM_HESSOPT, KNITRO_PARAM_HESSOPT	Specifies how to compute the (approximate) Hessian of the Lagrangian. [Derivatives, Knitro]	p. 465
XSLP_DELTA_A, SLPDELTA_A	Absolute perturbation of values for calculating numerical derivatives [Derivatives]	p. 183
XSLP_DELTA_INFINITY, SLPDELTA_INFINITY	Maximum value for partial derivatives [Derivatives]	p. 183
XSLP_DELTA_R, SLPDELTA_R	Relative perturbation of values for calculating numerical derivatives [Derivatives]	p. 183
XSLP_DELTA_X, SLPDELTA_X	Minimum absolute value of delta coefficients to be retained [Derivatives]	p. 184
XSLP_DELTA_Z, SLPDELTA_Z	Tolerance used when calculating derivatives [Derivatives, Tolerances]	p. 184
XSLP_DELTA_ZERO, SLPDELTA_ZERO	Absolute zero acceptance tolerance used when calculating derivatives [Derivatives, Tolerances]	p. 184
XSLP_DERIVATIVES, NLPDERIVATIVES	Bitmap describing the method of calculating derivatives [Derivatives]	p. 228
XSLP_FUNCEVAL, NLPFUNCEVAL	Bit map for determining the method of evaluating user functions and their derivatives [Derivatives, User Functions]	p. 232
XSLP_HESSIAN, NLPHESSIAN	Second order differentiation mode when using analytical derivatives [Derivatives]	p. 233
XSLP_JACOBIAN, NLPJACOBIAN	First order differentiation mode when using analytical derivatives [Derivatives]	p. 234

18.9.2 Derivatives attributes

XSLP_USEDERIVATIVES, NLPUSEDERIVATIVES	Indicates whether numeric or analytic derivatives were used to create the linear approximations and solve the problem	
	[Derivatives]	p. 159

18.10 File IO

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

18.10.1 File IO library functions

XSLPreadprob	Read an Xpress NonLinear extended MPS format matrix from a file into an SLP problem	
	[File IO, Problem Creation]	p. 371
XSLPrestore	Restore the Xpress NonLinear problem from a file created by XSLPsave	
	[File IO, Save Restore]	p. 374
XSLPsave	Save the Xpress NonLinear problem to file	
	[File IO, Save Restore]	p. 376
XSLPsaveas	Save the Xpress NonLinear problem to a named file	
	[File IO, Save Restore]	p. 377
XSLPsetlogfile	Define an output file to be used to receive messages from Xpress NonLinear	
	[File IO, Logging]	p. 413
XSLPwriteprob	Write the current problem to a file in extended MPS or text format	
	[File IO, Problem Information]	p. 423
XSLPwriteslxsol	Write the current solution to an MPS like file format	
	[File IO, Solution]	p. 424

18.10.2 File IO controls

XSLP_IVNAME, NLP_IVNAME	Name of the set of initial values to be used	
	[File IO]	p. 262
XSLP_TOLNAME, SLPTOLNAME	Name of the set of tolerance sets to be used	
	[File IO, SLP]	p. 265

18.11 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.

18.11.1 Heuristics controls

XKTR_PARAM_MIP_HEURISTIC, KNITRO_PARAM_MIP_HEURISTIC	Specifies which MIP heuristic search approach to apply to try to find an initial integer feasible point.	
	[Heuristics, Knitro-MINLP]	p. 468
XKTR_PARAM_MIP_HEURISTIC_MAXIT, KNITRO_PARAM_MIP_HEURISTIC_MAXIT	Specifies the maximum number of iterations to allow for MIP heuristic, if one is enabled.	
	[Heuristics, Knitro-MINLP]	p. 468
XSLP_FINDIV, NLP_FINDIV	Option for running a heuristic to find a feasible initial point	
	[Heuristics]	p. 231
XSLP_GRIDHEURSELECT, SLP_GRIDHEURSELECT	Bit map selectin which heuristics to run if the problem has variable with an integer delta	
	[Heuristics, SLP]	p. 233
XSLP_HEURSTRATEGY, SLP_HEURSTRATEGY	Branch and Bound: This specifies the MINLP heuristic strategy.	
	[Heuristics, SLP]	p. 233

18.12 Knitro

Reference section for controls related to the Knitro solver.

18.12.1 Knitro controls

XKTR_PARAM_ALGORITHM, KNITRO_PARAM_ALGORITHM	Indicates which algorithm to use to solve the problem	
	[Knitro, Solution Process]	p. 461
XKTR_PARAM_BAR_DIRECTINTERVAL, KNITRO_PARAM_BAR_DIRECTINTERVAL	Controls the maximum number of consecutive conjugate gradient (CG) steps before Knitro will try to enforce that a step is taken using direct linear algebra.	
	[Knitro, Limits]	p. 461
XKTR_PARAM_BAR_FEASIBLE, KNITRO_PARAM_BAR_FEASIBLE	Specifies whether special emphasis is placed on getting and staying feasible in the interior-point algorithms.	
	[Knitro]	p. 461
XKTR_PARAM_BAR_FEASMODETOL, KNITRO_PARAM_BAR_FEASMODETOL	Specifies the tolerance in equation that determines whether Knitro will force subsequent iterates to remain feasible.	
	[Knitro, Tolerances]	p. 457
XKTR_PARAM_BAR_INITMU, KNITRO_PARAM_BAR_INITMU	Specifies the initial value for the barrier parameter : μ used with the barrier algorithms.	
	[Knitro]	p. 457
XKTR_PARAM_BAR_INITPT, KNITRO_PARAM_BAR_INITPT	Indicates whether an initial point strategy is used with barrier algorithms.	
	[Knitro]	p. 462

XKTR_PARAM_BAR_MAXBACKTRACK, KNITRO_PARAM_BAR_MAXBACKTRACK	Indicates the maximum allowable number of backtracks during the linesearch of the Interior/Direct algorithm before reverting to a CG step. [Knitro, Limits]	p. 462
XKTR_PARAM_BAR_MAXCROSSIT, KNITRO_PARAM_BAR_MAXCROSSIT	Specifies the maximum number of crossover iterations before termination. [Knitro, Limits]	p. 463
XKTR_PARAM_BAR_MAXREFACTOR, KNITRO_PARAM_BAR_MAXREFACTOR	Indicates the maximum number of refactorizations of the KKT system per iteration of the Interior/Direct algorithm before reverting to a CG step. [Knitro, Limits]	p. 463
XKTR_PARAM_BAR_MURULE, KNITRO_PARAM_BAR_MURULE	Indicates which strategy to use for modifying the barrier parameter μ in the barrier algorithms. [Knitro]	p. 463
XKTR_PARAM_BAR_PENCONS, KNITRO_PARAM_BAR_PENCONS	Indicates whether a penalty approach is applied to the constraints. [Knitro]	p. 464
XKTR_PARAM_BAR_PENRULE, KNITRO_PARAM_BAR_PENRULE	Indicates which penalty parameter strategy to use for determining whether or not to accept a trial iterate. [Knitro]	p. 464
XKTR_PARAM_BAR_SWITCHRULE, KNITRO_PARAM_BAR_SWITCHRULE	Indicates whether or not the barrier algorithms will allow switching from an optimality phase to a pure feasibility phase. [Knitro]	p. 465
XKTR_PARAM_DELTA, KNITRO_PARAM_DELTA	Specifies the initial trust region radius scaling factor used to determine the initial trust region size. [Knitro]	p. 457
XKTR_PARAM_FEASTOL, KNITRO_PARAM_FEASTOL	Specifies the final relative stopping tolerance for the feasibility error. [Knitro, Tolerances]	p. 457
XKTR_PARAM_FEASTOLABS, KNITRO_PARAM_FEASTOLABS	Specifies the final absolute stopping tolerance for the feasibility error. [Knitro, Tolerances]	p. 458
XKTR_PARAM_GRADOPT, KNITRO_PARAM_GRADOPT	Specifies how to compute the gradients of the objective and constraint functions. [Derivatives, Knitro]	p. 465
XKTR_PARAM_HESSOPT, KNITRO_PARAM_HESSOPT	Specifies how to compute the (approximate) Hessian of the Lagrangian. [Derivatives, Knitro]	p. 465
XKTR_PARAM_HONORBND, KNITRO_PARAM_HONORBND	Indicates whether or not to enforce satisfaction of simple variable bounds throughout the optimization. [Knitro]	p. 466
XKTR_PARAM_INFEASTOL, KNITRO_PARAM_INFEASTOL	Specifies the (relative) tolerance used for declaring infeasibility of a model. [Knitro, Tolerances]	p. 458

<code>XKTR_PARAM_LMSIZE, KNITRO_PARAM_LMSIZE</code>	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option. [Knitro, Limits]	p. 466
<code>XKTR_PARAM_MAXCGIT, KNITRO_PARAM_MAXCGIT</code>	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option. [Knitro, Limits]	p. 467
<code>XKTR_PARAM_MAXIT, KNITRO_PARAM_MAXIT</code>	Specifies the maximum number of iterations before termination. [Knitro, Limits]	p. 467
<code>XKTR_PARAM_MIP_INTEGERTOL, KNITRO_PARAM_INTEGERTOL</code>	This value specifies the threshold for deciding whether or not a variable is determined to be an integer. [Knitro, Knitro-MINLP, Tolerances]	p. 458
<code>XKTR_PARAM_MIP_INTGAPABS, KNITRO_PARAM_INTGAPABS</code>	The absolute integrality gap stop tolerance for MIP. [Knitro, Knitro-MINLP, Tolerances]	p. 458
<code>XKTR_PARAM_MIP_INTGAPREL, KNITRO_PARAM_INTGAPREL</code>	The relative integrality gap stop tolerance for MIP. [Knitro, Knitro-MINLP, Tolerances]	p. 459
<code>XKTR_PARAM_OBJRANGE, KNITRO_PARAM_OBJRANGE</code>	Specifies the extreme limits of the objective function for purposes of determining unboundedness. [Knitro, Limits]	p. 459
<code>XKTR_PARAM_OPTTOL, KNITRO_PARAM_OPTTOL</code>	Specifies the final relative stopping tolerance for the KKT (optimality) error. [Knitro, Tolerances]	p. 459
<code>XKTR_PARAM_OPTTOLABS, KNITRO_PARAM_OPTTOLABS</code>	Specifies the final absolute stopping tolerance for the KKT (optimality) error. [Knitro, Tolerances]	p. 459
<code>XKTR_PARAM_OUTLEV, KNITRO_PARAM_OUTLEV</code>	Controls the level of output produced by Knitro. [Knitro, Logging]	p. 473
<code>XKTR_PARAM_PRESOLVE, KNITRO_PARAM_PRESOLVE</code>	Determine whether or not to use the Knitro presolver to try to simplify the model by removing variables or constraints. [Knitro, Presolve]	p. 473
<code>XKTR_PARAM_PRESOLVE_TOL, KNITRO_PARAM_PRESOLVE_TOL</code>	Determines the tolerance used by the Knitro presolver to remove variables and constraints from the model. [Knitro, Presolve, Tolerances]	p. 460
<code>XKTR_PARAM_SCALE, KNITRO_PARAM_SCALE</code>	Performs a scaling of the objective and constraint functions based on their values at the initial point. [Knitro, Numerics]	p. 473
<code>XKTR_PARAM_SOC, KNITRO_PARAM_SOC</code>	Specifies whether or not to try second order corrections (SOC). [Knitro]	p. 474

`XKTR_PARAM_XTOL`, `KNITRO_PARAM_XTOL` The optimization process will terminate if the relative change in all components of the solution point estimate is less than `xtol`.
[Knitro, Tolerances] p. 460

18.13 Knitro-MINLP

Reference section for controls related to the MINLP solver in Knitro.

18.13.1 Knitro-MINLP controls

`XKTR_PARAM_MIP_BRANCHRULE`, `KNITRO_PARAM_MIP_BRANCHRULE` Specifies which branching rule to use for MIP branch and bound procedure.
[Branching, Knitro-MINLP] p. 467

`XKTR_PARAM_MIP_GUB_BRANCH`, `KNITRO_PARAM_MIP_GUB_BRANCH` Specifies whether or not to branch on generalized upper bounds (GUBs).
[Branching, Knitro-MINLP] p. 468

`XKTR_PARAM_MIP_HEURISTIC`, `KNITRO_PARAM_MIP_HEURISTIC` Specifies which MIP heuristic search approach to apply to try to find an initial integer feasible point.
[Heuristics, Knitro-MINLP] p. 468

`XKTR_PARAM_MIP_HEURISTIC_MAXIT`, `KNITRO_PARAM_MIP_HEURISTIC_MAXIT` Specifies the maximum number of iterations to allow for MIP heuristic, if one is enabled.
[Heuristics, Knitro-MINLP] p. 468

`XKTR_PARAM_MIP_IMPLICATNS`, `KNITRO_PARAM_MIP_IMPLICATNS` Specifies whether or not to add constraints to the MIP derived from logical implications.
[Knitro-MINLP, Presolve] p. 468

`XKTR_PARAM_MIP_INTEGERTOL`, `KNITRO_PARAM_INTEGERTOL` This value specifies the threshold for deciding whether or not a variable is determined to be an integer.
[Knitro, Knitro-MINLP, Tolerances] p. 458

`XKTR_PARAM_MIP_INTGAPABS`, `KNITRO_PARAM_INTGAPABS` The absolute integrality gap stop tolerance for MIP.
[Knitro, Knitro-MINLP, Tolerances] p. 458

`XKTR_PARAM_MIP_INTGAPREL`, `KNITRO_PARAM_INTGAPREL` The relative integrality gap stop tolerance for MIP.
[Knitro, Knitro-MINLP, Tolerances] p. 459

`XKTR_PARAM_MIP_KNAPSACK`, `KNITRO_PARAM_MIP_KNAPSACK` Specifies rules for adding MIP knapsack cuts.
[Cuts, Knitro-MINLP] p. 469

`XKTR_PARAM_MIP_LPALG`, `KNITRO_PARAM_MIP_LPALG` Specifies which algorithm to use for any linear programming (LP) subproblem solves that may occur in the MIP branch and bound procedure.
[Knitro-MINLP, Solution Process] p. 469

`XKTR_PARAM_MIP_MAXNODES`, `KNITRO_PARAM_MIP_MAXNODES` Specifies the maximum number of nodes explored.
[Knitro-MINLP, Limits] p. 469

<code>XKTR_PARAM_MIP_MAXSOLVES</code> , <code>KNITRO_PARAM_MIP_MAXSOLVES</code>	Specifies the maximum number of subproblem solves allowed (0 means no limit). [Knitro-MINLP, Limits]	p. 470
<code>XKTR_PARAM_MIP_METHOD</code> , <code>KNITRO_PARAM_MIP_METHOD</code>	Specifies which MIP method to use. [Knitro-MINLP, Solution Process]	p. 470
<code>XKTR_PARAM_MIP_OUTINTERVAL</code> , <code>KNITRO_PARAM_MIP_OUTINTERVAL</code>	Specifies node printing interval for <code>XKTR_PARAM_MIP_OUTLEVEL</code> when <code>XKTR_PARAM_MIP_OUTLEVEL > 0</code> . [Knitro-MINLP, Logging]	p. 470
<code>XKTR_PARAM_MIP_OUTLEVEL</code> , <code>KNITRO_PARAM_MIP_OUTLEVEL</code>	Specifies how much MIP information to print. [Knitro-MINLP, Logging]	p. 470
<code>XKTR_PARAM_MIP_PSEUDOINIT</code> , <code>KNITRO_PARAM_MIP_PSEUDOINIT</code>	Specifies the method used to initialize pseudo-costs corresponding to variables that have not yet been branched on in the MIP method. [Branching, Knitro-MINLP]	p. 471
<code>XKTR_PARAM_MIP_ROOTALG</code> , <code>KNITRO_PARAM_MIP_ROOTALG</code>	Specifies which algorithm to use for the root node solve in MIP (same options as <code>XKTR_PARAM_ALGORITHM</code> user option). [Knitro-MINLP, Solution Process]	p. 471
<code>XKTR_PARAM_MIP_ROUNDING</code> , <code>KNITRO_PARAM_MIP_ROUNDING</code>	Specifies the MIP rounding rule to apply. [Knitro-MINLP]	p. 471
<code>XKTR_PARAM_MIP_SELECTRULE</code> , <code>KNITRO_PARAM_MIP_SELECTRULE</code>	Specifies the MIP select rule for choosing the next node in the branch and bound tree. [Knitro-MINLP]	p. 471
<code>XKTR_PARAM_MIP_STRONG_CANDLIM</code> , <code>KNITRO_PARAM_MIP_STRONG_CANDLIM</code>	Specifies the maximum number of candidates to explore for MIP strong branching. [Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_STRONG_LEVEL</code> , <code>KNITRO_PARAM_MIP_STRONG_LEVEL</code>	Specifies the maximum number of tree levels on which to perform MIP strong branching. [Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_STRONG_MAXIT</code> , <code>KNITRO_PARAM_MIP_STRONG_MAXIT</code>	Specifies the maximum number of iterations to allow for MIP strong branching solves. [Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_TERMINATE</code> , <code>KNITRO_PARAM_MIP_TERMINATE</code>	Specifies conditions for terminating the MIP algorithm. [Knitro-MINLP]	p. 472

18.14 Licensing

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

18.14.1 Licensing library functions

XSLPfree	Free any memory allocated by Xpress NonLinear and close any open Xpress NonLinear files [Licensing]	p. 320
XSLPinit	Initializes the Xpress NonLinear system [Licensing]	p. 349

18.15 Limits

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

18.15.1 Limits controls

XKTR_PARAM_BAR_DIRECTINTERVAL, KNITRO_PARAM_BAR_DIRECTINTERVAL	Controls the maximum number of consecutive conjugate gradient (CG) steps before Knitro will try to enforce that a step is taken using direct linear algebra. [Knitro, Limits]	p. 461
XKTR_PARAM_BAR_MAXBACKTRACK, KNITRO_PARAM_BAR_MAXBACKTRACK	Indicates the maximum allowable number of backtracks during the linesearch of the Interior/Direct algorithm before reverting to a CG step. [Knitro, Limits]	p. 462
XKTR_PARAM_BAR_MAXCROSSIT, KNITRO_PARAM_BAR_MAXCROSSIT	Specifies the maximum number of crossover iterations before termination. [Knitro, Limits]	p. 463
XKTR_PARAM_BAR_MAXREFACTOR, KNITRO_PARAM_BAR_MAXREFACTOR	Indicates the maximum number of refactorizations of the KKT system per iteration of the Interior/Direct algorithm before reverting to a CG step. [Knitro, Limits]	p. 463
XKTR_PARAM_LMSIZE, KNITRO_PARAM_LMSIZE	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option. [Knitro, Limits]	p. 466
XKTR_PARAM_MAXCGIT, KNITRO_PARAM_MAXCGIT	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option. [Knitro, Limits]	p. 467
XKTR_PARAM_MAXIT, KNITRO_PARAM_MAXIT	Specifies the maximum number of iterations before termination. [Knitro, Limits]	p. 467
XKTR_PARAM_MIP_MAXNODES, KNITRO_PARAM_MIP_MAXNODES	Specifies the maximum number of nodes explored. [Knitro-MINLP, Limits]	p. 469
XKTR_PARAM_MIP_MAXSOLVES, KNITRO_PARAM_MIP_MAXSOLVES	Specifies the maximum number of subproblem solves allowed (0 means no limit). [Knitro-MINLP, Limits]	p. 470

<code>XKTR_PARAM_MIP_STRONG_CANDLIM</code> , <code>KNITRO_PARAM_MIP_STRONG_CANDLIM</code>	Specifies the maximum number of candidates to explore for MIP strong branching. [Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_STRONG_LEVEL</code> , <code>KNITRO_PARAM_MIP_STRONG_LEVEL</code>	Specifies the maximum number of tree levels on which to perform MIP strong branching. [Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_STRONG_MAXIT</code> , <code>KNITRO_PARAM_MIP_STRONG_MAXIT</code>	Specifies the maximum number of iterations to allow for MIP strong branching solves. [Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_OBJRANGE</code> , <code>KNITRO_PARAM_OBJRANGE</code>	Specifies the extreme limits of the objective function for purposes of determining unboundedness. [Knitro, Limits]	p. 459
<code>XSLP_BARLIMIT</code> , <code>SLPBARLIMIT</code>	Number of initial SLP iterations using the barrier method [Limits, Linearizations, SLP]	p. 221
<code>XSLP_BARSTALLINGLIMIT</code> , <code>SLPBARSTALLINGLIMIT</code>	Number of iterations to allow numerical failures in barrier before switching to dual [Limits, Linearizations, SLP]	p. 222
<code>XSLP_BARSTALLINGOBJLIMIT</code> , <code>SLPBARSTALLINGOBJLIMIT</code>	Number of iterations over which to measure the objective change for barrier iterations with no crossover [Limits, Linearizations, SLP]	p. 222
<code>XSLP_CASCADENLIMIT</code> , <code>SLPCASCADENLIMIT</code>	Maximum number of iterations for cascading with non-linear determining rows [Cascading, Limits, SLP]	p. 224
<code>XSLP_INFEASLIMIT</code> , <code>SLPINFEASLIMIT</code>	The maximum number of consecutive infeasible SLP iterations which can occur before Xpress-SLP terminates [Limits, SLP]	p. 234
<code>XSLP_ITERLIMIT</code> , <code>SLPITERLIMIT</code>	The maximum number of SLP iterations [Limits, SLP]	p. 234
<code>XSLP_LSITERLIMIT</code> , <code>SLPLSITERLIMIT</code>	Number of iterations in the line search [Limits, SLP]	p. 236
<code>XSLP_LSPATTERNLIMIT</code> , <code>SLPLSPATTERNLIMIT</code>	Number of iterations in the pattern search preceding the line search [Limits, SLP]	p. 236
<code>XSLP_LSZEROLIMIT</code> , <code>SLPLSZEROLIMIT</code>	Maximum number of zero length line search steps before line search is deactivated [Limits, SLP]	p. 237
<code>XSLP_MAXTIME</code> , <code>NLPMAXTIME</code>	The maximum time in seconds that the SLP optimization will run before it terminates [Limits]	p. 237
<code>XSLP_MIPCUTOFFCOUNT</code> , <code>SLPMIPCUTOFFCOUNT</code>	Number of SLP iterations to check when considering a node for cutting off [Limits, MISLP]	p. 239

<code>XSLP_MIPCUTOFFLIMIT</code> , <code>SLPMIPCUTOFFLIMIT</code>	Number of SLP iterations to check when considering a node for cutting off [Limits, MISLP]	p. 239
<code>XSLP_MIPITERLIMIT</code> , <code>SLPMIPITERLIMIT</code>	Maximum number of SLP iterations at each node [Limits, MISLP]	p. 241
<code>XSLP_MULTISTART_MAXSOLVES</code> , <code>MULTISTART_MAXSOLVES</code>	The maximum number of jobs to create during the multistart search. [Limits, Multistart]	p. 243
<code>XSLP_MULTISTART_MAXTIME</code> , <code>MULTISTART_MAXTIME</code>	The maximum total time to be spent in the multistart search. [Limits, Multistart]	p. 243
<code>XSLP_MULTISTART_POOLSIZE</code> , <code>MULTISTART_POOLSIZE</code>	The maximum number of problem objects allowed to pool up before synchronization in the deterministic multistart. [Limits, Multistart]	p. 244
<code>XSLP_XLIMIT</code> , <code>SLPXLIMIT</code>	Number of SLP iterations up to which static objective (1) convergence testing is performed [Limits, SLP, SLP-convergence]	p. 258
<code>XSLP_ZEROCRITERIONCOUNT</code> , <code>SLPZEROCRITERIONCOUNT</code>	Number of consecutive times a placeholder entry is zero before being considered for deletion [Limits, SLP]	p. 259

18.16 Linearizations

Reference section for functions, controls, and attributes related to solving the linearizations in SLP.

18.16.1 Linearizations controls

<code>XSLP_BARCROSSOVERSTART</code> , <code>SLPBARCROSSOVERSTART</code>	Default crossover activation behaviour for barrier start [Linearizations, SLP]	p. 221
<code>XSLP_BARLIMIT</code> , <code>SLPBARLIMIT</code>	Number of initial SLP iterations using the barrier method [Limits, Linearizations, SLP]	p. 221
<code>XSLP_BARSTALLINGLIMIT</code> , <code>SLPBARSTALLINGLIMIT</code>	Number of iterations to allow numerical failures in barrier before switching to dual [Limits, Linearizations, SLP]	p. 222
<code>XSLP_BARSTALLINGOBJLIMIT</code> , <code>SLPBARSTALLINGOBJLIMIT</code>	Number of iterations over which to measure the objective change for barrier iterations with no crossover [Limits, Linearizations, SLP]	p. 222
<code>XSLP_BARSTALLINGTOL</code> , <code>SLPBARSTALLINGTOL</code>	Required change in the objective when progress is measured in barrier iterations without crossover [Linearizations, SLP]	p. 176
<code>XSLP_BARSTARTOPS</code> , <code>SLPBARSTARTOPS</code>	Controls behaviour when the barrier is used to solve the linearizations [Linearizations, SLP]	p. 222

<code>XSLP_FEASTOLTARGET</code> , <code>SLPFEASTOLTARGET</code>	When set, this defines a target feasibility tolerance to which the linearizations are solved to [Linearizations, SLP, Tolerances]	p. 193
<code>XSLP_ITERFALLBACKOPS</code> , <code>SLPITERFALLBACKOPS</code>	Alternative LP level control values for numerically challenging problems [Linearizations, Numerics, SLP]	p. 261
<code>XSLP_MIPDEFAULTALGORITHM</code> , <code>SLPMIPDEFAULTALGORITHM</code>	Default algorithm to be used during the tree search in MISLP [Linearizations, MISLP]	p. 240
<code>XSLP_OPTIMALITYTOLTARGET</code> , <code>SLPOPTIMALITYTOLTARGET</code>	When set, this defines a target optimality tolerance to which the linearizations are solved to [Linearizations, SLP, Tolerances]	p. 203
<code>XSLP_UNFINISHEDLIMIT</code> , <code>SLPUNFINISHEDLIMIT</code>	The number of consecutive SLP iterations that may have an unfinished status before the solve is terminated. [Linearizations, SLP]	p. 254

18.16.2 Linearizations attributes

<code>XSLP_ECFCOUNT</code> , <code>SLPECFCOUNT</code>	Number of infeasible constraints found at the point of linearization [Linearizations, SLP]	p. 146
---	---	--------

18.17 Logging

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

18.17.1 Logging library functions

<code>XSLPprntevalinfo</code> , <code>XPRSnlpprntevalinfo</code>	Print a summary of any evaluation errors that may have occurred during solving a problem [Logging]	p. 370
<code>XSLPprintmemory</code>	Print the dimensions and memory allocations for a problem [Logging]	p. 369
<code>XSLPsetcbmessage</code>	Set a user callback to be called whenever Xpress NonLinear outputs a line of text according to <code>XSLP_ECHOXPRSMESSAGES</code> . [Callback, Logging]	p. 392
<code>XSLPsetlogfile</code>	Define an output file to be used to receive messages from Xpress NonLinear [File IO, Logging]	p. 413

18.17.2 Logging controls

<code>XKTR_PARAM_MIP_OUTINTERVAL</code> , <code>KNITRO_PARAM_MIP_OUTINTERVAL</code>	Specifies node printing interval for <code>XKTR_PARAM_MIP_OUTLEVEL</code> when <code>XKTR_PARAM_MIP_OUTLEVEL > 0</code> . [Knitro-MINLP, Logging]	p. 470
---	---	--------

XKTR_PARAM_MIP_OUTLEVEL, KNITRO_PARAM_MIP_OUTLEVEL	Specifies how much MIP information to print. [Knitro-MINLP, Logging]	p. 470
XKTR_PARAM_OUTLEV, KNITRO_PARAM_OUTLEV	Controls the level of output produced by Knitro. [Knitro, Logging]	p. 473
XSLP_ANALYZE, SLPANALYZE	Bit map activating additional options supporting model / solution path analysis [Bit-vector, Logging, SLP]	p. 218
XSLP_AUTOSAVE, SLPAUTOSAVE	Frequency with which to save the model [Logging, SLP]	p. 220
XSLP_CASCADETOL_PA, SLPCASCADETOL_PA	Absolute cascading print tolerance [Cascading, Logging, SLP]	p. 177
XSLP_CASCADETOL_PR, SLPCASCADETOL_PR	Relative cascading print tolerance [Cascading, Logging, SLP]	p. 177
XSLP_DELTAFORMAT, SLPDELTAFORMAT	Formatting string for creation of names for SLP delta vectors [Logging, SLP]	p. 261
XSLP_DELTAOFFSET, SLPDELTAOFFSET	Position of first character of SLP variable name used to create name of delta vector [Logging, SLP]	p. 227
XSLP_ECHOXPRSMESSAGES	Controls if the XSLP message callback should relay messages from the XPRS library. [Logging]	p. 230
XSLP_ERROROFFSET, SLPERROROFFSET	Position of first character of constraint name used to create name of penalty error vectors [Logging, SLP]	p. 230
XSLP_LOG, NLPLOG	Level of printing during SLP iterations [Logging, SLP]	p. 236
XSLP_MINUSDELTAFORMAT, SLPMINUSDELTAFORMAT	Formatting string for creation of names for SLP negative penalty delta vectors [Logging, SLP]	p. 262
XSLP_MINUSERERRORFORMAT, SLPMINUSERERRORFORMAT	Formatting string for creation of names for SLP negative penalty error vectors [Logging, SLP]	p. 262
XSLP_MIPLOG, SLPMIPLOG	Frequency with which MIP status is printed [Logging, MISLP]	p. 241
XSLP_MULTISTART_LOG, MULTISTART_LOG	The level of logging during the multistart run. [Logging, Multistart]	p. 243
XSLP_PENALTYCOLFORMAT, SLPPENALTYCOLFORMAT	Formatting string for creation of the names of the SLP penalty transfer vectors [Logging, SLP]	p. 263
XSLP_PENALTYINFOSTART, SLPPENALTYINFOSTART	Iteration from which to record row penalty information [Logging, SLP]	p. 246

<code>XSLP_PENALTYROWFORMAT</code> , <code>SLPPENALTYROWFORMAT</code>	Formatting string for creation of the names of the SLP penalty rows [Logging, SLP]	p. 263
<code>XSLP_PLUSDELTAFORMAT</code> , <code>SLPLUSDELTAFORMAT</code>	Formatting string for creation of names for SLP positive penalty delta vectors [Logging, SLP]	p. 263
<code>XSLP_PLUSEXACTFORMAT</code> , <code>SLPLUSEXACTFORMAT</code>	Formatting string for creation of names for SLP positive penalty error vectors [Logging, SLP]	p. 264
<code>XSLP_PRIMALINTEGRALALPHA</code> , <code>NLPPRIMALINTEGRALALPHA</code>	Decay term for primal integral computation [Logging]	p. 205
<code>XSLP_PRIMALINTEGRALREF</code> , <code>NLPPRIMALINTEGRALREF</code>	Reference solution value to take into account when calculating the primal integral [Logging]	p. 206
<code>XSLP_SBLOROWFORMAT</code> , <code>SLPSBLOROWFORMAT</code>	Formatting string for creation of names for SLP lower step bound rows [Logging, SLP]	p. 264
<code>XSLP_SBNAME</code> , <code>SLPSBNAME</code>	Name of the set of initial step bounds to be used [Logging, SLP]	p. 265
<code>XSLP_SBROWOFFSET</code> , <code>SLPSBROWOFFSET</code>	Position of first character of SLP variable name used to create name of SLP lower and upper step bound rows [Logging, SLP]	p. 250
<code>XSLP_SBUPROWFORMAT</code> , <code>SLPSBUPROWFORMAT</code>	Formatting string for creation of names for SLP upper step bound rows [Logging, SLP]	p. 265
<code>XSLP_SLPLOG</code>	Frequency with which SLP status is printed [Logging, SLP]	p. 252
<code>XSLP_TRACEMASK</code> , <code>SLPTRACEMASK</code>	Mask of variable or row names that are to be traced through the SLP iterates [Logging, SLP]	p. 266
<code>XSLP_TRACEMASKOPS</code> , <code>SLPTRACEMASKOPS</code>	Controls the information printed for <code>XSLP_TRACEMASK</code> . [Bit-vector, Logging, SLP]	p. 253
<code>XSLP_UPDATEFORMAT</code> , <code>SLPUPDATEFORMAT</code>	Formatting string for creation of names for SLP update rows [Logging, SLP]	p. 266
<code>XSLP_UPDATEOFFSET</code> , <code>SLPUPDATEOFFSET</code>	Position of first character of SLP variable name used to create name of SLP update row [Logging, SLP]	p. 255

18.18 Memory

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

18.18.1 Memory controls

XSLP_MEMORYFACTOR	Factor for expanding size of dynamic arrays in memory [Memory]	p. 196
-------------------	---	--------

18.19 MISLP

Reference section for functions, controls, and attributes for using a combination of sequential linear programming, and branch and bound to solve MINLPs to local optimality.

18.19.1 MISLP library functions

XSLPsetcbintsol, XPRSsetcbslpintsol	Set a user callback to be called during MISLP when an integer solution is obtained [Callback, MISLP]	p. 388
XSLPsetcboptnode	Set a user callback to be called during MISLP when an optimal SLP solution is obtained at a node [Callback, MISLP]	p. 397
XSLPsetcbprenode	Set a user callback to be called during MISLP after the set-up of the SLP problem to be solved at a node, but before SLP optimization [Callback, MISLP]	p. 398
XSLPsetcbslpnode	Set a user callback to be called during MISLP after the SLP optimization at each node. [Callback, MISLP]	p. 403

18.19.2 MISLP controls

XSLP_CUTSTRATEGY, SLPCUTSTRATEGY	Determines which cuts to apply in the MISLP search when the default SLP-in-MIP strategy is used. [Cuts, MISLP]	p. 226
XSLP_DETERMINISTIC, NLPDETERMINISTIC	Determines if the parallel features of SLP should be guaranteed to be deterministic [MISLP, Parallel, SLP]	p. 229
XSLP_MIPALGORITHM, SLPMIPALGORITHM	Bitmap describing the MISLP algorithms to be used [Bit-vector, MISLP]	p. 238
XSLP_MIPCUTOFF_A, SLPMIPCUTOFF_A	Absolute objective function cutoff for MIP termination [MISLP, Tolerances]	p. 197
XSLP_MIPCUTOFF_R, SLPMIPCUTOFF_R	Absolute objective function cutoff for MIP termination [MISLP, Tolerances]	p. 198
XSLP_MIPCUTOFFCOUNT, SLPMIPCUTOFFCOUNT	Number of SLP iterations to check when considering a node for cutting off [Limits, MISLP]	p. 239
XSLP_MIPCUTOFFLIMIT, SLPMIPCUTOFFLIMIT	Number of SLP iterations to check when considering a node for cutting off [Limits, MISLP]	p. 239

XSLP_MIPDEFAULTALGORITHM, SLP_MIPDEFAULTALGORITHM	Default algorithm to be used during the tree search in MISLP [Linearizations, MISLP]	p. 240
XSLP_MIPERRORTOL_A, SLP_MIPERRORTOL_A	Absolute penalty error cost tolerance for MIP cut-off [MISLP, Tolerances]	p. 198
XSLP_MIPERRORTOL_R, SLP_MIPERRORTOL_R	Relative penalty error cost tolerance for MIP cut-off [MISLP, Tolerances]	p. 199
XSLP_MIPFIXSTEPBOUNDS, SLP_MIPFIXSTEPBOUNDS	Bitmap describing the step-bound fixing strategy during MISLP [MISLP]	p. 240
XSLP_MIPITERLIMIT, SLP_MIPITERLIMIT	Maximum number of SLP iterations at each node [Limits, MISLP]	p. 241
XSLP_MIPLOG, SLP_MIPLOG	Frequency with which MIP status is printed [Logging, MISLP]	p. 241
XSLP_MIPOCOUNT, SLP_MIPOCOUNT	Number of SLP iterations at each node over which to measure objective function variation [MISLP, SLP-convergence]	p. 241
XSLP_MIPOTOL_A, SLP_MIPOTOL_A	Absolute objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MIPOTOL_R, SLP_MIPOTOL_R	Relative objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MIPRELAXSTEPBOUNDS, SLP_MIPRELAXSTEPBOUNDS	Bitmap describing the step-bound relaxation strategy during MISLP [MISLP]	p. 242

18.19.3 MISLP attributes

XSLP_MIPITER, SLP_MIPITER	Total number of SLP iterations in MISLP [MISLP]	p. 149
XSLP_MIPNODES, SLP_MIPNODES	Number of nodes explored in MISLP. [MISLP]	p. 149
XSLP_MIPPROBLEM	The underlying Optimizer MIP problem. [MISLP]	p. 161
XSLP_MIPSOLS, SLP_MIPSOLS	Number of integer solutions found in MISLP. [MISLP, Solution]	p. 149

18.20 Misc

Reference section for further miscellaneous functionality.

18.20.1 Misc library functions

XSLPgetlasterror	Retrieve the error message corresponding to the last Xpress NonLinear error during an SLP run [Misc]	p. 337
XSLPsetfunctionerror, XPRSnlpsetfunctionerror	Set the function error flag for the problem [Misc]	p. 410

18.20.2 Misc controls

XSLP_KEEPEQUALSCOLUMN, NLPKEEPEQUALSCOLUMN	When set to a nonzero value, the MPS reader will keep the equals column in the problem [Misc]	p. 235
--	--	--------

18.20.3 Misc attributes

XSLP_IFS, NLPIFS	Number of internal functions [Misc]	p. 147
XSLP_JOBID, NLPJOBID	Unique identifier for the current job [Misc, Multistart]	p. 148
XSLP_PRIMALINTEGRAL, NLPPRIMALINTEGRAL	Local primal integral of the solve [Misc]	p. 144
XSLP_VERSIONDATE	Date of creation of Xpress NonLinear [Misc]	p. 162
XSLP_XPRSPROBLEM	The underlying Optimizer problem [Misc]	p. 161
XSLP_XSLPPROBLEM	The Xpress NonLinear problem [Misc]	p. 161

18.21 Multistart

Reference section for functions, controls, and attributes for using multistart with the nonlinear local solvers.

18.21.1 Multistart library functions

XSLPmsaddcustompreset, XPRSmaddcustompreset	A combined version of XSLPmsaddjob and XSLPmsaddpreset. [Data Input, Multistart]	p. 362
XSLPmsaddjob, XPRSmaddjob	Adds a multistart job to the multistart pool [Data Input, Multistart]	p. 363
XSLPmsaddpreset, XPRSmaddpreset	Loads a preset of jobs into the multistart job pool. [Data Input, Multistart]	p. 364
XSLPmsclear, XPRSmclear	Removes all scheduled jobs from the multistart job pool [Multistart]	p. 365

XSLPsetcbmsjobend, XPRSsetcbmsjobend	Set a user callback to be called every time a new multistart job finishes.	
	[Callback, Multistart]	p. 394
XSLPsetcbmsjobstart, XPRSsetcbmsjobstart	Set a user callback to be called every time a new multistart job is created, and the pre-loaded settings are applied	
	[Callback, Multistart]	p. 395
XSLPsetcbmswinner, XPRSsetcbmswinner	Set a user callback to be called every time a multistart winner has been declared	
	[Callback, Multistart]	p. 396

18.21.2 Multistart controls

XSLP_MS_MAXBOUNDRANGE, MS_MAXBOUNDRANGE	Defines the maximum range inside which initial points are generated by multistart presets	
	[Multistart]	p. 200
XSLP_MULTISTART, MULTISTART	The multistart master control.	
	[Multistart]	p. 242
XSLP_MULTISTART_LOG, MULTISTART_LOG	The level of logging during the multistart run.	
	[Logging, Multistart]	p. 243
XSLP_MULTISTART_MAXSOLVES, MULTISTART_MAXSOLVES	The maximum number of jobs to create during the multistart search.	
	[Limits, Multistart]	p. 243
XSLP_MULTISTART_MAXTIME, MULTISTART_MAXTIME	The maximum total time to be spent in the multistart search.	
	[Limits, Multistart]	p. 243
XSLP_MULTISTART_POOLSIZE, MULTISTART_POOLSIZE	The maximum number of problem objects allowed to pool up before synchronization in the deterministic multistart.	
	[Limits, Multistart]	p. 244
XSLP_MULTISTART_SEED, MULTISTART_SEED	Random seed used for the automatic generation of initial point when loading multistart presets	
	[Multistart]	p. 244
XSLP_MULTISTART_THREADS, MULTISTART_THREADS	The maximum number of threads to be used in multistart	
	[Multistart, Parallel]	p. 245

18.21.3 Multistart attributes

XSLP_JOBID, NLPJOBID	Unique identifier for the current job	
	[Misc, Multistart]	p. 148
XSLP_MSSTATUS	Status of the multistart search	
	[Multistart]	p. 150

18.22 Names Manager

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

18.22.1 Names Manager library functions

XSLPitemname	Retrieves the name of an Xpress NonLinear entity or the value of a function token as a character string. [Names Manager]	p. 351
--------------	---	--------

18.23 Numerics

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

18.23.1 Numerics library functions

XSLPscaling	Analyze the current matrix for largest/smallest coefficients and ratios [Numerics]	p. 378
XSLPsetcbcofevalerror, XPRSsetcbnlpcfevalerror	Set a user callback to be called when an evaluation of a coefficient fails during the solve [Callback, Numerics]	p. 383

18.23.2 Numerics controls

XKTR_PARAM_SCALE, KNITRO_PARAM_SCALE	Performs a scaling of the objective and constraint functions based on their values at the initial point. [Knitro, Numerics]	p. 473
XSLP_INFINITY, NLPINFINITY	Value returned by a divide-by-zero in a formula [Numerics]	p. 193
XSLP_ITERFALLBACKOPS, SLPITERFALLBACKOPS	Alternative LP level control values for numerically challenging problems [Linearizations, Numerics, SLP]	p. 261
XSLP_SCALE	When to re-scale the SLP problem [Numerics, SLP]	p. 251
XSLP_SCALECOUNT	Iteration limit used in determining when to re-scale the SLP matrix [Numerics, SLP]	p. 251

18.23.3 Numerics attributes

XSLP_TOTALEVALUATIONERRORS, NLPTOTALEVALUATIONERRORS	The total number of evaluation errors during the solve [Numerics]	p. 157
--	--	--------

18.24 Parallel

Reference section for functionality around modern multi-core CPUs. By default, Xpress 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 solver uses the parallel hardware.

18.24.1 Parallel controls

XSLP_CALCTHREADS, NLPCALCTHREADS	Number of threads used for formula and derivatives evaluations [Parallel]	p. 223
XSLP_DETERMINISTIC, NLPDETERMINISTIC	Determines if the parallel features of SLP should be guaranteed to be deterministic [MISLP, Parallel, SLP]	p. 229
XSLP_MULTISTART_THREADS, MULTISTART_THREADS	The maximum number of threads to be used in multistart [Multistart, Parallel]	p. 245
XSLP_THREADS, NLP_THREADS	Default number of threads to be used [Parallel]	p. 253
XSLP_THREADSAFEUSERFUNC, NLP_THREADSAFEUSERFUNC	Defines if user functions are allowed to be called in parallel [Parallel, User Functions]	p. 253

18.25 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.

18.25.1 Presolve library functions

XSLPpostsolve, XPRSnlppostsolve	Restores the problem to its pre-solve state [Presolve]	p. 367
XSLPpresolve	Perform a nonlinear presolve on the problem [Presolve]	p. 368

18.25.2 Presolve controls

XKTR_PARAM_MIP_IMPLICATNS, KNITRO_PARAM_MIP_IMPLICATNS	Specifies whether or not to add constraints to the MIP derived from logical implications. [Knitro-MINLP, Presolve]	p. 468
XKTR_PARAM_PRESOLVE, KNITRO_PARAM_PRESOLVE	Determine whether or not to use the Knitro presolver to try to simplify the model by removing variables or constraints. [Knitro, Presolve]	p. 473
XKTR_PARAM_PRESOLVE_TOL, KNITRO_PARAM_PRESOLVE_TOL	Determines the tolerance used by the Knitro presolver to remove variables and constraints from the model. [Knitro, Presolve, Tolerances]	p. 460
XSLP_BOUNDTHRESHOLD, SLPBOUNDTHRESHOLD	The maximum size of a bound that can be introduced by nonlinear presolve. [Presolve, SLP]	p. 177

XSLP_LINQUADBR, NLPLINQUADBR	Use linear and quadratic constraints and objective function to further reduce bounds on all variables [Presolve]	p. 235
XSLP_POSTSOLVE, NLPPOSTSOLVE	This control determines whether postsolving should be performed automatically [Presolve]	p. 246
XSLP_PRESOLVE, NLPPRESOLVE	This control determines whether presolving should be performed on the nonlinear problem prior to starting the main algorithm [Presolve]	p. 246
XSLP_PRESOLVE_ELIMTOL, NLPPRESOLVE_ELIMTOL	Tolerance for nonlinear eliminations during SLP presolve [Presolve, Tolerances]	p. 205
XSLP_PRESOLVELEVEL, NLPPRESOLVELEVEL	This control determines the level of changes presolve may carry out on the problem and whether column/row indices may change [Presolve]	p. 247
XSLP_PRESOLVEOPS, NLPPRESOLVEOPS	Bitmap indicating the SLP presolve actions to be taken [Bit-vector, Presolve]	p. 247
XSLP_PRESOLVEZERO, NLPPRESOLVEZERO	Minimum absolute value for a variable which is identified as nonzero during SLP presolve [Presolve, Tolerances]	p. 205
XSLP_PROBING, NLPPROBING	This control determines whether probing on a subset of variables should be performed prior to starting the main algorithm. [Presolve]	p. 248
XSLP_REFORMULATE, NLPREFORMULATE	Controls the problem reformulations carried out before augmentation. [Presolve]	p. 248

18.25.3 Presolve attributes

XSLP_PRESOLVEELIMINATIONS, NLPPRESOLVEELIMINATIONS	Number of SLP variables eliminated by XSLPpresolve [Presolve]	p. 153
XSLP_PRESOLVESTATE	Indicates if the problem is presolved [Presolve]	p. 154

18.26 Problem Creation

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

18.26.1 Problem Creation library functions

XSLPcopyprob	Copy an existing SLP problem to another [Problem Creation]	p. 309
XSLPcreateprob	Create a new SLP problem [Problem Creation]	p. 310

XSLPdestroyprob	Delete an SLP problem and release all the associated memory [Problem Creation]	p. 316
XSLPreadprob	Read an Xpress NonLinear extended MPS format matrix from a file into an SLP problem [File IO, Problem Creation]	p. 371

18.27 Problem Information

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

18.27.1 Problem Information library functions

XPRSnlpgetformulastr	Retrieve a single matrix formula in a character string. [Problem Information]	p. 331
XPRSslpgetcoefstr	Retrieve a single matrix coefficient as a formula in a character string. [Problem Information, SLP]	p. 323
XSLPgetcoeffformula, XPRSslpgetcoeffformula	Retrieve a single matrix coefficient as a formula split into tokens. [Problem Information, SLP]	p. 324
XSLPgetcoefs, XPRSslpgetcoefs	Retrieve the list of positions of the nonlinear coefficients in the problem. [Problem Information, SLP]	p. 325
XSLPgetformula, XPRSnlpgetformula	Retrieve a single matrix formula as a formula split into tokens. [Problem Information]	p. 330
XSLPgetformularows, XPRSnlpgetformularows	Retrieve the list of positions of the nonlinear formulas in the problem [Problem Information]	p. 333
XSLPgetindex	Retrieve the index of an Xpress NonLinear entity with a given name [Problem Information, User Functions]	p. 334
XSLPloadcoefs, XPRSslploadcoefs	Load non-linear coefficients into the SLP problem. [Problem Information, SLP]	p. 352
XSLPloadformulas, XPRSnlploadformulas	Load non-linear formulas into the SLP problem [Problem Information]	p. 355
XSLPwriteprob	Write the current problem to a file in extended MPS or text format [File IO, Problem Information]	p. 423

18.27.2 Problem Information attributes

XSLP_COEFFICIENTS, SLP_COEFFICIENTS	Number of nonlinear coefficients [Problem Information, SLP]	p. 146
XSLP_DELTAS, SLPDELTAS	Number of delta vectors created during augmentation [Problem Information, SLP]	p. 146

XSLP_EQUALSCOLUMN, NLP_EQUALSCOLUMN	Index of the reserved "=" column [Problem Information]	p. 147
XSLP_EXPLOREDELTAS, SLPEXPLOREDELTAS	Number of variables with an exploration-type delta set up in the problem [Problem Information, SLP]	p. 146
XSLP_IMPLICITVARIABLES, NLPIMPLICITVARIABLES	Number of SLP variables appearing only in coefficients [Problem Information, SLP]	p. 147
XSLP_INTEGERDELTAS, SLP_INTEGERDELTAS	Number of variables set up with an integer delta in the problem [Problem Information, SLP]	p. 147
XSLP_MINUSPENALTYERRORS, SLPMINUSPENALTYERRORS	Number of negative penalty error vectors [Problem Information, SLP]	p. 148
XSLP_NONLINEARCONSTRAINTS, NONLINEARCONSTRAINTS	Number of nonlinear constraints in the problem [Problem Information]	p. 151
XSLP_ORIGINALCOLS, NLP_ORIGINALCOLS	Number of model columns in the extended original problem [Problem Information]	p. 151
XSLP_ORIGINALROWS, NLP_ORIGINALROWS	Number of model rows in the extended original problem [Problem Information]	p. 151
XSLP_PENALTYDELTACOLUMN, SLPPENALTYDELTACOLUMN	Index of column costing the penalty delta row [Problem Information, SLP]	p. 152
XSLP_PENALTYDELTAROW, SLPPENALTYDELTAROW	Index of equality row holding the penalties for delta vectors [Problem Information, SLP]	p. 152
XSLP_PENALTYDELTAS, SLPPENALTYDELTAS	Number of penalty delta vectors [Problem Information, SLP]	p. 152
XSLP_PENALTYERRORCOLUMN, SLPPENALTYERRORCOLUMN	Index of column costing the penalty error row [Problem Information, SLP]	p. 152
XSLP_PENALTYERRORROW, SLPPENALTYERRORROW	Index of equality row holding the penalties for penalty error vectors [Problem Information, SLP]	p. 153
XSLP_PENALTYERRORS, SLPPENALTYERRORS	Number of penalty error vectors [Problem Information, SLP]	p. 153
XSLP_PLUSPENALTYERRORS, SLPPPLUSPENALTYERRORS	Number of positive penalty error vectors [Problem Information, SLP]	p. 153
XSLP_SEMICONTDLTAS, SLPSEMICONTDLTAS	Number of variables with a minimum perturbation step set up in the problem [Problem Information, SLP]	p. 154

18.28 Problem Modification

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

18.28.1 Problem Modification library functions

<code>XPRSnlpchgformulastr</code>	Add or replace a single matrix formula using a character string for the formula. [Problem Modification]	p. 298
<code>XPRSslpchgcoefstr</code>	Add or change a single matrix coefficient using a character string for the formula. [Problem Modification, SLP]	p. 291
<code>XSLPaddcoefs</code> , <code>XPRSslpaddcoefs</code>	Add non-linear coefficients to the SLP problem. [Problem Modification, SLP]	p. 277
<code>XSLPaddformulas</code> , <code>XPRSnlpaddformulas</code>	Add non-linear formulas to the SLP problem. [Problem Modification]	p. 279
<code>XSLPchgcoef</code> , <code>XPRSslpchgcoef</code>	Add or change a single matrix coefficient using a parsed or unparsed formula. [Problem Modification, SLP]	p. 293
<code>XSLPchgdeltype</code> , <code>XPRSslpchgdeltype</code>	Changes the type of the delta assigned to a nonlinear variable [Problem Modification, SLP]	p. 295
<code>XSLPchgformula</code> , <code>XPRSnlpchgformula</code>	Add or replace a single matrix formula using a parsed or unparsed formula [Problem Modification]	p. 299
<code>XSLPdelcoefs</code> , <code>XPRSslpdelcoefs</code>	Delete coefficients from the current problem. [Problem Modification, SLP]	p. 311
<code>XSLPdelformulas</code> , <code>XPRSnlpdelformulas</code>	Delete nonlinear formulas from the current problem [Problem Modification]	p. 312

18.29 Save Restore

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

18.29.1 Save Restore library functions

<code>XSLPrestore</code>	Restore the Xpress NonLinear problem from a file created by <code>XSLPsave</code> [File IO, Save Restore]	p. 374
<code>XSLPsave</code>	Save the Xpress NonLinear problem to file [File IO, Save Restore]	p. 376
<code>XSLPsaveas</code>	Save the Xpress NonLinear problem to a named file [File IO, Save Restore]	p. 377

18.30 SLP

Reference section for functions, controls, and attributes that are specific to using the sequential linear programming solver SLP.

18.30.1 SLP library functions

<code>XPRSslpchgcoefstr</code>	Add or change a single matrix coefficient using a character string for the formula. [Problem Modification, SLP]	p. 291
<code>XPRSslpgetcoefstr</code>	Retrieve a single matrix coefficient as a formula in a character string. [Problem Information, SLP]	p. 323
<code>XSLPaddcoefs</code> , <code>XPRSslpaddcoefs</code>	Add non-linear coefficients to the SLP problem. [Problem Modification, SLP]	p. 277
<code>XSLPcascade</code> , <code>XPRSslpcascade</code>	Re-calculate consistent values for SLP variables based on the current values of the remaining variables. [Cascading, SLP, Solution Process]	p. 288
<code>XSLPcascadeorder</code> , <code>XPRSslpcascadeorder</code>	Establish a re-calculation sequence for SLP variables with determining rows. [Data Input, SLP]	p. 289
<code>XSLPchgcascadenlimit</code> , <code>XPRSslpchgcascadenlimit</code>	Set a variable specific cascade iteration limit [Data Input, SLP]	p. 290
<code>XSLPchgcoef</code> , <code>XPRSslpchgcoef</code>	Add or change a single matrix coefficient using a parsed or unparsed formula. [Problem Modification, SLP]	p. 293
<code>XSLPchgdeltatype</code> , <code>XPRSslpchgdeltatype</code>	Changes the type of the delta assigned to a nonlinear variable [Problem Modification, SLP]	p. 295
<code>XSLPchgrowstatus</code> , <code>XPRSslpchgrowstatus</code>	Change the status setting of a constraint [Bit-vector, Data Input, SLP]	p. 300
<code>XSLPchgrowwt</code> , <code>XPRSslpchgrowwt</code>	Set or change the initial penalty error weight for a row [Data Input, SLP]	p. 301
<code>XSLPconstruct</code> , <code>XPRSslpconstruct</code>	Create the full augmented SLP matrix and data structures, ready for optimization [SLP, Solution Process]	p. 306
<code>XSLPdelcoefs</code> , <code>XPRSslpdelcoefs</code>	Delete coefficients from the current problem. [Problem Modification, SLP]	p. 311
<code>XSLPgetcoeffformula</code> , <code>XPRSslpgetcoeffformula</code>	Retrieve a single matrix coefficient as a formula split into tokens. [Problem Information, SLP]	p. 324
<code>XSLPgetcoefs</code> , <code>XPRSslpgetcoefs</code>	Retrieve the list of positions of the nonlinear coefficients in the problem. [Problem Information, SLP]	p. 325

XSLPgetcolinfo, XPRSslpgetcolinfo	Get current column information. [SLP, Solution]	p. 326
XSLPgetrowinfo, XPRSslpgetrowinfo	Get current row information. [SLP, Solution]	p. 339
XSLPgetrowstatus, XPRSslpgetrowstatus	Retrieve the status setting of a constraint [Bit-vector, SLP, Solution]	p. 340
XSLPgetrowwt, XPRSslpgetrowwt	Get the initial penalty error weight for a row [Data Information, SLP]	p. 341
XSLPloadcoefs, XPRSslploadcoefs	Load non-linear coefficients into the SLP problem. [Problem Information, SLP]	p. 352
XSLPreinitialize, XPRSslpreinitialize	Reset the SLP problem to match a just augmented system [SLP, Solution Process]	p. 375
XSLPsetcbcascadeend, XPRSsetcbslpcascadeend	Set a user callback to be called at the end of the cascading process, after the last variable has been cascaded [Callback, Cascading, SLP]	p. 379
XSLPsetcbcascadestart, XPRSsetcbslpcascadestart	Set a user callback to be called at the start of the cascading process, before any variables have been cascaded [Callback, Cascading, SLP]	p. 380
XSLPsetcbcascadevar, XPRSsetcbslpcascadevar	Set a user callback to be called after each column has been cascaded [Callback, Cascading, SLP]	p. 381
XSLPsetcbcascadevarfail, XPRSsetcbslpcascadevarfail	Set a user callback to be called after cascading a column was not successful [Callback, Cascading, SLP]	p. 382
XSLPsetcbconstruct, XPRSsetcbslpconstruct	Set a user callback to be called during the Xpress-SLP augmentation process [Callback, SLP]	p. 384
XSLPsetcbdrcol, XPRSsetcbslpdrcol	Set a user callback used to override the update of variables with small determining column [Callback, Cascading, SLP]	p. 387
XSLPsetcbiterend, XPRSsetcbslpiterend	Set a user callback to be called at the end of each SLP iteration [Callback, SLP]	p. 389
XSLPsetcbiterstart, XPRSsetcbslpiterstart	Set a user callback to be called at the start of each SLP iteration [Callback, SLP]	p. 390
XSLPsetcbitervar, XPRSsetcbslpitervar	Set a user callback to be called after each column has been tested for convergence [Callback, SLP, SLP-convergence]	p. 391
XSLPsetcbpresolved	Set a user callback to be called after the nonlinear presolver has been applied. [Callback, SLP]	p. 401

<code>XSLPsetcbpreupdatelinearization</code> , <code>XPRSsetcbslppreupdatelinearization</code>	Set a user callback to be called before the linearization is updated [Callback, SLP]	p. 399
<code>XSLPsetcbslpend</code>	Set a user callback to be called at the end of the SLP optimization [Callback, SLP]	p. 402
<code>XSLPsetcbslpstart</code>	Set a user callback to be called at the start of the SLP optimization [Callback, SLP]	p. 404
<code>XSLPsetdetrow</code> , <code>XPRSslpsetdetrow</code>	Set the determining row of a variable [Cascading, Data Input, SLP]	p. 409
<code>XSLPunconstruct</code> , <code>XPRSslpunconstruct</code>	Removes the augmentation and returns the problem to its pre-linearization state [SLP, Solution Process]	p. 416
<code>XSLPupdatelinearization</code> , <code>XPRSslpupdatelinearization</code>	Updates the current linearization [SLP, Solution Process]	p. 417

18.30.2 SLP controls

<code>XSLP_ALGORITHM</code> , <code>SLPALGORITHM</code>	Bit map describing the SLP algorithm(s) to be used [Bit-vector, SLP]	p. 216
<code>XSLP_ANALYZE</code> , <code>SLPANALYZE</code>	Bit map activating additional options supporting model / solution path analysis [Bit-vector, Logging, SLP]	p. 218
<code>XSLP_ATOL_A</code> , <code>SLPATOL_A</code>	Absolute delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
<code>XSLP_ATOL_R</code> , <code>SLPATOL_R</code>	Relative delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
<code>XSLP_AUGMENTATION</code> , <code>SLPAUGMENTATION</code>	Bit map describing the SLP augmentation method(s) to be used [Bit-vector, SLP]	p. 219
<code>XSLP_AUTOSAVE</code> , <code>SLPAUTOSAVE</code>	Frequency with which to save the model [Logging, SLP]	p. 220
<code>XSLP_BARCROSSOVERSTART</code> , <code>SLPBARCROSSOVERSTART</code>	Default crossover activation behaviour for barrier start [Linearizations, SLP]	p. 221
<code>XSLP_BARLIMIT</code> , <code>SLPBARLIMIT</code>	Number of initial SLP iterations using the barrier method [Limits, Linearizations, SLP]	p. 221
<code>XSLP_BARSTALLINGLIMIT</code> , <code>SLPBARSTALLINGLIMIT</code>	Number of iterations to allow numerical failures in barrier before switching to dual [Limits, Linearizations, SLP]	p. 222
<code>XSLP_BARSTALLINGOBJLIMIT</code> , <code>SLPBARSTALLINGOBJLIMIT</code>	Number of iterations over which to measure the objective change for barrier iterations with no crossover [Limits, Linearizations, SLP]	p. 222

XSLP_BARSTALLINGTOL, SLPBARSTALLINGTOL	Required change in the objective when progress is measured in barrier iterations without crossover [Linearizations, SLP]	p. 176
XSLP_BARSTARTOPS, SLPBARSTARTOPS	Controls behaviour when the barrier is used to solve the linearizations [Linearizations, SLP]	p. 222
XSLP_BOUNDTHRESHOLD, SLPBOUNDTHRESHOLD	The maximum size of a bound that can be introduced by nonlinear presolve. [Presolve, SLP]	p. 177
XSLP_CASCADE, SLPCASCADE	Bit map describing the cascading to be used [Cascading, SLP]	p. 223
XSLP_CASCADENLIMIT, SLPASCADENLIMIT	Maximum number of iterations for cascading with non-linear determining rows [Cascading, Limits, SLP]	p. 224
XSLP_CASCADETOL_PA, SLPCASCADETOL_PA	Absolute cascading print tolerance [Cascading, Logging, SLP]	p. 177
XSLP_CASCADETOL_PR, SLPCASCADETOL_PR	Relative cascading print tolerance [Cascading, Logging, SLP]	p. 177
XSLP_CLAMP SHRINK, SLPCLAMP SHRINK	Shrink ratio used to impose strict convergence on variables converged in extended criteria only [SLP]	p. 179
XSLP_CLAMPVALIDATIONTOL_A, SLPCLAMPVALIDATIONTOL_A	Absolute validation tolerance for applying XSLP_CLAMP SHRINK [SLP, Tolerances]	p. 179
XSLP_CLAMPVALIDATIONTOL_R, SLPCLAMPVALIDATIONTOL_R	Relative validation tolerance for applying XSLP_CLAMP SHRINK [SLP, Tolerances]	p. 179
XSLP_CONVERGENCEOPS, SLPCONVERGENCEOPS	Bit map describing which convergence tests should be carried out [Bit-vector, SLP, SLP-convergence]	p. 225
XSLP_CTOL, SLPCTOL	Closure convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 180
XSLP_DAMP, SLPDAMP	Damping factor for updating values of variables [SLP]	p. 180
XSLP_DAMPEXPAND, SLPDAMPEXPAND	Multiplier to increase damping factor during dynamic damping [SLP]	p. 180
XSLP_DAMP MAX, SLPDAMP MAX	Maximum value for the damping factor of a variable during dynamic damping [SLP]	p. 181
XSLP_DAMP MIN, SLPDAMP MIN	Minimum value for the damping factor of a variable during dynamic damping [SLP]	p. 181
XSLP_DAMP SHRINK, SLPDAMP SHRINK	Multiplier to decrease damping factor during dynamic damping [SLP]	p. 182

XSLP_DAMPSTART, SLPDAMPSTART	SLP iteration at which damping is activated [SLP]	p. 227
XSLP_DEFAULTSTEPBOUND, SLPDEFAULTSTEPBOUND	Minimum initial value for the step bound of an SLP variable if none is explicitly given [SLP]	p. 182
XSLP_DELAYUPDATEROWS, SLPDELAYUPDATEROWS	Number of SLP iterations before update rows are fully activated [SLP]	p. 227
XSLP_DELTACOST, SLPDELTACOST	Initial penalty cost multiplier for penalty delta vectors [SLP]	p. 185
XSLP_DELTACOSTFACTOR, SLPDELTACOSTFACTOR	Factor for increasing cost multiplier on total penalty delta vectors [SLP]	p. 185
XSLP_DELTAFORMAT, SLPDELTAFORMAT	Formatting string for creation of names for SLP delta vectors [Logging, SLP]	p. 261
XSLP_DELTAMAXCOST, SLPDELTAMAXCOST	Maximum penalty cost multiplier for penalty delta vectors [SLP]	p. 185
XSLP_DELTAOFFSET, SLPDELTAOFFSET	Position of first character of SLP variable name used to create name of delta vector [Logging, SLP]	p. 227
XSLP_DELTAZLIMIT, SLPDELTAZLIMIT	Number of SLP iterations during which to apply XSLP_DELTA_Z [SLP]	p. 228
XSLP_DETERMINISTIC, NLPDETERMINISTIC	Determines if the parallel features of SLP should be guaranteed to be deterministic [MISLP, Parallel, SLP]	p. 229
XSLP_DJTOL, SLPDJTOL	Tolerance on DJ value for determining if a variable is at its step bound [SLP, Tolerances]	p. 186
XSLP_DRCOLDJTOL, SLPDRCOLDJTOL	Reduced cost tolerance on the delta variable when fixing due to the determining column being below XSLP_DRCOLTOL. [Cascading, SLP, Tolerances]	p. 186
XSLP_DRCOLTOL, SLPDRCOLTOL	The minimum absolute magnitude of a determining column, for which the determined variable is still regarded as well defined [Cascading, SLP, Tolerances]	p. 186
XSLP_DRFIXRANGE, SLPDRFIXRANGE	The range around the previous value where variables are fixed in cascading if the determining column is below XSLP_DRCOLTOL. [Cascading, SLP]	p. 187
XSLP_ECFCHECK, SLPECFCHECK	Check feasibility at the point of linearization for extended convergence criteria [SLP, SLP-convergence]	p. 229

XSLP_ECFTOL_A, SLPECFTOL_A	Absolute tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187
XSLP_ECFTOL_R, SLPECFTOL_R	Relative tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187
XSLP_ENFORCECOSTSHRINK, SLPENFORCECOSTSHRINK	Factor by which to decrease the current penalty multiplier when enforcing rows. [SLP]	p. 188
XSLP_ENFORCEMAXCOST, SLPENFORCEMAXCOST	Maximum penalty cost in the objective before enforcing most violating rows [SLP]	p. 188
XSLP_ERRORCOST, SLPERRORCOST	Initial penalty cost multiplier for penalty error vectors [SLP]	p. 189
XSLP_ERRORCOSTFACTOR, SLPERRORCOSTFACTOR	Factor for increasing cost multiplier on total penalty error vectors [SLP]	p. 189
XSLP_ERRORMAXCOST, SLPERRORMAXCOST	Maximum penalty cost multiplier for penalty error vectors [SLP]	p. 189
XSLP_ERROROFFSET, SLPERROROFFSET	Position of first character of constraint name used to create name of penalty error vectors [Logging, SLP]	p. 230
XSLP_ERRORTOL_A, SLPERRORTOL_A	Absolute tolerance for error vectors [SLP, Tolerances]	p. 190
XSLP_ERRORTOL_P, SLPERRORTOL_P	Absolute tolerance for printing error vectors [SLP, Tolerances]	p. 190
XSLP_ESCALATION, SLPESCALATION	Factor for increasing cost multiplier on individual penalty error vectors [SLP]	p. 190
XSLP_ETOL_A, SLPETOL_A	Absolute tolerance on penalty vectors [SLP, Tolerances]	p. 191
XSLP_ETOL_R, SLPETOL_R	Relative tolerance on penalty vectors [SLP, Tolerances]	p. 191
XSLP_EVTOL_A, SLPEVTOL_A	Absolute tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 191
XSLP_EVTOL_R, SLPEVTOL_R	Relative tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 192
XSLP_EXPAND, SLPEXPAND	Multiplier to increase a step bound [SLP]	p. 192
XSLP_FEASTOLTARGET, SLPFEASTOLTARGET	When set, this defines a target feasibility tolerance to which the linearizations are solved to [Linearizations, SLP, Tolerances]	p. 193
XSLP_FILTER, SLPFILTER	Bit map for controlling solution updates [Bit-vector, SLP, Solution]	p. 231

XSLP_GRANULARITY, SLPGRANULARITY	Base for calculating penalty costs [SLP]	p. 193
XSLP_GRIDHEURSELECT, SLPGRIDHEURSELECT	Bit map selectin which heuristics to run if the problem has variable with an integer delta [Heuristics, SLP]	p. 233
XSLP_HEURSTRATEGY, SLPHEURSTRATEGY	Branch and Bound: This specifies the MINLP heuristic strategy. [Heuristics, SLP]	p. 233
XSLP_INFEASLIMIT, SLPINFEASLIMIT	The maximum number of consecutive infeasible SLP iterations which can occur before Xpress-SLP terminates [Limits, SLP]	p. 234
XSLP_ITERFALLBACKOPS, SLPITERFALLBACKOPS	Alternative LP level control values for numerically challengeing problems [Linearizations, Numerics, SLP]	p. 261
XSLP_ITERLIMIT, SLPITERLIMIT	The maximum number of SLP iterations [Limits, SLP]	p. 234
XSLP_ITOL_A, SLPITOL_A	Absolute impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_ITOL_R, SLPITOL_R	Relative impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_LOG, NLPLOG	Level of printing during SLP iterations [Logging, SLP]	p. 236
XSLP_LSITERLIMIT, SLPLSITERLIMIT	Number of iterations in the line search [Limits, SLP]	p. 236
XSLP_LSPATTERNLIMIT, SLPLSPATTERNLIMIT	Number of iterations in the pattern search preceding the line search [Limits, SLP]	p. 236
XSLP_LSSTART, SLPLSSTART	Iteration in which to active the line search [SLP]	p. 237
XSLP_LSZEROLIMIT, SLPLSZEROLIMIT	Maximum number of zero length line search steps before line search is deactivated [Limits, SLP]	p. 237
XSLP_MATRIXTOL, SLPMATRIXTOL	Nonzero tolerance for dropping coefficients from the linearization. [SLP]	p. 195
XSLP_MAXWEIGHT, SLPMAXWEIGHT	Maximum penalty weight for delta or error vectors [SLP]	p. 195
XSLP_MERITLAMBD, NLPMERITLAMBD	Factor by which the net objective is taken into account in the merit function [SLP, Solution]	p. 196
XSLP_MINSBFACTOR, SLPMINSBFACTOR	Factor by which step bounds can be decreased beneath XSLP_ATOL_A [SLP]	p. 197

XSLP_MINUSDELTAFORMAT, SLPMINUSDELTAFORMAT	Formatting string for creation of names for SLP negative penalty delta vectors [Logging, SLP]	p. 262
XSLP_MINUSERERRORFORMAT, SLPMINUSERERRORFORMAT	Formatting string for creation of names for SLP negative penalty error vectors [Logging, SLP]	p. 262
XSLP_MINWEIGHT, SLPMINWEIGHT	Minimum penalty weight for delta or error vectors [SLP]	p. 197
XSLP_MTOL_A, SLPMTOL_A	Absolute effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 200
XSLP_MTOL_R, SLPMTOL_R	Relative effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 201
XSLP_MVTOL, SLPMVTOL	Marginal value tolerance for determining if a constraint is slack [SLP, SLP-convergence, Tolerances]	p. 202
XSLP_OBJTHRESHOLD, SLPOBJTHRESHOLD	Assumed maximum value of the objective function in absolute value. [SLP]	p. 202
XSLP_OBJTOPENALTYCOST, SLPOBJTOPENALTYCOST	Factor to estimate initial penalty costs from objective function [SLP]	p. 203
XSLP_OCOUNT, SLPOCOUNT	Number of SLP iterations over which to measure objective function variation for static objective (2) convergence criterion [SLP, SLP-convergence]	p. 245
XSLP_OPTIMALITYTOLTARGET, SLOPTIMALITYTOLTARGET	When set, this defines a target optimality tolerance to which the linearizations are solved to [Linearizations, SLP, Tolerances]	p. 203
XSLP_OTOL_A, SLPOTOL_A	Absolute static objective (2) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 204
XSLP_OTOL_R, SLPOTOL_R	Relative static objective (2) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 204
XSLP_PENALTYCOLFORMAT, SLPPENALTYCOLFORMAT	Formatting string for creation of the names of the SLP penalty transfer vectors [Logging, SLP]	p. 263
XSLP_PENALTYINFOSTART, SLPPENALTYINFOSTART	Iteration from which to record row penalty information [Logging, SLP]	p. 246
XSLP_PENALTYROWFORMAT, SLPPENALTYROWFORMAT	Formatting string for creation of the names of the SLP penalty rows [Logging, SLP]	p. 263
XSLP_PLUSDELTAFORMAT, SLPPLUSDELTAFORMAT	Formatting string for creation of names for SLP positive penalty delta vectors [Logging, SLP]	p. 263
XSLP_PLUSERERRORFORMAT, SLPPLUSERERRORFORMAT	Formatting string for creation of names for SLP positive penalty error vectors [Logging, SLP]	p. 264

XSLP_SAMECOUNT, SLPSAMECOUNT	Number of steps reaching the step bound in the same direction before step bounds are increased [SLP]	p. 249
XSLP_SAMEDAMP, SLPSAMEDAMP	Number of steps in same direction before damping factor is increased [SLP]	p. 250
XSLP_SBLOROWFORMAT, SLPSBLOROWFORMAT	Formatting string for creation of names for SLP lower step bound rows [Logging, SLP]	p. 264
XSLP_SBNAME, SLPSBNAME	Name of the set of initial step bounds to be used [Logging, SLP]	p. 265
XSLP_SBROWOFFSET, SLPSBROWOFFSET	Position of first character of SLP variable name used to create name of SLP lower and upper step bound rows [Logging, SLP]	p. 250
XSLP_SBSTART, SLPSBSTART	SLP iteration after which step bounds are first applied [SLP]	p. 250
XSLP_SBUPROWFORMAT, SLPSBUPROWFORMAT	Formatting string for creation of names for SLP upper step bound rows [Logging, SLP]	p. 265
XSLP_SCALE	When to re-scale the SLP problem [Numerics, SLP]	p. 251
XSLP_SCALECOUNT	Iteration limit used in determining when to re-scale the SLP matrix [Numerics, SLP]	p. 251
XSLP_SHRINK, SLPSHRINK	Multiplier to reduce a step bound [SLP]	p. 206
XSLP_SHRINKBIAS, SLPSHRINKBIAS	Defines an overwrite / adjustment of step bounds for improving iterations [SLP]	p. 206
XSLP_SLPLOG	Frequency with which SLP status is printed [Logging, SLP]	p. 252
XSLP_STOL_A, SLPSTOL_A	Absolute slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_STOL_R, SLPSTOL_R	Relative slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_STOPOUTOFRANGE, NLPSTOPOUTOFRANGE	Stop optimization and return error code if internal function argument is out of range [SLP]	p. 252
XSLP_TOLNAME, SLPTOLNAME	Name of the set of tolerance sets to be used [File IO, SLP]	p. 265
XSLP_TRACEMASK, SLPTRACEMASK	Mask of variable or row names that are to be traced through the SLP iterates [Logging, SLP]	p. 266
XSLP_TRACEMASKOPS, SLPTRACEMASKOPS	Controls the information printed for XSLP_TRACEMASK. [Bit-vector, Logging, SLP]	p. 253

XSLP_UNFINISHEDLIMIT, SLPUNFINISHEDLIMIT	The number of consecutive SLP iterations that may have an unfinished status before the solve is terminated. [Linearizations, SLP]	p. 254
XSLP_UPDATEFORMAT, SLPUPDATEFORMAT	Formatting string for creation of names for SLP update rows [Logging, SLP]	p. 266
XSLP_UPDATEOFFSET, SLPUPDATEOFFSET	Position of first character of SLP variable name used to create name of SLP update row [Logging, SLP]	p. 255
XSLP_VALIDATIONFACTOR, NLPVALIDATIONFACTOR	Minimum improvement in validation targets to continue iterating [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_K, NLPVALIDATIONTARGET_K	Optimality target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_R, NLPVALIDATIONTARGET_R	Feasibility target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTOL_A, NLPVALIDATIONTOL_A	Absolute tolerance for the XSLPvalidate procedure [SLP, Tolerances]	p. 209
XSLP_VALIDATIONTOL_K, NLPVALIDATIONTOL_K	Relative tolerance for the XSLPvalidatekkt procedure [SLP, Tolerances]	p. 209
XSLP_VCOUNT, SLPVCOUNT	Number of SLP iterations over which to measure static objective (3) convergence [SLP, SLP-convergence]	p. 255
XSLP_VLIMIT, SLPVLIMIT	Number of SLP iterations after which static objective (3) convergence testing starts [SLP, SLP-convergence]	p. 256
XSLP_VTOL_A, SLPVTOL_A	Absolute static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 210
XSLP_VTOL_R, SLPVTOL_R	Relative static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WCOUNT, SLPWCOUNT	Number of SLP iterations over which to measure the objective for the extended convergence continuation criterion [SLP, SLP-convergence]	p. 256
XSLP_WTOL_A, SLPWTOL_A	Absolute extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WTOL_R, SLPWTOL_R	Relative extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 212
XSLP_XCOUNT, SLPXCOUNT	Number of SLP iterations over which to measure static objective (1) convergence [SLP, SLP-convergence]	p. 257
XSLP_XLIMIT, SLPXLIMIT	Number of SLP iterations up to which static objective (1) convergence testing is performed [Limits, SLP, SLP-convergence]	p. 258

XSLP_XTOL_A, SLPXTOL_A	Absolute static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 213
XSLP_XTOL_R, SLPXTOL_R	Relative static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 214
XSLP_ZEROCRITERION, SLPZEROCRITERION	Bitmap determining the behavior of the placeholder deletion procedure [Bit-vector, SLP]	p. 259
XSLP_ZEROCRITERIONCOUNT, SLPZEROCRITERIONCOUNT	Number of consecutive times a placeholder entry is zero before being considered for deletion [Limits, SLP]	p. 259
XSLP_ZEROCRITERIONSTART, SLPZEROCRITERIONSTART	SLP iteration at which criteria for deletion of placeholder entries are first activated. [SLP]	p. 260

18.30.3 SLP attributes

XSLP_COEFFICIENTS, SLP_COEFFICIENTS	Number of nonlinear coefficients [Problem Information, SLP]	p. 146
XSLP_CURRENTDELTACOST, SLP_CURRENTDELTACOST	Current value of penalty cost multiplier for penalty delta vectors [SLP, Solution]	p. 142
XSLP_CURRENTERRORCOST, SLP_CURRENTERRORCOST	Current value of penalty cost multiplier for penalty error vectors [SLP, Solution]	p. 142
XSLP_DELTAS, SLP_DELTAS	Number of delta vectors created during augmentation [Problem Information, SLP]	p. 146
XSLP_ECFCOUNT, SLP_ECFCOUNT	Number of infeasible constraints found at the point of linearization [Linearizations, SLP]	p. 146
XSLP_ERRORCOSTS, SLP_ERRORCOSTS	Total penalty costs in the solution [SLP, Solution]	p. 142
XSLP_EXPLOREDELTAS, SLP_EXPLOREDELTAS	Number of variables with an exploration-type delta set up in the problem [Problem Information, SLP]	p. 146
XSLP_IMPLICITVARIABLES, SLP_IMPLICITVARIABLES	Number of SLP variables appearing only in coefficients [Problem Information, SLP]	p. 147
XSLP_INTEGERDELTAS, SLP_INTEGERDELTAS	Number of variables set up with an integer delta in the problem [Problem Information, SLP]	p. 147
XSLP_KEEPPBESTITER, SLP_KEEPPBESTITER	The iteration in which the returned solution has been found. [SLP, Solution]	p. 148
XSLP_MINUSPENALTYERRORS, SLP_MINUSPENALTYERRORS	Number of negative penalty error vectors [Problem Information, SLP]	p. 148

XSLP_MODELCOLS, NLP_MODELCOLS	Number of model columns in the problem	p. 149
[SLP]		
XSLP_MODELROWS, NLP_MODELROWS	Number of model rows in the problem	p. 150
[SLP]		
XSLP_NONCONSTANTCOEFFS, SLP_NONCONSTANTCOEFFS	Number of coefficients in the augmented problem that might change between SLP iterations	p. 151
[SLP]		
XSLP_PENALTYDELTACOLUMN, SLPPENALTYDELTACOLUMN	Index of column costing the penalty delta row	p. 152
[Problem Information, SLP]		
XSLP_PENALTYDELTAROW, SLPPENALTYDELTAROW	Index of equality row holding the penalties for delta vectors	p. 152
[Problem Information, SLP]		
XSLP_PENALTYDELTAS, SLPPENALTYDELTAS	Number of penalty delta vectors	p. 152
[Problem Information, SLP]		
XSLP_PENALTYDELTATOTAL, SLPPENALTYDELTATOTAL	Total activity of penalty delta vectors	p. 143
[SLP, Solution]		
XSLP_PENALTYDELTAVALUE, SLPPENALTYDELTAVALUE	Total penalty cost attributed to penalty delta vectors	p. 143
[SLP, Solution]		
XSLP_PENALTYERRORCOLUMN, SLPPENALTYERRORCOLUMN	Index of column costing the penalty error row	p. 152
[Problem Information, SLP]		
XSLP_PENALTYERRORROW, SLPPENALTYERRORROW	Index of equality row holding the penalties for penalty error vectors	p. 153
[Problem Information, SLP]		
XSLP_PENALTYERRORS, SLPPENALTYERRORS	Number of penalty error vectors	p. 153
[Problem Information, SLP]		
XSLP_PENALTYERRORTOTAL, SLPPENALTYERRORTOTAL	Total activity of penalty error vectors	p. 143
[SLP, Solution]		
XSLP_PENALTYERRORVALUE, SLPPENALTYERRORVALUE	Total penalty cost attributed to penalty error vectors	p. 143
[SLP, Solution]		
XSLP_PLUSPENALTYERRORS, SLPPUSPENALTYERRORS	Number of positive penalty error vectors	p. 153
[Problem Information, SLP]		
XSLP_SBXCONVERGED, SLP_SBXCONVERGED	Number of step-bounded variables converged only on extended criteria	p. 154
[SLP, SLP-convergence]		
XSLP_SEMICONTDeltas, SLP_SEMICONTDeltas	Number of variables with a minimum perturbation step set up in the problem	p. 154
[Problem Information, SLP]		
XSLP_UCCONSTRAINEDCOUNT, SLP_UCCONSTRAINEDCOUNT	Number of unconverged variables with coefficients in constraining rows	p. 158
[SLP, SLP-convergence]		

XSLP_UNCONVERGED, SLPUNCONVERGED	Number of unconverged values [SLP, SLP-convergence]	p. 158
XSLP_VARIABLES, NLPVARIABLES	Number of SLP variables [Data Information, SLP]	p. 159
XSLP_VSOLINDEX	Vertex solution index [SLP, Solution]	p. 145
XSLP_ZEROESRESET, SLPZEROESRESET	Number of placeholder entries set to zero [SLP]	p. 160
XSLP_ZEROESRETAINED, SLPZEROESRETAINED	Number of potentially zero placeholders left untouched [SLP]	p. 160
XSLP_ZEROESTOTAL, SLPZEROESTOTAL	Number of potential zero placeholder entries [SLP]	p. 160

18.31 SLP-convergence

Reference section for functions, controls, and attributes for convergence criteria within the sequential linear programming solver SLP.

18.31.1 SLP-convergence library functions

XSLPsetcbitervar, XPRSsetcbslpitervar	Set a user callback to be called after each column has been tested for convergence [Callback, SLP, SLP-convergence]	p. 391
---------------------------------------	---	--------

18.31.2 SLP-convergence controls

XSLP_ATOL_A, SLPATOL_A	Absolute delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
XSLP_ATOL_R, SLPATOL_R	Relative delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
XSLP_CONVERGENCEOPS, SLPCONVERGENCEOPS	Bit map describing which convergence tests should be carried out [Bit-vector, SLP, SLP-convergence]	p. 225
XSLP_CTOL, SLPCTOL	Closure convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 180
XSLP_ECFCHECK, SLPECFCHECK	Check feasibility at the point of linearization for extended convergence criteria [SLP, SLP-convergence]	p. 229
XSLP_ECFTOL_A, SLPECFTOL_A	Absolute tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187
XSLP_ECFTOL_R, SLPECFTOL_R	Relative tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187

XSLP_EVTOL_A, SLPEVTOL_A	Absolute tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 191
XSLP_EVTOL_R, SLPEVTOL_R	Relative tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 192
XSLP_ITOL_A, SLPITOL_A	Absolute impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_ITOL_R, SLPITOL_R	Relative impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_MIPOCOUNT, SLPMIPOCOUNT	Number of SLP iterations at each node over which to measure objective function variation [MISLP, SLP-convergence]	p. 241
XSLP_MIPOTOL_A, SLPMIPOTOL_A	Absolute objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MIPOTOL_R, SLPMIPOTOL_R	Relative objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MTOL_A, SLPMTOL_A	Absolute effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 200
XSLP_MTOL_R, SLPMTOL_R	Relative effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 201
XSLP_MVTOL, SLPMVTOL	Marginal value tolerance for determining if a constraint is slack [SLP, SLP-convergence, Tolerances]	p. 202
XSLP_OCOUNT, SLPOCOUNT	Number of SLP iterations over which to measure objective function variation for static objective (2) convergence criterion [SLP, SLP-convergence]	p. 245
XSLP_OTOL_A, SLPOTOL_A	Absolute static objective (2) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 204
XSLP_OTOL_R, SLPOTOL_R	Relative static objective (2) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 204
XSLP_STOL_A, SLPSTOL_A	Absolute slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_STOL_R, SLPSTOL_R	Relative slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_VALIDATIONFACTOR, NLPVALIDATIONFACTOR	Minimum improvement in validation targets to continue iterating [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_K, NLPVALIDATIONTARGET_K	Optimality target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_R, NLPVALIDATIONTARGET_R	Feasibility target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VCOUNT, SLPVCOUNT	Number of SLP iterations over which to measure static objective (3) convergence [SLP, SLP-convergence]	p. 255

XSLP_VLIMIT, SLPVLIMIT	Number of SLP iterations after which static objective (3) convergence testing starts [SLP, SLP-convergence]	p. 256
XSLP_VTOL_A, SLPVTOL_A	Absolute static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 210
XSLP_VTOL_R, SLPVTOL_R	Relative static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WCOUNT, SLPWCOUNT	Number of SLP iterations over which to measure the objective for the extended convergence continuation criterion [SLP, SLP-convergence]	p. 256
XSLP_WTOL_A, SLPWTOL_A	Absolute extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WTOL_R, SLPWTOL_R	Relative extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 212
XSLP_XCOUNT, SLPXCOUNT	Number of SLP iterations over which to measure static objective (1) convergence [SLP, SLP-convergence]	p. 257
XSLP_XLIMIT, SLPXLIMIT	Number of SLP iterations up to which static objective (1) convergence testing is performed [Limits, SLP, SLP-convergence]	p. 258
XSLP_XTOL_A, SLPXTOL_A	Absolute static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 213
XSLP_XTOL_R, SLPXTOL_R	Relative static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 214

18.31.3 SLP-convergence attributes

XSLP_SBXCONVERGED, SLPSBXCONVERGED	Number of step-bounded variables converged only on extended criteria [SLP, SLP-convergence]	p. 154
XSLP_UCCONSTRAINEDCOUNT, SLPUCCONSTRAINEDCOUNT	Number of unconverged variables with coefficients in constraining rows [SLP, SLP-convergence]	p. 158
XSLP_UNCONVERGED, SLPUNCONVERGED	Number of unconverged values [SLP, SLP-convergence]	p. 158

18.32 Solution Process

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

18.32.1 Solution Process library functions

XSLPcascade, XPRSslpcascade	Re-calculate consistent values for SLP variables based on the current values of the remaining variables. [Cascading, SLP, Solution Process]	p. 288
-----------------------------	--	--------

XSLPconstruct, XPRSslpconstruct	Create the full augmented SLP matrix and data structures, ready for optimization [SLP, Solution Process]	p. 306
XSLPfixpenalties, XPRSslpfixpenalties	Fixe the values of the error vectors [Solution Process]	p. 319
XSLPinterrupt	Interrupts the current SLP optimization [Solution Process]	p. 350
XSLPnlptimize, XPRSslptimize	Maximize or minimize an SLP problem [Solution Process]	p. 366
XSLPreinitialize, XPRSslpreinitialize	Reset the SLP problem to match a just augmented system [SLP, Solution Process]	p. 375
XSLPunconstruct, XPRSslpunconstruct	Removes the augmentation and returns the problem to its pre-linearization state [SLP, Solution Process]	p. 416
XSLPupdatelinearization, XPRSslupdatelinearization	Updates the current linearization [SLP, Solution Process]	p. 417

18.32.2 Solution Process controls

XKTR_PARAM_ALGORITHM, KNITRO_PARAM_ALGORITHM	Indicates which algorithm to use to solve the problem [Knitro, Solution Process]	p. 461
XKTR_PARAM_MIP_LPALG, KNITRO_PARAM_MIP_LPALG	Specifies which algorithm to use for any linear programming (LP) subproblem solves that may occur in the MIP branch and bound procedure. [Knitro-MINLP, Solution Process]	p. 469
XKTR_PARAM_MIP_METHOD, KNITRO_PARAM_MIP_METHOD	Specifies which MIP method to use. [Knitro-MINLP, Solution Process]	p. 470
XKTR_PARAM_MIP_ROOTALG, KNITRO_PARAM_MIP_ROOTALG	Specifies which algorithm to use for the root node solve in MIP (same options as XKTR_PARAM_ALGORITHM user option). [Knitro-MINLP, Solution Process]	p. 471
XSLP_NLPSOLVER, NLPSOLVER	Controls whether to call FICO Xpress Global or one of the local solvers [Solution Process]	p. 244
XSLP_SOLVER, LOCALSOLVER	Selects the library to use for local solves [Solution Process]	p. 252

18.32.3 Solution Process attributes

XSLP_ITER, SLPITER	SLP iteration count [Solution Process]	p. 148
XSLP_NLPSTATUS, NLPSTATUS	The solution status of the problem. [Solution Process]	p. 150

XSLP_OPTTIME, NLPOPTTIME	Time spent in optimization [Solution Process]	p. 142
XSLP_SOLVERSELECTED, LOCALSOLVERSELECTED	Includes information of which Xpress solver has been used to solve the problem [Solution Process]	p. 154
XSLP_STATUS, SLPSTATUS	Bitmap holding the problem convergence status [Solution Process]	p. 155
XSLP_STOPSTATUS, NLPSTOPSTATUS	Status of the optimization process. [Solution Process]	p. 157

18.33 Solution

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

18.33.1 Solution library functions

XSLPcalclacks, XPRSnlpcalclacks	Calculate the slack values for the provided solution in the non-linear problem [Solution]	p. 287
XSLPevaluatecoef, XPRSslpevaluatecoef	Evaluate a coefficient using the current values of the variables [Solution]	p. 317
XSLPevaluateformula, XPRSnlpevaluateformula	Evaluate a formula using the current values of the variables [Solution]	p. 318
XSLPgetcolinfo, XPRSslpgetcolinfo	Get current column information. [SLP, Solution]	p. 326
XSLPgetrowinfo, XPRSslpgetrowinfo	Get current row information. [SLP, Solution]	p. 339
XSLPgetrowstatus, XPRSslpgetrowstatus	Retrieve the status setting of a constraint [Bit-vector, SLP, Solution]	p. 340
XSLPvalidate, XPRSnlpvalidate	Validate the feasibility of constraints in a converged solution [Solution]	p. 418
XSLPvalidatekkt, XPRSnlpvalidatekkt	Validates the first order optimality conditions also known as the Karush-Kuhn-Tucker (KKT) conditions versus the current solution [Solution]	p. 419
XSLPvalidateprob, XPRSnlpvalidateprob	Validates the current problem formulation and statement [Solution]	p. 420
XSLPvalidaterow, XPRSnlpvalidaterow	Prints an extensive analysis on a given constraint of the SLP problem [Solution]	p. 421

XSLPvalidatevector, XPRSnlpvalidatevector	Validate the feasibility of constraints for a given solution [Solution]	p. 422
XSLPwriteslxsol	Write the current solution to an MPS like file format [File IO, Solution]	p. 424

18.33.2 Solution controls

XSLP_FILTER, SLPFILTER	Bit map for controlling solution updates [Bit-vector, SLP, Solution]	p. 231
XSLP_MERITLAMBDA, NLPMERITLAMBDA	Factor by which the net objective is taken into account in the merit function [SLP, Solution]	p. 196

18.33.3 Solution attributes

XSLP_CURRENTDELTACOST, SLPCURRENTDELTACOST	Current value of penalty cost multiplier for penalty delta vectors [SLP, Solution]	p. 142
XSLP_CURRENTERRORCOST, SLPCURRENTERRORCOST	Current value of penalty cost multiplier for penalty error vectors [SLP, Solution]	p. 142
XSLP_ERRORCOSTS, SLPERRORCOSTS	Total penalty costs in the solution [SLP, Solution]	p. 142
XSLP_KEEPESTITER, NLPKEEPESTITER	The iteration in which the returned solution has been found. [SLP, Solution]	p. 148
XSLP_MIPSOLS, SLPMIPSOLS	Number of integer solutions found in MISLP. [MISLP, Solution]	p. 149
XSLP_OBJVAL, NLPOBJVAL	Objective function value excluding any penalty costs [Solution]	p. 143
XSLP_PENALTYDELTATOTAL, SLPPENALTYDELTATOTAL	Total activity of penalty delta vectors [SLP, Solution]	p. 143
XSLP_PENALTYDELTAVALUE, SLPPENALTYDELTAVALUE	Total penalty cost attributed to penalty delta vectors [SLP, Solution]	p. 143
XSLP_PENALTYERRORTOTAL, SLPPENALTYERRORTOTAL	Total activity of penalty error vectors [SLP, Solution]	p. 143
XSLP_PENALTYERRORVALUE, SLPPENALTYERRORVALUE	Total penalty cost attributed to penalty error vectors [SLP, Solution]	p. 143
XSLP_SOLSTATUS, NLPSTATUS	Indicates the type of solution returned by the solver. [Solution]	p. 155
XSLP_VALIDATIONINDEX_A, NLPVALIDATIONINDEX_A	Absolute validation index [Solution]	p. 144

XSLP_VALIDATIONINDEX_K, NLPVALIDATIONINDEX_K	Relative first order optimality validation index [Solution]	p. 144
XSLP_VALIDATIONINDEX_R, NLPVALIDATIONINDEX_R	Relative validation index [Solution]	p. 144
XSLP_VALIDATIONNETOBJ, NLPVALIDATIONNETOBJ	Net objective as calculated by validation [Solution]	p. 144
XSLP_VALIDATIONSTATUS, NLPVALIDATIONSTATUS	Feasibility status of the current solution. [Solution]	p. 159
XSLP_VSOLINDEX	Vertex solution index [SLP, Solution]	p. 145

18.34 Tolerances

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

18.34.1 Tolerances controls

XKTR_PARAM_BAR_FEASMODETOL, KNITRO_PARAM_BAR_FEASMODETOL	Specifies the tolerance in equation that determines whether Knitro will force subsequent iterates to remain feasible. [Knitro, Tolerances]	p. 457
XKTR_PARAM_FEASTOL, KNITRO_PARAM_FEASTOL	Specifies the final relative stopping tolerance for the feasibility error. [Knitro, Tolerances]	p. 457
XKTR_PARAM_FEASTOLABS, KNITRO_PARAM_FEASTOLABS	Specifies the final absolute stopping tolerance for the feasibility error. [Knitro, Tolerances]	p. 458
XKTR_PARAM_INFEASTOL, KNITRO_PARAM_INFEASTOL	Specifies the (relative) tolerance used for declaring infeasibility of a model. [Knitro, Tolerances]	p. 458
XKTR_PARAM_MIP_INTEGERTOL, KNITRO_PARAM_INTEGERTOL	This value specifies the threshold for deciding whether or not a variable is determined to be an integer. [Knitro, Knitro-MINLP, Tolerances]	p. 458
XKTR_PARAM_MIP_INTGAPABS, KNITRO_PARAM_INTGAPABS	The absolute integrality gap stop tolerance for MIP. [Knitro, Knitro-MINLP, Tolerances]	p. 458
XKTR_PARAM_MIP_INTGAPREL, KNITRO_PARAM_INTGAPREL	The relative integrality gap stop tolerance for MIP. [Knitro, Knitro-MINLP, Tolerances]	p. 459
XKTR_PARAM_OPTTOL, KNITRO_PARAM_OPTTOL	Specifies the final relative stopping tolerance for the KKT (optimality) error. [Knitro, Tolerances]	p. 459
XKTR_PARAM_OPTTOLABS, KNITRO_PARAM_OPTTOLABS	Specifies the final absolute stopping tolerance for the KKT (optimality) error. [Knitro, Tolerances]	p. 459

XKTR_PARAM_PRESOLVE_TOL, KNITRO_PARAM_PRESOLVE_TOL	Determines the tolerance used by the Knitro presolver to remove variables and constraints from the model. [Knitro, Presolve, Tolerances]	p. 460
XKTR_PARAM_XTOL, KNITRO_PARAM_XTOL	The optimization process will terminate if the relative change in all components of the solution point estimate is less than xtol. [Knitro, Tolerances]	p. 460
XSLP_ATOL_A, SLPATOL_A	Absolute delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
XSLP_ATOL_R, SLPATOL_R	Relative delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
XSLP_CDTOL_A, SLP_CDTOL_A	Absolute tolerance for deducing constant derivatives [Tolerances]	p. 178
XSLP_CDTOL_R, SLP_CDTOL_R	Relative tolerance for deducing constant derivatives [Tolerances]	p. 178
XSLP_CLAMPVALIDATIONTOL_A, SLP_CLAMPVALIDATIONTOL_A	Absolute validation tolerance for applying XSLP_CLAMP_SHRINK [SLP, Tolerances]	p. 179
XSLP_CLAMPVALIDATIONTOL_R, SLP_CLAMPVALIDATIONTOL_R	Relative validation tolerance for applying XSLP_CLAMP_SHRINK [SLP, Tolerances]	p. 179
XSLP_CTOL, SLP_CTOL	Closure convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 180
XSLP_DELTA_Z, SLP_DELTA_Z	Tolerance used when calculating derivatives [Derivatives, Tolerances]	p. 184
XSLP_DELTA_ZERO, SLP_DELTA_ZERO	Absolute zero acceptance tolerance used when calculating derivatives [Derivatives, Tolerances]	p. 184
XSLP_DJTOL, SLP_DJTOL	Tolerance on DJ value for determining if a variable is at its step bound [SLP, Tolerances]	p. 186
XSLP_DRCOLDJTOL, SLP_DRCOLDJTOL	Reduced cost tolerance on the delta variable when fixing due to the determining column being below XSLP_DRCOLTOL. [Cascading, SLP, Tolerances]	p. 186
XSLP_DRCOLTOL, SLP_DRCOLTOL	The minimum absolute magnitude of a determining column, for which the determined variable is still regarded as well defined [Cascading, SLP, Tolerances]	p. 186
XSLP_ECFTOL_A, SLP_ECFTOL_A	Absolute tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187
XSLP_ECFTOL_R, SLP_ECFTOL_R	Relative tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187

XSLP_ERRORTOL_A, SLPERRORTOL_A	Absolute tolerance for error vectors [SLP, Tolerances]	p. 190
XSLP_ERRORTOL_P, SLPERRORTOL_P	Absolute tolerance for printing error vectors [SLP, Tolerances]	p. 190
XSLP_ETOL_A, SLPETOL_A	Absolute tolerance on penalty vectors [SLP, Tolerances]	p. 191
XSLP_ETOL_R, SLPETOL_R	Relative tolerance on penalty vectors [SLP, Tolerances]	p. 191
XSLP_EVTOL_A, SLPEVTOL_A	Absolute tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 191
XSLP_EVTOL_R, SLPEVTOL_R	Relative tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 192
XSLP_FEASTOLTARGET, SLPFEASTOLTARGET	When set, this defines a target feasibility tolerance to which the linearizations are solved to [Linearizations, SLP, Tolerances]	p. 193
XSLP_ITOL_A, SLPITOL_A	Absolute impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_ITOL_R, SLPITOL_R	Relative impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_MIPCUTOFF_A, SLPMIPCUTOFF_A	Absolute objective function cutoff for MIP termination [MISLP, Tolerances]	p. 197
XSLP_MIPCUTOFF_R, SLPMIPCUTOFF_R	Absolute objective function cutoff for MIP termination [MISLP, Tolerances]	p. 198
XSLP_MIPERRORTOL_A, SLPMIPERRORTOL_A	Absolute penalty error cost tolerance for MIP cut-off [MISLP, Tolerances]	p. 198
XSLP_MIPERRORTOL_R, SLPMIPERRORTOL_R	Relative penalty error cost tolerance for MIP cut-off [MISLP, Tolerances]	p. 199
XSLP_MIPOTOL_A, SLPMIPOTOL_A	Absolute objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MIPOTOL_R, SLPMIPOTOL_R	Relative objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MTOL_A, SLPMTOL_A	Absolute effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 200
XSLP_MTOL_R, SLPMTOL_R	Relative effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 201
XSLP_MVTOL, SLPMVTOL	Marginal value tolerance for determining if a constraint is slack [SLP, SLP-convergence, Tolerances]	p. 202
XSLP_OPTIMALITYTOLTARGET, SLP OptimalityTolTarget	When set, this defines a target optimality tolerance to which the linearizations are solved to [Linearizations, SLP, Tolerances]	p. 203
XSLP_OTOL_A, SLPOTOL_A	Absolute static objective (2) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 204

XSLP_OTOL_R, SLPOTOL_R	Relative static objective (2) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 204
XSLP_PRESOLVE_ELIMTOL, NLPPRESOLVE_ELIMTOL	Tolerance for nonlinear eliminations during SLP presolve [Presolve, Tolerances]	p. 205
XSLP_PRESOLVEZERO, NLPPRESOLVEZERO	Minimum absolute value for a variable which is identified as nonzero during SLP presolve [Presolve, Tolerances]	p. 205
XSLP_STOL_A, SLPSTOL_A	Absolute slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_STOL_R, SLPSTOL_R	Relative slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_VALIDATIONFACTOR, NLPVALIDATIONFACTOR	Minimum improvement in validation targets to continue iterating [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_K, NLPVALIDATIONTARGET_K	Optimality target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_R, NLPVALIDATIONTARGET_R	Feasibility target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTOL_A, NLPVALIDATIONTOL_A	Absolute tolerance for the XSLPvalidate procedure [SLP, Tolerances]	p. 209
XSLP_VALIDATIONTOL_K, NLPVALIDATIONTOL_K	Relative tolerance for the XSLPvalidatekkt procedure [SLP, Tolerances]	p. 209
XSLP_VALIDATIONTOL_R, NLPVALIDATIONTOL_R	Relative tolerance for the XSLPvalidate procedure [Tolerances]	p. 209
XSLP_VTOL_A, SLPVTOL_A	Absolute static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 210
XSLP_VTOL_R, SLPVTOL_R	Relative static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WTOL_A, SLPWTOL_A	Absolute extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WTOL_R, SLPWTOL_R	Relative extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 212
XSLP_XTOL_A, SLPXTOL_A	Absolute static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 213
XSLP_XTOL_R, SLPXTOL_R	Relative static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 214
XSLP_ZERO, NLPZERO	Absolute tolerance [Tolerances]	p. 215

18.35 User Functions

Reference section for functions, controls, and attributes related to user functions (compare User Functions).

18.35.1 User Functions library functions

XSLPadduserfunction, XPRSnlpadduserfunction	Add user function definitions to an SLP problem.	
	[User Functions]	p. 283
XSLPdeluserfunction, XPRSnlpdeluserfunction	Delete a user function from the current problem	
	[User Functions]	p. 314
XSLPgetindex	Retrieve the index of an Xpress NonLinear entity with a given name	
	[Problem Information, User Functions]	p. 334
XSLPimportlibfunc, XPRSnlpimportlibfunc	Imports a function from a library file to be called as a user function	
	[User Functions]	p. 348

18.35.2 User Functions controls

XSLP_EVALUATE, NLPEVALUATE	Evaluation strategy for user functions	
	[User Functions]	p. 230
XSLP_FUNCEVAL, NLPFUNCEVAL	Bit map for determining the method of evaluating user functions and their derivatives	
	[Derivatives, User Functions]	p. 232
XSLP_THREADSAFEUSERFUNC, NLP_THREADSAFEUSERFUNC	Defines if user functions are allowed to be called in parallel	
	[Parallel, User Functions]	p. 253

18.35.3 User Functions attributes

XSLP_UFINSTANCES	Number of user function instances	
	[User Functions]	p. 158
XSLP_UFS, NLPUFS	Number of user functions	
	[User Functions]	p. 158
XSLP_USERFUNCCALLS, NLPUSERFUNCCALLS	Number of calls made to user functions	
	[User Functions]	p. 159

CHAPTER 19

Problem Attributes

During the optimization process, various properties of the problem being solved are stored and made available to users of the Xpress NonLinear Libraries in the form of *problem attributes*. These can be accessed in much the same manner as 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.

Library users are provided with the following functions for obtaining the values of attributes:

XSLPgetintattrib	XSLPgetdblattrib
XSLPgetptrattrib	XSLPgetstrattrib

The attributes listed in this chapter are all prefixed with `XSLP_`. Most of them also exist within the Optimizer library with an NLP or SLP prefix, e.g., `XPRS_NLPSOLSTATUS`. It is possible to use the above functions with other attributes for the Xpress Optimizer (attributes prefixed with `XPRS_`). For details of the Optimizer attributes, see the Optimizer manual.

Example of the usage of the functions:

```
XSLPgetintattrib(Prob, XSLP_ITER, &nIter);
printf("The number of SLP iterations is %d\n", nIter);
XSLPgetdblattrib(Prob, XSLP_ERRORCOSTS, &Errors);
printf("and the total error cost is %lg\n", Errors);
```

The following is a list of all the Xpress NonLinear attributes:

XSLP_COEFFICIENTS, SLPCOEFFICIENTS	Number of nonlinear coefficients [Problem Information, SLP]	p. 146
XSLP_CURRENTDELTACOST, SLPCURRENTDELTACOST	Current value of penalty cost multiplier for penalty delta vectors [SLP, Solution]	p. 142
XSLP_CURRENTERRORCOST, SLPCURRENTERRORCOST	Current value of penalty cost multiplier for penalty error vectors [SLP, Solution]	p. 142
XSLP_DELTAS, SLPDELTAS	Number of delta vectors created during augmentation [Problem Information, SLP]	p. 146
XSLP_ECFCOUNT, SLPECFCOUNT	Number of infeasible constraints found at the point of linearization [Linearizations, SLP]	p. 146
XSLP_EQUALSCOLUMN, NLPEQUALSCOLUMN	Index of the reserved "=" column [Problem Information]	p. 147

XSLP_ERRORCOSTS, SLPERRORCOSTS	Total penalty costs in the solution [SLP, Solution]	p. 142
XSLP_EXPLOREDELTAS, SLPEXPLOREDELTAS	Number of variables with an exploration-type delta set up in the problem [Problem Information, SLP]	p. 146
XSLP_IFS, NLPIFS	Number of internal functions [Misc]	p. 147
XSLP_IMPLICITVARIABLES, NLPIMPLICITVARIABLES	Number of SLP variables appearing only in coefficients [Problem Information, SLP]	p. 147
XSLP_INTEGERDELTAS, SLPINTEGERDELTAS	Number of variables set up with an integer delta in the problem [Problem Information, SLP]	p. 147
XSLP_ITER, SLPITER	SLP iteration count [Solution Process]	p. 148
XSLP_JOBID, NLPJOBID	Unique identifier for the current job [Misc, Multistart]	p. 148
XSLP_KEEPPBESTITER, NLPKEEPPBESTITER	The iteration in which the returned solution has been found. [SLP, Solution]	p. 148
XSLP_MINUSPENALTYERRORS, SLPMINUSPENALTYERRORS	Number of negative penalty error vectors [Problem Information, SLP]	p. 148
XSLP_MIPITER, SLPMIPITER	Total number of SLP iterations in MISLP [MISLP]	p. 149
XSLP_MIPNODES, SLPMIPNODES	Number of nodes explored in MISLP. [MISLP]	p. 149
XSLP_MIPPROBLEM	The underlying Optimizer MIP problem. [MISLP]	p. 161
XSLP_MIPSOLS, SLPMIPSOLS	Number of integer solutions found in MISLP. [MISLP, Solution]	p. 149
XSLP_MODELCOLS, NLPMODELCOLS	Number of model columns in the problem [SLP]	p. 149
XSLP_MODELROWS, NLPMODELROWS	Number of model rows in the problem [SLP]	p. 150
XSLP_MSSTATUS	Status of the multistart search [Multistart]	p. 150
XSLP_NLPSTATUS, NLPSTATUS	The solution status of the problem. [Solution Process]	p. 150
XSLP_NONCONSTANTCOEFFS, SLPNONCONSTANTCOEFFS	Number of coefficients in the augmented problem that might change between SLP iterations [SLP]	p. 151

XSLP_NONLINEARCONSTRAINTS, NONLINEARCONSTRAINTS	Number of nonlinear constraints in the problem [Problem Information]	p. 151
XSLP_OBJVAL, NLPOBJVAL	Objective function value excluding any penalty costs [Solution]	p. 143
XSLP_OPTTIME, NLPOPTTIME	Time spent in optimization [Solution Process]	p. 142
XSLP_ORIGINALCOLS, NLPORIGINALCOLS	Number of model columns in the extended original problem [Problem Information]	p. 151
XSLP_ORIGINALROWS, NLPORIGINALROWS	Number of model rows in the extended original problem [Problem Information]	p. 151
XSLP_PENALTYDELTACOLUMN, SLPPENALTYDELTACOLUMN	Index of column costing the penalty delta row [Problem Information, SLP]	p. 152
XSLP_PENALTYDELTAROW, SLPPENALTYDELTAROW	Index of equality row holding the penalties for delta vectors [Problem Information, SLP]	p. 152
XSLP_PENALTYDELTAS, SLPPENALTYDELTAS	Number of penalty delta vectors [Problem Information, SLP]	p. 152
XSLP_PENALTYDELTATOTAL, SLPPENALTYDELTATOTAL	Total activity of penalty delta vectors [SLP, Solution]	p. 143
XSLP_PENALTYDELTAVALUE, SLPPENALTYDELTAVALUE	Total penalty cost attributed to penalty delta vectors [SLP, Solution]	p. 143
XSLP_PENALTYERRORCOLUMN, SLPPENALTYERRORCOLUMN	Index of column costing the penalty error row [Problem Information, SLP]	p. 152
XSLP_PENALTYERRORROW, SLPPENALTYERRORROW	Index of equality row holding the penalties for penalty error vectors [Problem Information, SLP]	p. 153
XSLP_PENALTYERRORS, SLPPENALTYERRORS	Number of penalty error vectors [Problem Information, SLP]	p. 153
XSLP_PENALTYERRORTOTAL, SLPPENALTYERRORTOTAL	Total activity of penalty error vectors [SLP, Solution]	p. 143
XSLP_PENALTYERRORVALUE, SLPPENALTYERRORVALUE	Total penalty cost attributed to penalty error vectors [SLP, Solution]	p. 143
XSLP_PLUSPENALTYERRORS, SLPPUSPENALTYERRORS	Number of positive penalty error vectors [Problem Information, SLP]	p. 153
XSLP_PRESOLVEELIMINATIONS, NLPPRESOLVEELIMINATIONS	Number of SLP variables eliminated by XSLPresolve [Presolve]	p. 153

XSLP_PRESOLVSTATE	Indicates if the problem is presolved [Presolve]	p. 154
XSLP_PRIMALINTEGRAL, NLPPRIMALINTEGRAL	Local primal integral of the solve [Misc]	p. 144
XSLP_SBXCONVERGED, SLPSTBXCONVERGED	Number of step-bounded variables converged only on extended criteria [SLP, SLP-convergence]	p. 154
XSLP_SEMICONTDeltas, SLPSEMICONTDeltas	Number of variables with a minimum perturbation step set up in the problem [Problem Information, SLP]	p. 154
XSLP_SOLSTATUS, NLPSTATUS	Indicates the type of solution returned by the solver. [Solution]	p. 155
XSLP_SOLVERSELECTED, LOCALSOLVERSELECTED	Includes information of which Xpress solver has been used to solve the problem [Solution Process]	p. 154
XSLP_STATUS, SLPSTATUS	Bitmap holding the problem convergence status [Solution Process]	p. 155
XSLP_STOPSTATUS, NLPSTOPSTATUS	Status of the optimization process. [Solution Process]	p. 157
XSLP_TOTALEVALUATIONERRORS, NLP_TOTALEVALUATIONERRORS	The total number of evaluation errors during the solve [Numerics]	p. 157
XSLP_UCCONSTRAINEDCOUNT, SLPUCCONSTRAINEDCOUNT	Number of unconverged variables with coefficients in constraining rows [SLP, SLP-convergence]	p. 158
XSLP_UFINSTANCES	Number of user function instances [User Functions]	p. 158
XSLP_UFS, NLP_UFS	Number of user functions [User Functions]	p. 158
XSLP_UNCONVERGED, SLPUNCONVERGED	Number of unconverged values [SLP, SLP-convergence]	p. 158
XSLP_USEDERIVATIVES, NLPUSEDERIVATIVES	Indicates whether numeric or analytic derivatives were used to create the linear approximations and solve the problem [Derivatives]	p. 159
XSLP_USERFUNCCALLS, NLPUSERFUNCCALLS	Number of calls made to user functions [User Functions]	p. 159
XSLP_VALIDATIONINDEX_A, NLPVALIDATIONINDEX_A	Absolute validation index [Solution]	p. 144
XSLP_VALIDATIONINDEX_K, NLPVALIDATIONINDEX_K	Relative first order optimality validation index [Solution]	p. 144
XSLP_VALIDATIONINDEX_R, NLPVALIDATIONINDEX_R	Relative validation index [Solution]	p. 144

XSLP_VALIDATIONNETOBJ, NLPVALIDATIONNETOBJ	Net objective as calculated by validation	
[Solution]		p. 144
XSLP_VALIDATIONSTATUS, NLPVALIDATIONSTATUS	Feasibility status of the current solution.	
[Solution]		p. 159
XSLP_VARIABLES, NLPVARIABLES	Number of SLP variables	
[Data Information, SLP]		p. 159
XSLP_VERSIONDATE	Date of creation of Xpress NonLinear	
[Misc]		p. 162
XSLP_VSOLINDEX	Vertex solution index	
[SLP, Solution]		p. 145
XSLP_XPRSPROBLEM	The underlying Optimizer problem	
[Misc]		p. 161
XSLP_XSLPPROBLEM	The Xpress NonLinear problem	
[Misc]		p. 161
XSLP_ZEROESRESET, SLPZEROESRESET	Number of placeholder entries set to zero	
[SLP]		p. 160
XSLP_ZEROESRETAINED, SLPZEROESRETAINED	Number of potentially zero placeholders left untouched	
[SLP]		p. 160
XSLP_ZEROESTOTAL, SLPZEROESTOTAL	Number of potential zero placeholder entries	
[SLP]		p. 160

19.1 Double problem attributes

XSLP_CURRENTDELTACOST, SLPCURRENTDELTACOST

Description	Current value of penalty cost multiplier for penalty delta vectors
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Set by routines	XSLPmaxim, XSLPminim
See also	XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_CURRENTERRORCOST

XSLP_CURRENTERRORCOST, SLPCURRENTERRORCOST

Description	Current value of penalty cost multiplier for penalty error vectors
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Set by routines	XSLPmaxim, XSLPminim
See also	XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_CURRENTDELTACOST

XSLP_ERRORCOSTS, SLPERRORCOSTS

Description	Total penalty costs in the solution
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Set by routines	XSLPmaxim, XSLPminim

XSLP_OPTTIME, NLPOPTTIME

Description	Time spent in optimization
Type	Double
Category	Attribute
Topic area	Solution Process
Set by routines	XSLPmaxim, XSLPminim

XSLP_OBJVAL, NLPOBJVAL

Description	Objective function value excluding any penalty costs
Type	Double
Category	Attribute
Topic area	Solution
Set by routines	XSLPmaxim, XSLPminim

XSLP_PENALTYDELTATOTAL, SLPPENALTYDELTATOTAL

Description	Total activity of penalty delta vectors
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Set by routines	XSLPmaxim, XSLPminim

XSLP_PENALTYDELTAVALUE, SLPPENALTYDELTAVALUE

Description	Total penalty cost attributed to penalty delta vectors
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Set by routines	XSLPmaxim, XSLPminim

XSLP_PENALTYERRORTOTAL, SLPPENALTYERRORTOTAL

Description	Total activity of penalty error vectors
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Set by routines	XSLPmaxim, XSLPminim

XSLP_PENALTYERRORVALUE, SLPPENALTYERRORVALUE

Description	Total penalty cost attributed to penalty error vectors
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Set by routines	XSLPmaxim, XSLPminim

XSLP_PRIMALINTEGRAL, NLP_PRIMALINTEGRAL

Description	Local primal integral of the solve
Type	Double
Category	Attribute
Topic area	Misc
Set by routines	XSLPmaxim, XSLPminim

XSLP_VALIDATIONINDEX_A, NLP_VALIDATIONINDEX_A

Description	Absolute validation index
Type	Double
Category	Attribute
Topic area	Solution
Set by routines	XSLPvalidate

XSLP_VALIDATIONINDEX_K, NLP_VALIDATIONINDEX_K

Description	Relative first order optimality validation index
Type	Double
Category	Attribute
Topic area	Solution
Set by routines	XSLPvalidatekkt

XSLP_VALIDATIONINDEX_R, NLP_VALIDATIONINDEX_R

Description	Relative validation index
Type	Double
Category	Attribute
Topic area	Solution
Set by routines	XSLPvalidate

XSLP_VALIDATIONNETOBJ, NLP_VALIDATIONNETOBJ

Description	Net objective as calculated by validation
Type	Double
Category	Attribute
Topic area	Solution
Set by routines	XSLPvalidate

XSLP_VSOLINDEX

Description	Vertex solution index
Type	Double
Category	Attribute
Topic areas	SLP, Solution
Notes	<p>The <i>vertex solution index</i> (VSOLINDEX) is a measure of how nearly the converged solution to a problem is at a vertex (that is, at the intersection of a set of constraints) of the feasible region.</p> <p>Where the solution is in the middle of a face, the solution will in general have been achieved through the use of step bounds. The VSOLINDEX is the fraction of delta vectors which are <i>not</i> at a bound in the solution. Therefore, a value of 1.0 means that no delta is at a step bound and therefore the solution is at a vertex of the feasible region. Smaller values indicate that there are deltas at step bounds and so the solution is further from being a vertex solution.</p>

19.2 Integer problem attributes

XSLP_COEFFICIENTS, SLP_COEFFICIENTS

Description	Number of nonlinear coefficients
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Note	XSLP_COEFFICIENTS includes both coefficients (nonlinear expressions multiplying a variable) and formulas (nonlinear expressions which do not multiply a variable). See 6.1 for more information.
Set by routines	XSLPaddcoefs, XSLPchgcoef, XSLPloadcoefs, XSLPreadprob

XSLP_DELTAS, SLPDELTAS

Description	Number of delta vectors created during augmentation
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	XSLPconstruct

XSLP_ECFCOUNT, SLPECFCOUNT

Description	Number of infeasible constraints found at the point of linearization
Type	Integer
Category	Attribute
Topic areas	SLP, Linearizations
Set by routines	XSLPmaxim, XSLPminim
See also	XSLP_ECFCHECK, XSLP_ECFTOL_A, XSLP_ECFTOL_R

XSLP_EXPLOREDELTAS, SLPEXPLOREDELTAS

Description	Number of variables with an exploration-type delta set up in the problem
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	XSLPconstruct

XSLP_EQUALSCOLUMN, NLPEQUALSCOLUMN

Description	Index of the reserved "=" column
Type	Integer
Category	Attribute
Topic area	Problem Information
Note	If there had been no "=" column present, it will be assumed that the user needs the index to add nonlinear terms into the problem that are not coefficients, and an "=" columns will be added to the problem, whose index is then returned. Please note, that this means that a call to XSLPgetintattrib with this attribute might make a slight modification to the problem itself.
Set by routines	XSLPconstruct, XSLPreadprob

XSLP_IFS, NLPIFS

Description	Number of internal functions
Type	Integer
Category	Attribute
Topic area	Misc
Set by routines	XSLPcreateprob

XSLP_IMPLICITVARIABLES, NLPIMPLICITVARIABLES

Description	Number of SLP variables appearing only in coefficients
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	XSLPconstruct

XSLP_INTEGERDELTAS, SLPINTEGERDELTAS

Description	Number of variables set up with an integer delta in the problem
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	XSLPconstruct

XSLP_ITER, SLPITER

Description	SLP iteration count
Type	Integer
Category	Attribute
Topic area	Solution Process
Set by routines	XSLPmaxim, XSLPminim

XSLP_JOBID, NLPJOBID

Description	Unique identifier for the current job
Type	Integer
Category	Attribute
Topic areas	Misc, Multistart
Note	Assigned when a job is created, and can be used to identify jobs in callbacks. Note that all callback receives an optional job name that can be assigned at job creation time.
Set by routines	XSLPmaxim, XSLPminim

XSLP_KEEPESTITER, NLPKEEPESTITER

Description	The iteration in which the returned solution has been found.
Type	Integer
Category	Attribute
Topic areas	SLP, Solution
Note	A zero value indicates no solution or the filter option is off. A value of '-1' indicates the initial solution has been returned.
Set by routines	XSLPmaxim, XSLPminim

XSLP_MINUSPENALTYERRORS, SLPMINUSPENALTYERRORS

Description	Number of negative penalty error vectors
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	XSLPconstruct

XSLP_MIPITER, SLPMIPITER

Description	Total number of SLP iterations in MISLP
Type	Integer
Category	Attribute
Topic area	MISLP
Set by routines	XSLPnlpoptimize, XSLPmaxim, XSLPminim.

XSLP_MIPNODES, SLPMIPNODES

Description	Number of nodes explored in MISLP. This includes any nodes for which a non-linear solve has been carried out.
Type	Integer
Category	Attribute
Topic area	MISLP
Set by routines	XSLPnlpoptimize, XSLPmaxim, XSLPminim.

XSLP_MIPSOLS, SLPMIPSOLS

Description	Number of integer solutions found in MISLP. This includes solutions found during the tree search or any heuristics.
Type	Integer
Category	Attribute
Topic areas	MISLP, Solution
Set by routines	XSLPnlpoptimize, XSLPmaxim, XSLPminim.

XSLP_MODELCOLS, NLPMODELCOLS

Description	Number of model columns in the problem
Type	Integer
Category	Attribute
Topic area	SLP
Note	This is the number of columns currently in the problem without any augmentation, i.e. the number of columns that describe the algebraic definition of the problem. These columns always precede the augmentation columns in order. If the problem is presolved, this may be smaller than the number of original columns in the problem. To access the number of original columns, use XPRS_INPUTCOLS.
See also	XSLP_MODELROWS, XPRS_INPUTROWS, XPRS_INPUTCOLS.

XSLP_MODELROWS, NLPMODELROWS

Description	Number of model rows in the problem
Type	Integer
Category	Attribute
Topic area	SLP
Note	This is the number of rows currently in the problem without any augmentation, i.e. the number of rows that describe the algebraic definition of the problem. These rows always precede the augmentation rows in order. If the problem is presolved, this may be smaller than the number of original rows in the problem. To access the number of original rows, use XPRS_INPUTROWS.
See also	XSLP_MODELCOLS, XPRS_INPUTROWS, XPRS_INPUTCOLS.

XSLP_MSSTATUS

Description	Status of the multistart search
Type	Integer
Category	Attribute
Topic area	Multistart
Note	The value matches that of the winner job if the multistart search completes and a feasible solution has been found. If no solution is found, it is set to XSLP_NLPSTATUS_INFEASIBLE. If the search is terminated early, it is set to XSLP_NLPSTATUS_UNFINISHED (though in which case the winner if any is still synchronized to the base problem and the solution and XSLP_NLPSTATUS is available).

XSLP_NLPSTATUS, NLPSTATUS

Description	The solution status of the problem.																
Type	Integer																
Category	Attribute																
Topic area	Solution Process																
Values	<table> <tr> <td>0</td><td>Optimization unstarted (XSLP_NLPSTATUS_UNSTARTED)</td></tr> <tr> <td>1</td><td>Solution found (XSLP_NLPSTATUS_SOLUTION)</td></tr> <tr> <td>2</td><td>Globally optimal (XSLP_NLPSTATUS_OPTIMAL)</td></tr> <tr> <td>3</td><td>No solution found (XSLP_NLPSTATUS_NOSOLUTION)</td></tr> <tr> <td>4</td><td>Proven infeasible (XSLP_NLPSTATUS_INFEASIBLE)</td></tr> <tr> <td>5</td><td>Locally unbounded (XSLP_NLPSTATUS_UNBOUNDED)</td></tr> <tr> <td>6</td><td>Not yet solved to completion (XSLP_NLPSTATUS_UNFINISHED)</td></tr> <tr> <td>7</td><td>Could not be solved due to numerical issues (XSLP_NLPSTATUS_UNUNSOLVED)</td></tr> </table>	0	Optimization unstarted (XSLP_NLPSTATUS_UNSTARTED)	1	Solution found (XSLP_NLPSTATUS_SOLUTION)	2	Globally optimal (XSLP_NLPSTATUS_OPTIMAL)	3	No solution found (XSLP_NLPSTATUS_NOSOLUTION)	4	Proven infeasible (XSLP_NLPSTATUS_INFEASIBLE)	5	Locally unbounded (XSLP_NLPSTATUS_UNBOUNDED)	6	Not yet solved to completion (XSLP_NLPSTATUS_UNFINISHED)	7	Could not be solved due to numerical issues (XSLP_NLPSTATUS_UNUNSOLVED)
0	Optimization unstarted (XSLP_NLPSTATUS_UNSTARTED)																
1	Solution found (XSLP_NLPSTATUS_SOLUTION)																
2	Globally optimal (XSLP_NLPSTATUS_OPTIMAL)																
3	No solution found (XSLP_NLPSTATUS_NOSOLUTION)																
4	Proven infeasible (XSLP_NLPSTATUS_INFEASIBLE)																
5	Locally unbounded (XSLP_NLPSTATUS_UNBOUNDED)																
6	Not yet solved to completion (XSLP_NLPSTATUS_UNFINISHED)																
7	Could not be solved due to numerical issues (XSLP_NLPSTATUS_UNUNSOLVED)																
Note	XSLP_NLPSTATUS_OPTIMAL is only set when a globally optimal solution has been found, e.g., because the problem was convex, or because the global solver was used.																

Note	XSLP_NLPSTATUS_SOLUTION is set whenever a feasible solution was found. The solve may have been interrupted, or it completed without proving global optimality.
Note	XSLP_NLPSTATUS_UNFINISHED indicates that the solve was interrupted without finding a feasible solution.
Set by routines	XSLPnlpoptimize, XSLPmaxim, XSLPminim.

XSLP_NONCONSTANTCOEFFS, SLPNONCONSTANTCOEFFS

Description	Number of coefficients in the augmented problem that might change between SLP iterations
Type	Integer
Category	Attribute
Topic area	SLP
Set by routines	XSLPconstruct

XSLP_NONLINEARCONSTRAINTS, NONLINEARCONSTRAINTS

Description	Number of nonlinear constraints in the problem
Type	Integer
Category	Attribute
Topic area	Problem Information
Set by routines	XSLPconstruct

XSLP_ORIGINALCOLS, NLPORIGINALCOLS

Description	Number of model columns in the extended original problem
Type	Integer
Category	Attribute
Topic area	Problem Information
Note	This includes columns introduced by transformations and reformulations carried out during nonlinear presolve, but does not include augmentation columns. To access the number of original columns, use XPRS_INPUTCOLS.
See also	XSLP_ORIGINALROWS, XSLP_MODELROWS, XSLP_MODELCOLS, XPRS_INPUTROWS, XPRS_INPUTCOLS.

XSLP_ORIGINALROWS, NLPORIGINALROWS

Description	Number of model rows in the extended original problem
Type	Integer
Category	Attribute
Topic area	Problem Information

Note	This includes rows introduced by transformations and reformulations carried out during nonlinear presolve, but does not include augmentation rows. To access the number of original rows, use <code>XPRS_INPUTROWS</code> .
See also	<code>XSLP_ORIGINALCOLS</code> , <code>XSLP_MODELROWS</code> , <code>XSLP_MODELCOLS</code> , <code>XPRS_INPUTROWS</code> , <code>XPRS_INPUTCOLS</code> .

XSLP_PENALTYDELTACOLUMN, SLPPENALTYDELTACOLUMN

Description	Index of column costing the penalty delta row
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Note	This index always counts from 1. It is zero if there is no penalty delta row.
Set by routines	<code>XSLPconstruct</code>

XSLP_PENALTYDELTAROW, SLPPENALTYDELTAROW

Description	Index of equality row holding the penalties for delta vectors
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Note	This index always counts from 1. It is zero if there are no penalty delta vectors.
Set by routines	<code>XSLPconstruct</code>

XSLP_PENALTYDELTAS, SLPPENALTYDELTAS

Description	Number of penalty delta vectors
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	<code>XSLPconstruct</code>

XSLP_PENALTYERRORCOLUMN, SLPPENALTYERRORCOLUMN

Description	Index of column costing the penalty error row
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Note	This index always counts from 1. It is zero if there is no penalty error row.
Set by routines	<code>XSLPconstruct</code>

XSLP_PENALTYERRORROW, SLPPENALTYERRORROW

Description	Index of equality row holding the penalties for penalty error vectors
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Note	This index always counts from 1. It is zero if there are no penalty error vectors.
Set by routines	XSLPconstruct

XSLP_PENALTYERRORS, SLPPENALTYERRORS

Description	Number of penalty error vectors
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	XSLPconstruct

XSLP_PLUSPENALTYERRORS, SLPPLUSPENALTYERRORS

Description	Number of positive penalty error vectors
Type	Integer
Category	Attribute
Topic areas	SLP, Problem Information
Set by routines	XSLPconstruct

XSLP_PRESOLVEELIMINATIONS, NLPPRESOLVEELIMINATIONS

Description	Number of SLP variables eliminated by XSLPpresolve
Type	Integer
Category	Attribute
Topic area	Presolve
Set by routines	XSLPpresolve

XSLP_PRESOLVSTATE

Description	Indicates if the problem is presolved	
Type	Integer	
Category	Attribute	
Topic area	Presolve	
Values	0	The problem is not presolved
	1	The problem is presolved, but no columns or rows have been removed from the problem
	2	The problem is fully presolved, and the column and row indices do not match the original problem
Set by routines	XSLPmaxim, XSLPminim, XSLPpresolve	

XSLP_SBXCONVERGED, SLPSBXCONVERGED

Description	Number of step-bounded variables converged only on extended criteria	
Type	Integer	
Category	Attribute	
Topic areas	SLP, SLP-convergence	
Set by routines	XSLPmaxim, XSLPminim	

XSLP_SEMICONTDeltas, SLPSEMICONTDeltas

Description	Number of variables with a minimum perturbation step set up in the problem	
Type	Integer	
Category	Attribute	
Topic areas	SLP, Problem Information	
Set by routines	XSLPconstruct	

XSLP_SOLVERSELECTED, LOCALSOLVERSELECTED

Description	Includes information of which Xpress solver has been used to solve the problem	
Type	Integer	
Category	Attribute	
Topic area	Solution Process	
Values	-1	Unset
	0	Xpress-SLP
	1	Knitro (Artelys)
	2	Xpress Optimizer

Set by routines XSLPmaxim, XSLPminim

Note The following constants are provided:

- 0 XSLP_SOLVER_XSLP
- 1 XSLP_SOLVER_KNITRO
- 2 XSLP_SOLVER_OPTIMIZER

XSLP_SOLSTATUS, NLPSOLSTATUS

Description Indicates the type of solution returned by the solver.

Type Integer

Category Attribute

Topic area Solution

Values

- 0 No solution available.
- 1 A solution with no dual information.
- 2 A locally optimal solution with dual information.
- 3 A globally optimal solution without dual information.
- 4 A globally optimal solution with dual information.

Set by routines XSLPnlpoptimize, XSLPmaxim, XSLPminim.

Note The following constants are provided:

- 0 XSLP_SOLSTATUS_NONE
- 1 XSLP_SOLSTATUS_SOLUTION_NODUALS
- 2 XSLP_SOLSTATUS_LOCALLYOPTIMAL_WITHDUALS
- 3 XSLP_SOLSTATUS_GLOBALLYOPTIMAL_NODUALS
- 4 XSLP_SOLSTATUS_GLOBALLYOPTIMAL_WITHDUALS

XSLP_STATUS, SLPSTATUS

Description Bitmap holding the problem convergence status

Type Integer

Category Attribute

Topic area Solution Process

Values

Bit	Meaning
0	Converged on objective function with no unconverged values in active constraints.
1	Converged on objective function with some variables converged on extended criteria only.
2	LP solution is infeasible.
3	LP solution is unfinished (not optimal or infeasible).
4	SLP terminated on maximum SLP iterations.
5	SLP is integer infeasible.
6	SLP converged with residual penalty errors.
7	Converged on objective.
9	SLP terminated on max time.
10	SLP terminated by user.
11	Some variables are linked to active constraints.
12	No unconverged values in active constraints.
13	OTOL is satisfied - range of objective change small, active step bounds.
14	VTOL is satisfied - range of objective change is small.
15	XTOL is satisfied - range of objective change small, no unconverged in active.
16	WTOL is satisfied - convergence continuation.
17	ERRORTOL satisfied - penalties not increased further.
18	EVTOL satisfied - penalties not increased further.
19	There were iterations where the solution had to be polished.
20	There were iterations where the solution polishing failed.
21	There were iterations where rows were enforced.
22	Terminated due to XSLP_INFEASLIMIT.

Note

A value of zero after SLP optimization means that the solution is fully converged.

The following constants are provided for checking these bits:

Setting bit 0	XSLP_STATUS_CONVERGEDOBJUCC
Setting bit 1	XSLP_STATUS_CONVERGEDOBJSBX
Setting bit 2	XSLP_STATUS_LPINFEASIBLE
Setting bit 3	XSLP_STATUS_LPUNFINISHED
Setting bit 4	XSLP_STATUS_MAXSLPITERATIONS
Setting bit 5	XSLP_STATUS_INTEGERINFEASIBLE
Setting bit 6	XSLP_STATUS_RESIDUALPENALTIES
Setting bit 7	XSLP_STATUS_CONVERGEDOBJOBJ
Setting bit 9	XSLP_STATUS_MAXTIME
Setting bit 10	XSLP_STATUS_USER
Setting bit 11	XSLP_STATUS_VARSLINKEDINACTIVE
Setting bit 12	XSLP_STATUS_NOVARSINACTIVE
Setting bit 13	XSLP_STATUS_OTOL
Setting bit 14	XSLP_STATUS_VTOL
Setting bit 15	XSLP_STATUS_XTOL
Setting bit 16	XSLP_STATUS_WTOL
Setting bit 17	XSLP_STATUS_ERROTOL
Setting bit 18	XSLP_STATUS_EVTOL
Setting bit 19	XSLP_STATUS_POLISHED
Setting bit 20	XSLP_STATUS_POLISH_FAILURE
Setting bit 21	XSLP_STATUS_ENFORCED
Setting bit 22	XSLP_STATUS_CONSECUTIVE_INFEAS

Set by routines XSLPmaxim, XSLPminim

XSLP_STOPSTATUS, NLPSTOPSTATUS

Description Status of the optimization process.

Type Integer

Category Attribute

Topic area Solution Process

Note Possible values are:

Value	Description
XSLP_STOP_NONE	no interruption – the solve completed normally
XSLP_STOP_TIMELIMIT	time limit hit
XSLP_STOP_CTRLC	control C hit
XSLP_STOP_NODELIMIT	node limit hit
XSLP_STOP_ITERLIMIT	iteration limit hit
XSLP_STOP_MIPGAP	MIP gap is sufficiently small
XSLP_STOP_SOLLIMIT	solution limit hit
XSLP_STOP_USER	user interrupt.

Set by routines XSLPnlpoptimize, XSLPmaxim, XSLPminim.

XSLP_TOLSETS, SLPTOLSETS

Description *This parameter is deprecated and will be removed in a future release.*
Number of tolerance sets.

Type Integer

Category Attribute

Topic area SLP-convergence

Set by routines XSLPaddtolsets, XSLPchgtolset, XSLPloadtolsets, XSLPreadprob

XSLP_TOTALEVALUATIONERRORS, NLPOTALEVALUATIONERRORS

Description The total number of evaluation errors during the solve

Type Integer

Category Attribute

Topic area Numerics

Set by routines XSLPnlpoptimize, XSLPmaxim, XSLPminim.

XSLP_UCONSTRAINEDCOUNT, SLPUCCONSTRAINEDCOUNT

Description	Number of unconverged variables with coefficients in constraining rows
Type	Integer
Category	Attribute
Topic areas	SLP, SLP-convergence
Set by routines	XSLPmaxim, XSLPminim

XSLP_UFINSTANCES

Description	Number of user function instances
Type	Integer
Category	Attribute
Topic area	User Functions
Set by routines	XSLPconstruct

XSLP_UFS, NLPUFS

Description	Number of user functions
Type	Integer
Category	Attribute
Topic area	User Functions
Set by routines	XSLPadduserfunction, XSLPdeluserfunction, XSLPreadprob

XSLP_UNCONVERGED, SLPUNCONVERGED

Description	Number of unconverged values
Type	Integer
Category	Attribute
Topic areas	SLP, SLP-convergence
Note	Prior to the first iteration this will return -1.
Set by routines	XSLPmaxim, XSLPminim

XSLP_USEDERIVATIVES, NLPUSEDERIVATIVES

Description	Indicates whether numeric or analytic derivatives were used to create the linear approximations and solve the problem	
Type	Integer	
Category	Attribute	
Topic area	Derivatives	
Values	0	numeric derivatives.
	1	analytic derivatives for all formulae unless otherwise specified.
Set by routines	XSLPconstruct	

XSLP_USERFUNCCALLS, NLPUSERFUNCCALLS

Description	Number of calls made to user functions	
Type	Integer	
Category	Attribute	
Topic area	User Functions	
Set by routines	XSLPcascade, XSLPconstruct, XSLPevaluatecoef, XSLPevaluateformula, XSLPmaxim, XSLPminim	

XSLP_VALIDATIONSTATUS, NLPVALIDATIONSTATUS

Description	Feasibility status of the current solution.	
Type	Integer	
Category	Attribute	
Topic area	Solution	
Values	1	Feasible.
	4	Infeasible.
Set by routines	XSLPvalidate.	
Note	The following constants are provided:	
	1	XSLP_NLPSTATUS_SOLUTION
	4	XSLP_NLPSTATUS_INFEASIBLE

XSLP_VARIABLES, NLPVARIABLES

Description	Number of SLP variables	
Type	Integer	
Category	Attribute	
Topic areas	SLP, Data Information	
Set by routines	XSLPconstruct	

XSLP_ZEROESRESET, SLPZEROESRESET

Description	Number of placeholder entries set to zero
Type	Integer
Category	Attribute
Topic area	SLP
Note	For an explanation of deletion of placeholder entries in the matrix see <i>Management of zero placeholder entries</i> .
Set by routines	XSLPmaxim, XSLPminim
See also	XSLP_ZEROCRITERIONCOUNT, XSLP_ZEROCRITERIONSTART, XSLP_ZEROESRETAINED, XSLP_ZEROESTOTAL, <i>Management of zero placeholder entries</i>

XSLP_ZEROESRETAINED, SLPZEROESRETAINED

Description	Number of potentially zero placeholders left untouched
Type	Integer
Category	Attribute
Topic area	SLP
Note	For an explanation of deletion of placeholder entries in the matrix see <i>Management of zero placeholder entries</i> .
Set by routines	XSLPmaxim, XSLPminim
See also	XSLP_ZEROCRITERIONCOUNT, XSLP_ZEROCRITERIONSTART, XSLP_ZEROESRESET, XSLP_ZEROESTOTAL, <i>Management of zero placeholder entries</i>

XSLP_ZEROESTOTAL, SLPZEROESTOTAL

Description	Number of potential zero placeholder entries
Type	Integer
Category	Attribute
Topic area	SLP
Note	For an explanation of deletion of placeholder entries in the matrix see <i>Management of zero placeholder entries</i> .
Set by routines	XSLPmaxim, XSLPminim
See also	XSLP_ZEROCRITERIONCOUNT, XSLP_ZEROCRITERIONSTART, XSLP_ZEROESRESET, XSLP_ZEROESRETAINED, <i>Management of zero placeholder entries</i>

19.3 Reference (pointer) problem attributes

The reference attributes are void pointers whose size (32 or 64 bit) depends on the platform.

XSLP_MIPPROBLEM

Description	The underlying Optimizer MIP problem. <code>XSLP_MIPPROBLEM</code> is a reference of type <code>XPRSPprob</code> , and should be used in MISLP callbacks to access MIP-specific Optimizer values (such as node and parent numbers).
Type	Reference
Category	Attribute
Topic area	MISLP
Set by routines	<code>XSLPnlpoptimize</code>

XSLP_XPRSPROBLEM

Description	The underlying Optimizer problem
Type	Reference
Category	Attribute
Topic area	Misc
Set by routines	<code>XSLPcreateprob</code>

XSLP_XSLPPROBLEM

Description	The Xpress NonLinear problem
Type	Reference
Category	Attribute
Topic area	Misc
Set by routines	<code>XSLPcreateprob</code>

19.4 String problem attributes

XSLP_VERSIONDATE

Description	Date of creation of Xpress NonLinear
Type	String
Category	Attribute
Topic area	Misc
Note	The format of the date is dd mmm yyyy.
Set by routines	XSLPinit

CHAPTER 20

Control Parameters

Various controls exist within Xpress NonLinear to govern the solution procedure and the form of the output. Some of these take integer values and act as switches between various types of behavior. Many are tolerances on values related to the convergence criteria; these are all double precision. There are also a few controls which are character strings, setting names for structures. Any of these may be altered by the user to enhance performance of the SLP algorithm. In most cases, 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.

Users of the Xpress NonLinear function library are provided with the following set of functions for setting and obtaining control values:

XSLPgetintcontrol	XSLPgetdblcontrol	XSLPgetstrcontrol
XSLPsetintcontrol	XSLPsetdblcontrol	XSLPsetstrcontrol

All the controls as listed in this chapter are prefixed with XSLP_. Most of them also exist within the Optimizer library with an NLP or SLP prefix, e.g., XPRS_NLPPRESOLVE. It is possible to use the above functions with other control parameters for the Xpress Optimizer (controls prefixed with XPRS_). For details of the Optimizer controls, see the Optimizer manual.

Example of the usage of the functions:

```
XSLPgetintcontrol(Prob, XSLP_PRESOLVE, &presolve);
printf("The value of PRESOLVE was %d\n", presolve);
XSLPsetintcontrol(Prob, XSLP_PRESOLVE, 1-presolve);
printf("The value of PRESOLVE is now %d\n", 1-presolve);
```

The following is a list of all the Xpress NonLinear controls:

XSLP_ALGORITHM, SLPALGORITHM	Bit map describing the SLP algorithm(s) to be used [Bit-vector, SLP]	p. 216
XSLP_ANALYZE, SLPANALYZE	Bit map activating additional options supporting model / solution path analysis [Bit-vector, Logging, SLP]	p. 218
XSLP_ATOL_A, SLPATOL_A	Absolute delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
XSLP_ATOL_R, SLPATOL_R	Relative delta convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 176
XSLP_AUGMENTATION, SLPAUGMENTATION	Bit map describing the SLP augmentation method(s) to be used [Bit-vector, SLP]	p. 219

XSLP_AUTOSAVE, SLPAUTOSAVE	Frequency with which to save the model [Logging, SLP]	p. 220
XSLP_BARCROSSOVERSTART, SLPBARCROSSOVERSTART	Default crossover activation behaviour for barrier start [Linearizations, SLP]	p. 221
XSLP_BARLIMIT, SLPBARLIMIT	Number of initial SLP iterations using the barrier method [Limits, Linearizations, SLP]	p. 221
XSLP_BARSTALLINGLIMIT, SLPBARSTALLINGLIMIT	Number of iterations to allow numerical failures in barrier before switching to dual [Limits, Linearizations, SLP]	p. 222
XSLP_BARSTALLINGOBJLIMIT, SLPBARSTALLINGOBJLIMIT	Number of iterations over which to measure the objective change for barrier iterations with no crossover [Limits, Linearizations, SLP]	p. 222
XSLP_BARSTALLINGTOL, SLPBARSTALLINGTOL	Required change in the objective when progress is measured in barrier iterations without crossover [Linearizations, SLP]	p. 176
XSLP_BARSTARTOPS, SLPBARSTARTOPS	Controls behaviour when the barrier is used to solve the linearizations [Linearizations, SLP]	p. 222
XSLP_BOUNDTHRESHOLD, SLPBOUNDTHRESHOLD	The maximum size of a bound that can be introduced by nonlinear presolve. [Presolve, SLP]	p. 177
XSLP_CALCTHREADS, NLPCALCTHREADS	Number of threads used for formula and derivatives evaluations [Parallel]	p. 223
XSLP_CASCADE, SLPCASCADE	Bit map describing the cascading to be used [Cascading, SLP]	p. 223
XSLP_CASCADENLIMIT, SLPASCADENLIMIT	Maximum number of iterations for cascading with non-linear determining rows [Cascading, Limits, SLP]	p. 224
XSLP_CASCADETOL_PA, SLPCASCADETOL_PA	Absolute cascading print tolerance [Cascading, Logging, SLP]	p. 177
XSLP_CASCADETOL_PR, SLPCASCADETOL_PR	Relative cascading print tolerance [Cascading, Logging, SLP]	p. 177
XSLP_CDTOL_A, SLPCTOL_A	Absolute tolerance for deducing constant derivatives [Tolerances]	p. 178
XSLP_CDTOL_R, SLPCTOL_R	Relative tolerance for deducing constant derivatives [Tolerances]	p. 178
XSLP_CLAMP SHRINK, SLPCLAMP SHRINK	Shrink ratio used to impose strict convergence on variables converged in extended criteria only [SLP]	p. 179
XSLP_CLAMPVALIDATIONTOL_A, SLPCLAMPVALIDATIONTOL_A	Absolute validation tolerance for applying XSLP_CLAMP SHRINK [SLP, Tolerances]	p. 179

XSLP_CLAMPVALIDATIONTOL_R, SLP_CLAMPVALIDATIONTOL_R	Relative validation tolerance for applying XSLP_CLAMP_SHRINK [SLP, Tolerances]	p. 179
XSLP_CONTROL	Bit map describing which Xpress NonLinear functions also activate the corresponding Optimizer Library function [Bit-vector]	p. 225
XSLP_CONVERGENCEOPS, SLP_CONVERGENCEOPS	Bit map describing which convergence tests should be carried out [Bit-vector, SLP, SLP-convergence]	p. 225
XSLP_CTOL, SLP_CTOL	Closure convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 180
XSLP_CUTSTRATEGY, SLP_CUTSTRATEGY	Determines which cuts to apply in the MISLP search when the default SLP-in-MIP strategy is used. [Cuts, MISLP]	p. 226
XSLP_DAMP, SLP_DAMP	Damping factor for updating values of variables [SLP]	p. 180
XSLP_DAMP_EXPAND, SLP_DAMP_EXPAND	Multiplier to increase damping factor during dynamic damping [SLP]	p. 180
XSLP_DAMP_MAX, SLP_DAMP_MAX	Maximum value for the damping factor of a variable during dynamic damping [SLP]	p. 181
XSLP_DAMP_MIN, SLP_DAMP_MIN	Minimum value for the damping factor of a variable during dynamic damping [SLP]	p. 181
XSLP_DAMP_SHRINK, SLP_DAMP_SHRINK	Multiplier to decrease damping factor during dynamic damping [SLP]	p. 182
XSLP_DAMP_START, SLP_DAMP_START	SLP iteration at which damping is activated [SLP]	p. 227
XSLP_DEFAULT_IV, SLP_DEFAULT_IV	Default initial value for an SLP variable if none is explicitly given [Data Input]	p. 182
XSLP_DEFAULT_STEPBOUND, SLP_DEFAULT_STEPBOUND	Minimum initial value for the step bound of an SLP variable if none is explicitly given [SLP]	p. 182
XSLP_DELAY_UPDATEROWS, SLP_DELAY_UPDATEROWS	Number of SLP iterations before update rows are fully activated [SLP]	p. 227
XSLP_DELTA_A, SLP_DELTA_A	Absolute perturbation of values for calculating numerical derivatives [Derivatives]	p. 183
XSLP_DELTA_INFINITY, SLP_DELTA_INFINITY	Maximum value for partial derivatives [Derivatives]	p. 183
XSLP_DELTA_R, SLP_DELTA_R	Relative perturbation of values for calculating numerical derivatives [Derivatives]	p. 183

XSLP_DELTA_X, SLPDELTA_X	Minimum absolute value of delta coefficients to be retained [Derivatives]	p. 184
XSLP_DELTA_Z, SLPDELTA_Z	Tolerance used when calculating derivatives [Derivatives, Tolerances]	p. 184
XSLP_DELTA_ZERO, SLPDELTA_ZERO	Absolute zero acceptance tolerance used when calculating derivatives [Derivatives, Tolerances]	p. 184
XSLP_DELTACOST, SLPDELTACOST	Initial penalty cost multiplier for penalty delta vectors [SLP]	p. 185
XSLP_DELTACOSTFACTOR, SLPDELTACOSTFACTOR	Factor for increasing cost multiplier on total penalty delta vectors [SLP]	p. 185
XSLP_DELTAFORMAT, SLPDELTAFORMAT	Formatting string for creation of names for SLP delta vectors [Logging, SLP]	p. 261
XSLP_DELTAMAXCOST, SLPDELTAMAXCOST	Maximum penalty cost multiplier for penalty delta vectors [SLP]	p. 185
XSLP_DELTAOFFSET, SLPDELTAOFFSET	Position of first character of SLP variable name used to create name of delta vector [Logging, SLP]	p. 227
XSLP_DELTAZLIMIT, SLPDELTAZLIMIT	Number of SLP iterations during which to apply XSLP_DELTA_Z [SLP]	p. 228
XSLP_DERIVATIVES, NLPDERIVATIVES	Bitmap describing the method of calculating derivatives [Derivatives]	p. 228
XSLP_DETERMINISTIC, NLPDETERMINISTIC	Determines if the parallel features of SLP should be guaranteed to be deterministic [MISLP, Parallel, SLP]	p. 229
XSLP_DJTOL, SLPDJTOL	Tolerance on DJ value for determining if a variable is at its step bound [SLP, Tolerances]	p. 186
XSLP_DRCOLDJTOL, SLPDRCOLDJTOL	Reduced cost tolerance on the delta variable when fixing due to the determining column being below XSLP_DRCOLTOL. [Cascading, SLP, Tolerances]	p. 186
XSLP_DRCOLTOL, SLPDRCOLTOL	The minimum absolute magnitude of a determining column, for which the determined variable is still regarded as well defined [Cascading, SLP, Tolerances]	p. 186
XSLP_DRFIXRANGE, SLPDRFIXRANGE	The range around the previous value where variables are fixed in cascading if the determining column is below XSLP_DRCOLTOL. [Cascading, SLP]	p. 187
XSLP_ECFCHECK, SLPECFCHECK	Check feasibility at the point of linearization for extended convergence criteria [SLP, SLP-convergence]	p. 229
XSLP_ECFTOL_A, SLPECFTOL_A	Absolute tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187

XSLP_ECFTOL_R, SLPECFTOL_R	Relative tolerance on testing feasibility at the point of linearization [SLP, SLP-convergence, Tolerances]	p. 187
XSLP_ECHOXPRSMESSAGES	Controls if the XSLP message callback should relay messages from the XPRS library. [Logging]	p. 230
XSLP_ENFORCECOSTSHRINK, SLPENFORCECOSTSHRINK	Factor by which to decrease the current penalty multiplier when enforcing rows. [SLP]	p. 188
XSLP_ENFORCEMAXCOST, SLPENFORCEMAXCOST	Maximum penalty cost in the objective before enforcing most violating rows [SLP]	p. 188
XSLP_ERRORCOST, SLPERRORCOST	Initial penalty cost multiplier for penalty error vectors [SLP]	p. 189
XSLP_ERRORCOSTFACTOR, SLPERRORCOSTFACTOR	Factor for increasing cost multiplier on total penalty error vectors [SLP]	p. 189
XSLP_ERRORMAXCOST, SLPERRORMAXCOST	Maximum penalty cost multiplier for penalty error vectors [SLP]	p. 189
XSLP_ERROROFFSET, SLPERROROFFSET	Position of first character of constraint name used to create name of penalty error vectors [Logging, SLP]	p. 230
XSLP_ERRORTOL_A, SLPERRORTOL_A	Absolute tolerance for error vectors [SLP, Tolerances]	p. 190
XSLP_ERRORTOL_P, SLPERRORTOL_P	Absolute tolerance for printing error vectors [SLP, Tolerances]	p. 190
XSLP_ESCALATION, SLPESCALATION	Factor for increasing cost multiplier on individual penalty error vectors [SLP]	p. 190
XSLP_ETOL_A, SLPETOL_A	Absolute tolerance on penalty vectors [SLP, Tolerances]	p. 191
XSLP_ETOL_R, SLPETOL_R	Relative tolerance on penalty vectors [SLP, Tolerances]	p. 191
XSLP_EVALUATE, NLPEVALUATE	Evaluation strategy for user functions [User Functions]	p. 230
XSLP_EVTOL_A, SLPEVTOL_A	Absolute tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 191
XSLP_EVTOL_R, SLPEVTOL_R	Relative tolerance on total penalty costs [SLP, SLP-convergence, Tolerances]	p. 192
XSLP_EXPAND, SLPEXPAND	Multiplier to increase a step bound [SLP]	p. 192
XSLP_FEAOSTLTARGET, SLPFEAOSTLTARGET	When set, this defines a target feasibility tolerance to which the linearizations are solved to [Linearizations, SLP, Tolerances]	p. 193

XSLP_FILTER, SLPFILTER	Bit map for controlling solution updates [Bit-vector, SLP, Solution]	p. 231
XSLP_FINDIV, NLPFINDIV	Option for running a heuristic to find a feasible initial point [Heuristics]	p. 231
XSLP_FUNCEVAL, NLPFUNCEVAL	Bit map for determining the method of evaluating user functions and their derivatives [Derivatives, User Functions]	p. 232
XSLP_GRANULARITY, SLPGRANULARITY	Base for calculating penalty costs [SLP]	p. 193
XSLP_GRIDHEURSELECT, SLPGRIDHEURSELECT	Bit map selectin which heuristics to run if the problem has variable with an integer delta [Heuristics, SLP]	p. 233
XSLP_HESSIAN, NLPHESSIAN	Second order differentiation mode when using analytical derivatives [Derivatives]	p. 233
XSLP_HEURSTRATEGY, SLPHEURSTRATEGY	Branch and Bound: This specifies the MINLP heuristic strategy. [Heuristics, SLP]	p. 233
XSLP_INFEASLIMIT, SLPINFEASLIMIT	The maximum number of consecutive infeasible SLP iterations which can occur before Xpress-SLP terminates [Limits, SLP]	p. 234
XSLP_INFINITY, NLPINFINITY	Value returned by a divide-by-zero in a formula [Numerics]	p. 193
XSLP_ITERFALLBACKOPS, SLPITERFALLBACKOPS	Alternative LP level control values for numerically challenging problems [Linearizations, Numerics, SLP]	p. 261
XSLP_ITERLIMIT, SLPITERLIMIT	The maximum number of SLP iterations [Limits, SLP]	p. 234
XSLP_ITOL_A, SLPITOL_A	Absolute impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_ITOL_R, SLPITOL_R	Relative impact convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 194
XSLP_IVNAME, NLPIVNAME	Name of the set of initial values to be used [File IO]	p. 262
XSLP_JACOBIAN, NLPJACOBIAN	First order differentiation mode when using analytical derivatives [Derivatives]	p. 234
XSLP_KEEPEQUALSCOLUMN, NLPKEEPEQUALSCOLUMN	When set to a nonzero value, the MPS reader will keep the equals column in the problem [Misc]	p. 235
XSLP_LINQUADBR, NLP_LINQUADBR	Use linear and quadratic constraints and objective function to further reduce bounds on all variables [Presolve]	p. 235
XSLP_LOG, NLPLOG	Level of printing during SLP iterations [Logging, SLP]	p. 236

XSLP_LSITERLIMIT, SLPLSITERLIMIT	Number of iterations in the line search [Limits, SLP]	p. 236
XSLP_LSPATTERNLIMIT, SLPLSPATTERNLIMIT	Number of iterations in the pattern search preceding the line search [Limits, SLP]	p. 236
XSLP_LSSTART, SLPLSSTART	Iteration in which to active the line search [SLP]	p. 237
XSLP_LSZEROLIMIT, SLPLSZEROLIMIT	Maximum number of zero length line search steps before line search is deactivated [Limits, SLP]	p. 237
XSLP_MATRIXTOL, SLPMATRIXTOL	Nonzero tolerance for dropping coefficients from the linearization. [SLP]	p. 195
XSLP_MAXTIME, NLPMAXTIME	The maximum time in seconds that the SLP optimization will run before it terminates [Limits]	p. 237
XSLP_MAXWEIGHT, SLPMAXWEIGHT	Maximum penalty weight for delta or error vectors [SLP]	p. 195
XSLP_MEMORYFACTOR	Factor for expanding size of dynamic arrays in memory [Memory]	p. 196
XSLP_MERITLAMBDA, NLPMERITLAMBDA	Factor by which the net objective is taken into account in the merit function [SLP, Solution]	p. 196
XSLP_MINSBFACTOR, SLPMINSBFACTOR	Factor by which step bounds can be decreased beneath XSLP_ATOL_A [SLP]	p. 197
XSLP_MINUSDELTAFORMAT, SLPMINUSDELTAFORMAT	Formatting string for creation of names for SLP negative penalty delta vectors [Logging, SLP]	p. 262
XSLP_MINUSERERRORFORMAT, SLPMINUSERERRORFORMAT	Formatting string for creation of names for SLP negative penalty error vectors [Logging, SLP]	p. 262
XSLP_MINWEIGHT, SLPMINWEIGHT	Minimum penalty weight for delta or error vectors [SLP]	p. 197
XSLP_MIPALGORITHM, SLPMIPALGORITHM	Bitmap describing the MISLP algorithms to be used [Bit-vector, MISLP]	p. 238
XSLP_MIPCUTOFF_A, SLPMIPCUTOFF_A	Absolute objective function cutoff for MIP termination [MISLP, Tolerances]	p. 197
XSLP_MIPCUTOFF_R, SLPMIPCUTOFF_R	Absolute objective function cutoff for MIP termination [MISLP, Tolerances]	p. 198
XSLP_MIPCUTOFFCOUNT, SLPMIPCUTOFFCOUNT	Number of SLP iterations to check when considering a node for cutting off [Limits, MISLP]	p. 239

XSLP_MIPCUTOFFLIMIT, SLPMIPCUTOFFLIMIT	Number of SLP iterations to check when considering a node for cutting off [Limits, MISLP]	p. 239
XSLP_MIPDEFAULTALGORITHM, SLPMIPDEFAULTALGORITHM	Default algorithm to be used during the tree search in MISLP [Linearizations, MISLP]	p. 240
XSLP_MIPERRORTOL_A, SLPMIPERRORTOL_A	Absolute penalty error cost tolerance for MIP cut-off [MISLP, Tolerances]	p. 198
XSLP_MIPERRORTOL_R, SLPMIPERRORTOL_R	Relative penalty error cost tolerance for MIP cut-off [MISLP, Tolerances]	p. 199
XSLP_MIPFIXSTEPBOUNDS, SLPMIPFIXSTEPBOUNDS	Bitmap describing the step-bound fixing strategy during MISLP [MISLP]	p. 240
XSLP_MIPITERLIMIT, SLPMIPITERLIMIT	Maximum number of SLP iterations at each node [Limits, MISLP]	p. 241
XSLP_MIPLOG, SLPMIPLOG	Frequency with which MIP status is printed [Logging, MISLP]	p. 241
XSLP_MIPOCOUNT, SLPMIPOCOUNT	Number of SLP iterations at each node over which to measure objective function variation [MISLP, SLP-convergence]	p. 241
XSLP_MIPOTOL_A, SLPMIPOTOL_A	Absolute objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MIPOTOL_R, SLPMIPOTOL_R	Relative objective function tolerance for MIP termination [MISLP, SLP-convergence, Tolerances]	p. 199
XSLP_MIPRELAXSTEPBOUNDS, SLPMIPRELAXSTEPBOUNDS	Bitmap describing the step-bound relaxation strategy during MISLP [MISLP]	p. 242
XSLP_MS_MAXBOUNDRANGE, MS_MAXBOUNDRANGE	Defines the maximum range inside which initial points are generated by multistart presets [Multistart]	p. 200
XSLP_MTOL_A, SLPMTOL_A	Absolute effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 200
XSLP_MTOL_R, SLPMTOL_R	Relative effective matrix element convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 201
XSLP_MULTISTART, MULTISTART	The multistart master control. [Multistart]	p. 242
XSLP_MULTISTART_LOG, MULTISTART_LOG	The level of logging during the multistart run. [Logging, Multistart]	p. 243
XSLP_MULTISTART_MAXSOLVES, MULTISTART_MAXSOLVES	The maximum number of jobs to create during the multistart search. [Limits, Multistart]	p. 243
XSLP_MULTISTART_MAXTIME, MULTISTART_MAXTIME	The maximum total time to be spent in the multistart search. [Limits, Multistart]	p. 243

XSLP_MULTISTART_POOLSIZE, MULTISTART_POOLSIZE	The maximum number of problem objects allowed to pool up before synchronization in the deterministic multistart.	[Limits, Multistart]	p. 244
XSLP_MULTISTART_SEED, MULTISTART_SEED	Random seed used for the automatic generation of initial point when loading multistart presets	[Multistart]	p. 244
XSLP_MULTISTART_THREADS, MULTISTART_THREADS	The maximum number of threads to be used in multistart	[Multistart, Parallel]	p. 245
XSLP_MVTOL, SLPMVTOL	Marginal value tolerance for determining if a constraint is slack	[SLP, SLP-convergence, Tolerances]	p. 202
XSLP_NLPSOLVER, NLPSOLVER	Controls whether to call FICO Xpress Global or one of the local solvers	[Solution Process]	p. 244
XSLP_OBJTHRESHOLD, SLPOBJTHRESHOLD	Assumed maximum value of the objective function in absolute value.	[SLP]	p. 202
XSLP_OBJTOPENALTYCOST, SLPOBJTOPENALTYCOST	Factor to estimate initial penalty costs from objective function	[SLP]	p. 203
XSLP_OCOUNT, SLPOCOUNT	Number of SLP iterations over which to measure objective function variation for static objective (2) convergence criterion	[SLP, SLP-convergence]	p. 245
XSLP_OPTIMALITYTOLTARGET, SLPOPTIMALITYTOLTARGET	When set, this defines a target optimality tolerance to which the linearizations are solved to	[Linearizations, SLP, Tolerances]	p. 203
XSLP_OTOL_A, SLPOTOL_A	Absolute static objective (2) convergence tolerance	[SLP, SLP-convergence, Tolerances]	p. 204
XSLP_OTOL_R, SLPOTOL_R	Relative static objective (2) convergence tolerance	[SLP, SLP-convergence, Tolerances]	p. 204
XSLP_PENALTYCOLFORMAT, SLPPENALTYCOLFORMAT	Formatting string for creation of the names of the SLP penalty transfer vectors	[Logging, SLP]	p. 263
XSLP_PENALTYINFOSTART, SLPPENALTYINFOSTART	Iteration from which to record row penalty information	[Logging, SLP]	p. 246
XSLP_PENALTYROWFORMAT, SLPPENALTYROWFORMAT	Formatting string for creation of the names of the SLP penalty rows	[Logging, SLP]	p. 263
XSLP_PLUSDELTAFORMAT, SLPPPLUSDELTAFORMAT	Formatting string for creation of names for SLP positive penalty delta vectors	[Logging, SLP]	p. 263
XSLP_PLUSEXPORTFORMAT, SLPPPLUSEXPORTFORMAT	Formatting string for creation of names for SLP positive penalty error vectors	[Logging, SLP]	p. 264

XSLP_POSTSOLVE, NLPPOSTSOLVE	This control determines whether postsolving should be performed automatically	[Presolve]	p. 246
XSLP_PRESOLVE, NLPPRESOLVE	This control determines whether presolving should be performed on the nonlinear problem prior to starting the main algorithm	[Presolve]	p. 246
XSLP_PRESOLVE_ELIMTOL, NLPPRESOLVE_ELIMTOL	Tolerance for nonlinear eliminations during SLP presolve	[Presolve, Tolerances]	p. 205
XSLP_PRESOLVELEVEL, NLPPRESOLVELEVEL	This control determines the level of changes presolve may carry out on the problem and whether column/row indices may change	[Presolve]	p. 247
XSLP_PRESOLVEOPS, NLPPRESOLVEOPS	Bitmap indicating the SLP presolve actions to be taken	[Bit-vector, Presolve]	p. 247
XSLP_PRESOLVEZERO, NLPPRESOLVEZERO	Minimum absolute value for a variable which is identified as nonzero during SLP presolve	[Presolve, Tolerances]	p. 205
XSLP_PRIMALINTEGRALALPHA, NLPPRIMALINTEGRALALPHA	Decay term for primal integral computation	[Logging]	p. 205
XSLP_PRIMALINTEGRALREF, NLPPRIMALINTEGRALREF	Reference solution value to take into account when calculating the primal integral	[Logging]	p. 206
XSLP_PROBING, NLPPROBING	This control determines whether probing on a subset of variables should be performed prior to starting the main algorithm.	[Presolve]	p. 248
XSLP_REFORMULATE, NLPREFORMULATE	Controls the problem reformulations carried out before augmentation.	[Presolve]	p. 248
XSLP_SAMECOUNT, SLPSAMECOUNT	Number of steps reaching the step bound in the same direction before step bounds are increased	[SLP]	p. 249
XSLP_SAMEDAMP, SLPSAMEDAMP	Number of steps in same direction before damping factor is increased	[SLP]	p. 250
XSLP_SBLOROWFORMAT, SLPSBLOROWFORMAT	Formatting string for creation of names for SLP lower step bound rows	[Logging, SLP]	p. 264
XSLP_SBNAME, SLPSBNAME	Name of the set of initial step bounds to be used	[Logging, SLP]	p. 265
XSLP_SBROWOFFSET, SLPSBROWOFFSET	Position of first character of SLP variable name used to create name of SLP lower and upper step bound rows	[Logging, SLP]	p. 250
XSLP_SBSTART, SLPSBSTART	SLP iteration after which step bounds are first applied	[SLP]	p. 250

XSLP_SBUPROWFORMAT, SLP_SBUPROWFORMAT	Formatting string for creation of names for SLP upper step bound rows [Logging, SLP]	p. 265
XSLP_SCALE	When to re-scale the SLP problem [Numerics, SLP]	p. 251
XSLP_SCALECOUNT	Iteration limit used in determining when to re-scale the SLP matrix [Numerics, SLP]	p. 251
XSLP_SHRINK, SLP_SHRINK	Multiplier to reduce a step bound [SLP]	p. 206
XSLP_SHRINKBIAS, SLP_SHRINKBIAS	Defines an overwrite / adjustment of step bounds for improving iterations [SLP]	p. 206
XSLP_SLPLOG	Frequency with which SLP status is printed [Logging, SLP]	p. 252
XSLP_SOLVER, LOCALSOLVER	Selects the library to use for local solves [Solution Process]	p. 252
XSLP_STOL_A, SLP_STOL_A	Absolute slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_STOL_R, SLP_STOL_R	Relative slack convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 207
XSLP_STOPOUTOFRANGE, NLP_STOPOUTOFRANGE	Stop optimization and return error code if internal function argument is out of range [SLP]	p. 252
XSLP_THREADS, NLP_THREADS	Default number of threads to be used [Parallel]	p. 253
XSLP_THREADSAFEUSERFUNC, NLP_THREADSAFEUSERFUNC	Defines if user functions are allowed to be called in parallel [Parallel, User Functions]	p. 253
XSLP_TOLNAME, SLP_TOLNAME	Name of the set of tolerance sets to be used [File IO, SLP]	p. 265
XSLP_TRACEMASK, SLP_TRACEMASK	Mask of variable or row names that are to be traced through the SLP iterates [Logging, SLP]	p. 266
XSLP_TRACEMASKOPS, SLP_TRACEMASKOPS	Controls the information printed for XSLP_TRACEMASK. [Bit-vector, Logging, SLP]	p. 253
XSLP_UNFINISHEDLIMIT, SLP_UNFINISHEDLIMIT	The number of consecutive SLP iterations that may have an unfinished status before the solve is terminated. [Linearizations, SLP]	p. 254
XSLP_UPDATEFORMAT, SLP_UPDATEFORMAT	Formatting string for creation of names for SLP update rows [Logging, SLP]	p. 266
XSLP_UPDATEOFFSET, SLP_UPDATEOFFSET	Position of first character of SLP variable name used to create name of SLP update row [Logging, SLP]	p. 255

XSLP_VALIDATIONFACTOR, NLPVALIDATIONFACTOR	Minimum improvement in validation targets to continue iterating [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_K, NLPVALIDATIONTARGET_K	Optimality target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTARGET_R, NLPVALIDATIONTARGET_R	Feasibility target tolerance [SLP, SLP-convergence, Tolerances]	p. 208
XSLP_VALIDATIONTOL_A, NLPVALIDATIONTOL_A	Absolute tolerance for the XSLPvalidate procedure [SLP, Tolerances]	p. 209
XSLP_VALIDATIONTOL_K, NLPVALIDATIONTOL_K	Relative tolerance for the XSLPvalidate procedure [SLP, Tolerances]	p. 209
XSLP_VALIDATIONTOL_R, NLPVALIDATIONTOL_R	Relative tolerance for the XSLPvalidate procedure [Tolerances]	p. 209
XSLP_VCOUNT, SLPVCOUNT	Number of SLP iterations over which to measure static objective (3) convergence [SLP, SLP-convergence]	p. 255
XSLP_VLIMIT, SLPVLIMIT	Number of SLP iterations after which static objective (3) convergence testing starts [SLP, SLP-convergence]	p. 256
XSLP_VTOL_A, SLPVTOL_A	Absolute static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 210
XSLP_VTOL_R, SLPVTOL_R	Relative static objective (3) convergence tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WCOUNT, SLPWCOUNT	Number of SLP iterations over which to measure the objective for the extended convergence continuation criterion [SLP, SLP-convergence]	p. 256
XSLP_WTOL_A, SLPWTOL_A	Absolute extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 211
XSLP_WTOL_R, SLPWTOL_R	Relative extended convergence continuation tolerance [SLP, SLP-convergence, Tolerances]	p. 212
XSLP_XCOUNT, SLPXCOUNT	Number of SLP iterations over which to measure static objective (1) convergence [SLP, SLP-convergence]	p. 257
XSLP_XLIMIT, SLPXLIMIT	Number of SLP iterations up to which static objective (1) convergence testing is performed [Limits, SLP, SLP-convergence]	p. 258
XSLP_XTOL_A, SLPXTOL_A	Absolute static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 213
XSLP_XTOL_R, SLPXTOL_R	Relative static objective function (1) tolerance [SLP, SLP-convergence, Tolerances]	p. 214
XSLP_ZERO, NLPZERO	Absolute tolerance [Tolerances]	p. 215

XSLP_ZEROCRITERION, SLPZEROCRITERION	Bitmap determining the behavior of the placeholder deletion procedure [Bit-vector, SLP]	p. 259
XSLP_ZEROCRITERIONCOUNT, SLPZEROCRITERIONCOUNT	Number of consecutive times a placeholder entry is zero before being considered for deletion [Limits, SLP]	p. 259
XSLP_ZEROCRITERIONSTART, SLPZEROCRITERIONSTART	SLP iteration at which criteria for deletion of placeholder entries are first activated. [SLP]	p. 260

20.1 Double control parameters

XSLP_ATOL_A, SLPATOL_A

Description	Absolute delta convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	<p>The absolute delta convergence criterion assesses the change in value of a variable (δX) against the absolute delta convergence tolerance. If $\delta X < \text{XSLP_ATOL_A}$ then the variable has converged on the absolute delta convergence criterion. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target $\text{XSLP_VALIDATIONTARGET_R}$. Good values for the control are usually fall between $1\text{e-}3$ and $1\text{e-}6$.</p>
Default value	-1.0
See also	Convergence Criteria, XSLP_ATOL_R

XSLP_ATOL_R, SLPATOL_R

Description	Relative delta convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	<p>The relative delta convergence criterion assesses the change in value of a variable (δX) relative to the value of the variable (X), against the relative delta convergence tolerance. If $\delta X < X * \text{XSLP_ATOL_R}$ then the variable has converged on the relative delta convergence criterion. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target $\text{XSLP_VALIDATIONTARGET_R}$. Good values for the control are usually fall between $1\text{e-}3$ and $1\text{e-}6$.</p>
Default value	-1.0
See also	Convergence Criteria, XSLP_ATOL_A

XSLP_BARSTALLINGTOL, SLPBARSTALLINGTOL

Description	Required change in the objective when progress is measured in barrier iterations without crossover
Type	Double
Category	Control
Topic areas	SLP, Linearizations

Note	Minumum objective variability change required in relation to control XSLP_BARSTALLINGOBLIMIT for the iterations to be regarded as making progress. The net objective, error cost and error sum are taken into account.
Default value	0.05
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_BARCROSSOVERSTART, XSLP_BARLIMIT, XSLP_BARSTARTOPS, XSLP_BARSTALLINGLIMIT, XSLP_BARSTALLINGOBLIMIT

XSLP_BOUNDTHRESHOLD, SLPBOUNDTHRESHOLD

Description	The maximum size of a bound that can be introduced by nonlinear presolve.
Type	Double
Category	Control
Topic areas	SLP, Presolve
Default value	1.0e+10
Affects routines	XSLPmaxim, XSLPminim

XSLP_CASCADETOL_PA, SLPCASCADETOL_PA

Description	Absolute cascading print tolerance
Type	Double
Category	Control
Topic areas	SLP, Cascading, Logging
Note	The change to the value of a variable as a result of cascading is only printed if the change is deemed significant. The change is tested against: absolute and relative convergence tolerance and absolute and relative cascading print tolerance. The change is printed only if all tests fail. The absolute cascading print criterion measures the change in value of a variable (δX) against the absolute cascading print tolerance. If $\delta X < \text{XSLP_CASCADETOL_PA}$ then the change is within the absolute cascading print tolerance and will not be printed. XSLP_LOG must be at least 5 for this control to have an effect.
Default value	0.01
See also	Cascading, XSLP_CASCADETOL_PR
Affects routines	XSLPcascade

XSLP_CASCADETOL_PR, SLPCASCADETOL_PR

Description	Relative cascading print tolerance
Type	Double
Category	Control
Topic areas	SLP, Cascading, Logging

Note	The change to the value of a variable as a result of cascading is only printed if the change is deemed significant. The change is tested against: absolute and relative convergence tolerance and absolute and relative cascading print tolerance. The change is printed only if all tests fail. The relative cascading print criterion measures the change in value of a variable (δX) relative to the value of the variable (X), against the relative cascading print tolerance. If $\delta X < X * \text{XSLP_CASCADETOL_PR}$ then the change is within the relative cascading print tolerance and will not be printed. <code>XSLP_LOG</code> must be at least 5 for this control to have an effect.
Default value	0.01
See also	Cascading, <code>XSLP_CASCADETOL_PA</code>
Affects routines	<code>XSLPcascade</code>

XSLP_CDTOL_A, SLPCDTOL_A

Description	Absolute tolerance for deducing constant derivatives
Type	Double
Category	Control
Topic area	Tolerances
Note	<p>The absolute tolerance test for constant derivatives is used as follows: If the value of the user function at point X_0 is Y_0 and the values at $(X_0 - \delta X)$ and $(X_0 + \delta X)$ are Y_d and Y_u respectively, then the numerical derivatives at X_0 are: "down" derivative $D_d = (Y_0 - Y_d)/\delta X$ "up" derivative $D_u = (Y_u - Y_0)/\delta X$ If $\text{abs}(D_d - D_u) \leq \text{XSLP_CDTOL_A}$ then the derivative is regarded as constant.</p>
Default value	1.0e-08
See also	<code>XSLP_CDTOL_R</code>

XSLP_CDTOL_R, SLPCDTOL_R

Description	Relative tolerance for deducing constant derivatives
Type	Double
Category	Control
Topic area	Tolerances
Note	<p>The relative tolerance test for constant derivatives is used as follows: If the value of the user function at point X_0 is Y_0 and the values at $(X_0 - \delta X)$ and $(X_0 + \delta X)$ are Y_d and Y_u respectively, then the numerical derivatives at X_0 are: "down" derivative $D_d = (Y_0 - Y_d)/\delta X$ "up" derivative $D_u = (Y_u - Y_0)/\delta X$ If $\text{abs}(D_d - D_u) \leq \text{XSLP_CDTOL_R} * \text{abs}(Y_d + Y_u)/2$ then the derivative is regarded as constant.</p>
Default value	1.0e-08
See also	<code>XSLP_CDTOL_A</code>

XSLP_CLAMPSHRINK, SLPCLAMPSHRINK

Description	Shrink ratio used to impose strict convergence on variables converged in extended criteria only
Type	Double
Category	Control
Topic area	SLP
Note	If the solution has converged but there are variables converged on extended criteria only, the XSLP_CLAMPSHRINK acts as a shrinking ratio on the step bounds and the problem is optimized (if necessary multiple times), with the purpose of expediting strict convergence on all variables. XSLP_ALGORITHM controls if this shrinking is applied at all, and if shrinking is applied to of the variables converged on extended criteria only with active step bounds only, or if on all variables.
Default value	0.3
See also	XSLP_ALGORITHM, XSLP_CLAMPVALIDATIONTOL_A, XSLP_CLAMPVALIDATIONTOL_R

XSLP_CLAMPVALIDATIONTOL_A, SLPCLAMPVALIDATIONTOL_A

Description	Absolute validation tolerance for applying XSLP_CLAMPSHRINK
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	If set and the absolute validation value is larger than this value, then control XSLP_CLAMPSHRINK is checked once the solution has converged, but there are variables converged on extended criteria only.
Default value	1.0e-06
See also	XSLP_ALGORITHM, XSLP_CLAMPSHRINK, XSLP_CLAMPVALIDATIONTOL_R

XSLP_CLAMPVALIDATIONTOL_R, SLPCLAMPVALIDATIONTOL_R

Description	Relative validation tolerance for applying XSLP_CLAMPSHRINK
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	If set and the relative validation value is larger than this value, then control XSLP_CLAMPSHRINK is checked once the solution has converged, but there are variables converged on extended criteria only.
Default value	1.0e-06
See also	XSLP_ALGORITHM, XSLP_CLAMPSHRINK, XSLP_CLAMPVALIDATIONTOL_A

XSLP_CTOL, SLPCTOL

Description	Closure convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Notes	<p>The closure convergence criterion measures the change in value of a variable (δX) relative to the value of its initial step bound (B), against the closure convergence tolerance. If $\delta X < B * XSLP_CTOL$ then the variable has converged on the closure convergence criterion.</p> <p>If no explicit initial step bound is provided, then the test will not be applied and the variable can never converge on the closure criterion. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target <code>XSLP_VALIDATIONTARGET_R</code>. Good values for the control are usually fall between $1e-3$ and $1e-6$.</p>
Default value	-1.0
See also	Convergence Criteria, <code>XSLP_ATOL_A</code> , <code>XSLP_ATOL_R</code>

XSLP_DAMP, SLPDAMP

Description	Damping factor for updating values of variables
Type	Double
Category	Control
Topic area	SLP
Note	<p>The damping factor sets the next <i>assumed value</i> for a variable based on the previous assumed value (X_0) and the <i>actual value</i> (X_1). The new assumed value is given by $X_1 * XSLP_DAMP + X_0 * (1 - XSLP_DAMP)$</p>
Default value	1
See also	Xpress-SLP Solution Process, <code>XSLP_DAMPEXPAND</code> <code>XSLP_DAMPMAX</code> , <code>XSLP_DAMPMIN</code> , <code>XSLP_DAMP SHRINK</code> , <code>XSLP_DAMPSTART</code>
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_DAMPEXPAND, SLPDAMPEXPAND

Description	Multiplier to increase damping factor during dynamic damping
Type	Double
Category	Control
Topic area	SLP
Note	<p>If dynamic damping is enabled, the damping factor for a variable will be increased if successive changes are in the same direction. More precisely, if there are <code>XSLP_SAMEDAMP</code> successive changes in the same direction for a variable, then the damping factor (D) for the variable will be reset to $D * XSLP_DAMPEXPAND + XSLP_DAMPMAX * (1 - XSLP_DAMPEXPAND)$</p>

Default value	1
See also	Xpress-SLP Solution Process, XSLP_ALGORITHM, XSLP_DAMP, XSLP_DAMPMAX, XSLP_DAMPMIN, XSLP_DAMPSHRINK, XSLP_DAMPSTART, XSLP_SAMEDAMP
Affects routines	XSLPmaxim, XSLPminim

XSLP_DAMPMAX, SLPDAMPMAX

Description	Maximum value for the damping factor of a variable during dynamic damping
Type	Double
Category	Control
Topic area	SLP
Note	If dynamic damping is enabled, the damping factor for a variable will be increased if successive changes are in the same direction. More precisely, if there are XSLP_SAMEDAMP successive changes in the same direction for a variable, then the damping factor (D) for the variable will be reset to $D * XSLP_DAMPEXPAND + XSLP_DAMPMAX * (1 - XSLP_DAMPEXPAND)$
Default value	1
See also	Xpress-SLP Solution Process, XSLP_ALGORITHM, XSLP_DAMP, XSLP_DAMPEXPAND, XSLP_DAMPMIN, XSLP_DAMPSHRINK, XSLP_DAMPSTART, XSLP_SAMEDAMP
Affects routines	XSLPmaxim, XSLPminim

XSLP_DAMPMIN, SLPDAMPMIN

Description	Minimum value for the damping factor of a variable during dynamic damping
Type	Double
Category	Control
Topic area	SLP
Note	If dynamic damping is enabled, the damping factor for a variable will be decreased if successive changes are in the opposite direction. More precisely, the damping factor (D) for the variable will be reset to $D * XSLP_DAMPSHRINK + XSLP_DAMPMIN * (1 - XSLP_DAMPEXPAND)$
Default value	1
See also	Xpress-SLP Solution Process, XSLP_ALGORITHM, XSLP_DAMP, XSLP_DAMPEXPAND, XSLP_DAMPMAX, XSLP_DAMPSHRINK, XSLP_DAMPSTART
Affects routines	XSLPmaxim, XSLPminim

XSLP_DAMPSHRINK, SLPDAMPSHRINK

Description	Multiplier to decrease damping factor during dynamic damping
Type	Double
Category	Control
Topic area	SLP
Note	If dynamic damping is enabled, the damping factor for a variable will be decreased if successive changes are in the opposite direction. More precisely, the damping factor (D) for the variable will be reset to $D * XSLP_DAMPSHRINK + XSLP_DAMPMIN * (1 - XSLP_DAMPEXPAND)$
Default value	1
See also	Xpress-SLP Solution Process, XSLP_ALGORITHM, XSLP_DAMP, XSLP_DAMPEXPAND, XSLP_DAMPMAX, XSLP_DAMPMIN, XSLP_DAMPSTART
Affects routines	XSLPmaxim, XSLPminim

XSLP_DEFAULTIV, NLPDEFAULTIV

Description	Default initial value for an SLP variable if none is explicitly given
Type	Double
Category	Control
Topic area	Data Input
Note	If no initial value is given for an SLP variable, then the initial value provided for the "equals column" will be used. If no such value has been provided, then XSLP_DEFAULTIV will be used. If this is above the upper bound for the variable, then the upper bound will be used; if it is below the lower bound for the variable, then the lower bound will be used.
Default value	100
Affects routines	XSLPconstruct

XSLP_DEFAULTSTEPBOUND, SLPDEFAULTSTEPBOUND

Description	Minimum initial value for the step bound of an SLP variable if none is explicitly given
Type	Double
Category	Control
Topic area	SLP
Notes	If no initial step bound value is given for an SLP variable, this will be used as a minimum value. If the algorithm is estimating step bounds, then the step bound actually used for a variable may be larger than the default. A default initial step bound is ignored when testing for the closure tolerance XSLP_CTOL: if there is no specific value, then the test will not be applied.
Default value	16
See also	XSLP_CTOL
Affects routines	XSLPconstruct

XSLP_DELTA_A, SLPDELTA_A

Description	Absolute perturbation of values for calculating numerical derivatives
Type	Double
Category	Control
Topic area	Derivatives
Note	First-order derivatives are calculated by perturbing the value of each variable in turn by a small amount. The amount is determined by the absolute and relative delta factors as follows: $\text{XSLP_DELTA_A} + \text{abs}(X) * \text{XSLP_DELTA_R}$ where (X) is the current value of the variable. If the perturbation takes the variable outside a bound, then the perturbation normally made only in the opposite direction.
Default value	0.001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_DELTA_R

XSLP_DELTA_INFINITY, SLPDELTA_INFINITY

Description	Maximum value for partial derivatives
Type	Double
Category	Control
Topic area	Derivatives
Default value	1.0e+15
Affects routines	XSLPmaxim, XSLPminim

XSLP_DELTA_R, SLPDELTA_R

Description	Relative perturbation of values for calculating numerical derivatives
Type	Double
Category	Control
Topic area	Derivatives
Note	First-order derivatives are calculated by perturbing the value of each variable in turn by a small amount. The amount is determined by the absolute and relative delta factors as follows: $\text{XSLP_DELTA_A} + \text{abs}(X) * \text{XSLP_DELTA_R}$ where (X) is the current value of the variable. If the perturbation takes the variable outside a bound, then the perturbation normally made only in the opposite direction.
Default value	0.001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_DELTA_A

XSLP_DELTA_X, SLPDELTA_X

Description	Minimum absolute value of delta coefficients to be retained
Type	Double
Category	Control
Topic area	Derivatives
Notes	If the value of a coefficient in a delta column is less than this value, it will be reset to zero. Larger values of XSLP_DELTA_X will result in matrices with fewer elements, which may be easier to solve. However, there will be increased likelihood of local optima as some of the small relationships between variables and constraints are deleted. There may also be increased difficulties with singular bases resulting from deletion of pivot elements from the matrix.
Default value	1.0e-6
Affects routines	XSLPmaxim, XSLPminim

XSLP_DELTA_Z, SLPDELTA_Z

Description	Tolerance used when calculating derivatives
Type	Double
Category	Control
Topic areas	Derivatives, Tolerances
Notes	If the absolute value of a variable is less than this value, then a value of XSLP_DELTA_Z will be used instead for calculating derivatives. If a nonzero derivative is calculated for a formula which always results in a matrix coefficient less than XSLP_DELTA_Z, then a larger value will be substituted so that at least one of the coefficients is XSLP_DELTA_Z in magnitude. If XSLP_DELTAZLIMIT is set to a positive number, then when that number of iterations have passed, values smaller than XSLP_DELTA_Z will be set to zero.
Default value	0.00001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_DELTAZLIMIT, XSLP_DELTA_ZERO

XSLP_DELTA_ZERO, SLPDELTA_ZERO

Description	Absolute zero acceptance tolerance used when calculating derivatives
Type	Double
Category	Control
Topic areas	Derivatives, Tolerances
Notes	Provides an override value for the XSLP_DELTA_Z behavior. Derivatives smaller than XSLP_DELTA_ZERO will not be substituted by XSLP_DELTA_Z, defining a range in which derivatives are deemed nonzero and are affected by XSLP_DELTA_Z. A negative value means that this tolerance will not be applied.

Default value	-1.0 (not applied)
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_DELTAZLIMIT, XSLP_DELTA_Z

XSLP_DELTACOST, SLPDELTACOST

Description	Initial penalty cost multiplier for penalty delta vectors
Type	Double
Category	Control
Topic area	SLP
Note	If penalty delta vectors are used, this parameter sets the initial cost factor. If there are active penalty delta vectors, then the penalty cost may be increased.
Default value	200
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_DELTACOSTFACTOR, XSLP_DELTAMAXCOST, XSLP_ERRORCOST

XSLP_DELTACOSTFACTOR, SLPDELTACOSTFACTOR

Description	Factor for increasing cost multiplier on total penalty delta vectors
Type	Double
Category	Control
Topic area	SLP
Note	If there are active penalty delta vectors, then the penalty cost multiplier will be increased by a factor of XSLP_DELTACOSTFACTOR up to a maximum of XSLP_DELTAMAXCOST
Default value	1.3
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_DELTACOST, XSLP_DELTAMAXCOST, XSLP_ERRORCOST

XSLP_DELTAMAXCOST, SLPDELTAMAXCOST

Description	Maximum penalty cost multiplier for penalty delta vectors
Type	Double
Category	Control
Topic area	SLP
Note	If there are active penalty delta vectors, then the penalty cost multiplier will be increased by a factor of XSLP_DELTACOSTFACTOR up to a maximum of XSLP_DELTAMAXCOST
Default value	XPRS_PLUSINFINITY
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_DELTACOST, XSLP_DELTACOSTFACTOR, XSLP_ERRORCOST

XSLP_DJTOL, SLPDJTOL

Description	Tolerance on DJ value for determining if a variable is at its step bound
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	If a variable is at its step bound and within the absolute delta tolerance XSLP_ATOL_A or closure tolerance XSLP_CTOL then the step bounds will not be further reduced. If the DJ is greater in magnitude than XSLP_DJTOL then the step bound may be relaxed if it meets the necessary criteria.
Default value	1.0e-6
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ATOL_A, XSLP_CTOL

XSLP_DRCOLDJTOL, SLPDRCOLDJTOL

Description	Reduced cost tolerance on the delta variable when fixing due to the determining column being below XSLP_DRCOLTOL.
Type	Double
Category	Control
Topic areas	SLP, Cascading, Tolerances
Default value	0.0
Affects routines	XSLPconstruct XSLPcascade

XSLP_DRCOLTOL, SLPDRCOLTOL

Description	The minimum absolute magnitude of a determining column, for which the determined variable is still regarded as well defined
Type	Double
Category	Control
Topic areas	SLP, Cascading, Tolerances
Notes	This control affects the cascading procedure. Please see Chapter <i>Cascading</i> for more information.
Default value	1.0e-6
See also	XSLP_CASCADE
Affects routines	XSLPconstruct XSLPcascade

XSLP_DRFIXRANGE, SLPDRFIXRANGE

Description	The range around the previous value where variables are fixed in cascading if the determining column is below XSLP_DRCOLTOL.
Type	Double
Category	Control
Topic areas	SLP, Cascading
Notes	This control affects the cascading procedure. Please see Chapter <i>Cascading</i> for more information.
Default value	0.1
See also	XSLP_CASCADE XSLP_DRCOLTOL
Affects routines	XSLPconstruct XSLPcascade

XSLP_ECFTOL_A, SLPECFTOL_A

Description	Absolute tolerance on testing feasibility at the point of linearization
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Notes	<p>The extended convergence criteria test how well the linearization approximates the true problem. They depend on the point of linearization being a reasonable approximation — in particular, that it should be reasonably close to feasibility. Each constraint is tested at the point of linearization, and the total positive and negative contributions to the constraint from the columns in the problem are calculated. A feasibility tolerance is calculated as the largest of $XSLP_{ECFTOL_A}$ and $\max(\text{abs(Positive)}, \text{abs(Negative)}) * XSLP_{ECFTOL_R}$.</p> <p>If the calculated infeasibility is greater than the tolerance, the point of linearization is regarded as infeasible and the extended convergence criteria will not be applied. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target $XSLP_{VALIDATIONTARGET_R}$. Good values for the control are usually fall between $1e-1$ and $1e-6$.</p>
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	Convergence criteria, XSLP_ECFCHECK, XSLP_ECFCOUNT, XSLP_ECFTOL_R

XSLP_ECFTOL_R, SLPECFTOL_R

Description	Relative tolerance on testing feasibility at the point of linearization
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances

Notes	<p>The extended convergence criteria test how well the linearization approximates the true problem. They depend on the point of linearization being a reasonable approximation — in particular, that it should be reasonably close to feasibility. Each constraint is tested at the point of linearization, and the total positive and negative contributions to the constraint from the columns in the problem are calculated. A feasibility tolerance is calculated as the largest of $XSLP_{ECFTOL_A}$ and $\max(\text{abs(Positive)}, \text{abs(Negative)}) * XSLP_{ECFTOL_R}$. If the calculated infeasibility is greater than the tolerance, the point of linearization is regarded as infeasible and the extended convergence criteria will not be applied. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target $XSLP_{VALIDATIONTARGET_R}$. Good values for the control are usually fall between $1e-1$ and $1e-6$.</p>
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	Convergence criteria, XSLP_ECFCHECK, XSLP_ECFCOUNT, XSLP_ECFTOL_A

XSLP_ENFORCECOSTSHRINK, SLPENFORCECOSTSHRINK

Description	Factor by which to decrease the current penalty multiplier when enforcing rows.
Type	Double
Category	Control
Topic area	SLP
Notes	<p>When feasibility of a row cannot be achieved by increasing the penalty cost on its error variable, removing the variable (fixing it to zero) can force the row to be satisfied, as set by $XSLP_{ENFORCEMAXCOST}$. After the error variables have been removed (which is equivalent to setting to row to be enforced) the penalties on the remaining error variables are rebalanced to allow for a reduction in the size of the penalties in the objective in order to achieve better numerical behaviour.</p>
Default value	0.00001
Affects routines	XSLPmaxim, XSLPminim
See also	$XSLP_{ENFORCEMAXCOST}$

XSLP_ENFORCEMAXCOST, SLPENFORCEMAXCOST

Description	Maximum penalty cost in the objective before enforcing most violating rows
Type	Double
Category	Control
Topic area	SLP
Notes	<p>When feasibility of a row cannot be achieved by increasing the penalty cost on its error variable, removing the variable (fixing it to zero) can force the row to be satisfied. After the error variables have been removed (which is equivalent to setting to row to be enforced) the penalties on the remaining error variables are rebalanced to allow for a reduction in the size of the penalties in the objective in order to achieve better numerical behaviour, controlled by $XSLP_{ENFORCECOSTSHRINK}$.</p>

Default value	1.0e+11
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ENFORCECOSTSHRINK

XSLP_ERRORCOST, SLPERRORCOST

Description	Initial penalty cost multiplier for penalty error vectors
Type	Double
Category	Control
Topic area	SLP
Note	If penalty error vectors are used, this parameter sets the initial cost factor. If there are active penalty error vectors, then the penalty cost may be increased.
Default value	200
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_DELTACOST, XSLP_ERRORCOSTFACTOR, XSLP_ERRORMAXCOST

XSLP_ERRORCOSTFACTOR, SLPERRORCOSTFACTOR

Description	Factor for increasing cost multiplier on total penalty error vectors
Type	Double
Category	Control
Topic area	SLP
Note	If there are active penalty error vectors, then the penalty cost multiplier will be increased by a factor of XSLP_ERRORCOSTFACTOR up to a maximum of XSLP_ERRORMAXCOST
Default value	1.3
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_ERRORMAXCOST

XSLP_ERRORMAXCOST, SLPERRORMAXCOST

Description	Maximum penalty cost multiplier for penalty error vectors
Type	Double
Category	Control
Topic area	SLP
Note	If there are active penalty error vectors, then the penalty cost multiplier will be increased by a factor of XSLP_ERRORCOSTFACTOR up to a maximum of XSLP_ERRORMAXCOST
Default value	XPRS_PLUSINFINITY
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_ERRORCOSTFACTOR

XSLP_ERRORTOL_A, SLPERRORTOL_A

Description	Absolute tolerance for error vectors
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	The solution will be regarded as having no active error vectors if one of the following applies: every penalty error vector and penalty delta vector has an activity less than XSLP_ERRORTOL_A; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than XSLP_EVTOL_A; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than XSLP_EVTOL_R * Obj where Obj is the current objective function value.
Default value	0.00001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_EVTOL_A, XSLP_EVTOL_R

XSLP_ERRORTOL_P, SLPERRORTOL_P

Description	Absolute tolerance for printing error vectors
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	The solution log includes a print of penalty delta and penalty error vectors with an activity greater than XSLP_ERRORTOL_P.
Default value	0.0001
Affects routines	XSLPmaxim, XSLPminim

XSLP_ESCALATION, SLPESCALATION

Description	Factor for increasing cost multiplier on individual penalty error vectors
Type	Double
Category	Control
Topic area	SLP
Note	If penalty cost escalation is activated in XSLP_ALGORITHM then the penalty cost multiplier will be increased by a factor of XSLP_ESCALATION for any active error vector up to a maximum of XSLP_MAXWEIGHT.
Default value	1.25
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ALGORITHM, XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_MAXWEIGHT

XSLP_ETOL_A, SLPETOL_A

Description	Absolute tolerance on penalty vectors
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	For each penalty error vector, the contribution to its constraint is calculated, together with the total positive and negative contributions to the constraint from other vectors. If its contribution is less than XSLP_ETOL_A or less than $\text{Positive} * \text{XSLP_ETOL_R}$ or less than $\text{abs}(\text{Negative}) * \text{XSLP_ETOL_R}$ then it will be regarded as insignificant and will not have its penalty increased. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target XSLP_VALIDATIONTARGET_R. Good values for the control are usually fall between $1e-3$ and $1e-6$.
Default value	0.0001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ETOL_R XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_ESCALATION

XSLP_ETOL_R, SLPETOL_R

Description	Relative tolerance on penalty vectors
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	For each penalty error vector, the contribution to its constraint is calculated, together with the total positive and negative contributions to the constraint from other vectors. If its contribution is less than XSLP_ETOL_A or less than $\text{Positive} * \text{XSLP_ETOL_R}$ or less than $\text{abs}(\text{Negative}) * \text{XSLP_ETOL_R}$ then it will be regarded as insignificant and will not have its penalty increased. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target XSLP_VALIDATIONTARGET_R. Good values for the control are usually fall between $1e-3$ and $1e-6$.
Default value	0.0001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ETOL_A XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_ESCALATION

XSLP_EVTOL_A, SLPEVTOL_A

Description	Absolute tolerance on total penalty costs
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances

Note	The solution will be regarded as having no active error vectors if one of the following applies: every penalty error vector and penalty delta vector has an activity less than $XSLP_ERRORTOL_A$; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than $XSLP_EVTOL_A$; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than $XSLP_EVTOL_R * Obj$ where Obj is the current objective function value. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target $XSLP_VALIDATIONTARGET_R$. Good values for the control are usually fall between $1e-2$ and $1e-6$, but normally a magnitude larger than $XSLP_ETOL_A$.
Default value	-1.0
Affects routines	$XSLPmaxim$, $XSLPminim$
See also	$XSLP_ERRORTOL_A$, $XSLP_EVTOL_R$

$XSLP_EVTOL_R$, $SLPEVTOL_R$

Description	Relative tolerance on total penalty costs
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	The solution will be regarded as having no active error vectors if one of the following applies: every penalty error vector and penalty delta vector has an activity less than $XSLP_ERRORTOL_A$; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than $XSLP_EVTOL_A$; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than $XSLP_EVTOL_R * Obj$ where Obj is the current objective function value. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target $XSLP_VALIDATIONTARGET_R$. Good values for the control are usually fall between $1e-2$ and $1e-6$, but normally a magnitude larger than $XSLP_ETOL_R$.
Default value	-1.0
Affects routines	$XSLPmaxim$, $XSLPminim$
See also	$XSLP_ERRORTOL_A$, $XSLP_EVTOL_A$

$XSLP_EXPAND$, $SLPEXPAND$

Description	Multiplier to increase a step bound
Type	Double
Category	Control
Topic area	SLP
Note	If step bounding is enabled, the step bound for a variable will be increased if successive changes are in the same direction. More precisely, if there are $XSLP_SAMECOUNT$ successive changes reaching the step bound and in the same direction for a variable, then the step bound (B) for the variable will be reset to $B * XSLP_EXPAND$.

Default value	2
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_SHRINK, XSLP_SHRINKBIAS, XSLP_SAMECOUNT

XSLP_FEASTOLTARGET, SLPFEASTOLTARGET

Description	When set, this defines a target feasibility tolerance to which the linearizations are solved to
Type	Double
Category	Control
Topic areas	SLP, Linearizations, Tolerances
Note	This is a soft version of XPRS_FEASTOL, and will dynamically revert back to XPRS_FEASTOL if the desired accuracy could not be achieved.
Default value	0 (ignored, not set)
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_OPTIMALITYTOLTARGET,

XSLP_GRANULARITY, SLPGRANULARITY

Description	Base for calculating penalty costs
Type	Double
Category	Control
Topic area	SLP
Note	If XSLP_GRANULARITY >1, then initial penalty costs will be powers of XSLP_GRANULARITY.
Default value	4
Affects routines	XSLPconstruct
See also	XSLP_MAXWEIGHT, XSLP_MINWEIGHT

XSLP_INFINITY, NLPINFINITY

Description	Value returned by a divide-by-zero in a formula
Type	Double
Category	Control
Topic area	Numerics
Default value	1.0e+10

XSLP_ITOL_A, SLPITOL_A

Description	Absolute impact convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	The absolute impact convergence criterion assesses the change in the effect of a coefficient in a constraint. The <i>effect</i> of a coefficient is its value multiplied by the activity of the column in which it appears.

$$E = X * C$$

where X is the activity of the matrix column in which the coefficient appears, and C is the value of the coefficient. The linearization approximates the effect of the coefficient as

$$E_1 = X * C_0 + \delta X * C'_0$$

where X is as before, C_0 is the value of the coefficient C calculated using the assumed values for the variables and C'_0 is the value of $\frac{\partial C}{\partial X}$ calculated using the assumed values for the variables.

If C_1 is the value of the coefficient C calculated using the actual values for the variables, then the error in the effect of the coefficient is given by

$$\delta E = X * C_1 - (X * C_0 + \delta X * C'_0)$$

If $\delta E < \text{XSLP_ITOL_A}$

then the variable has passed the absolute impact convergence criterion for this coefficient.

If a variable which has not converged on strict (closure or delta) criteria passes the (relative or absolute) impact or matrix criteria for all the coefficients in which it appears, then it is deemed to have converged. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target `XSLP_VALIDATIONTARGET_R`. Good values for the control are usually fall between $1e-3$ and $1e-6$.

Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ITOL_R, XSLP_MTOL_A, XSLP_MTOL_R, XSLP_STOL_A, XSLP_STOL_R

XSLP_ITOL_R, SLPITOL_R

Description	Relative impact convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	The relative impact convergence criterion assesses the change in the effect of a coefficient in a constraint in relation to the magnitude of the constituents of the constraint. The <i>effect</i> of a coefficient is its value multiplied by the activity of the column in which it appears.

$$E = X * C$$

where X is the activity of the matrix column in which the coefficient appears, and C is the value of the coefficient. The linearization approximates the effect of the coefficient as

$$E_1 = X * C_0 + \delta X * C'_0$$

where X is as before, C_0 is the value of the coefficient C calculated using the assumed values for the variables and C'_0 is the value of $\frac{\partial C}{\partial X}$ calculated using the assumed values for the variables.

If C_1 is the value of the coefficient C calculated using the actual values for the variables, then the error in the effect of the coefficient is given by

$$\delta E = X * C_1 - (X * C_0 + \delta X * C'_0)$$

All the elements of the constraint are examined, excluding delta and error vectors: for each, the contribution to the constraint is evaluated as the element multiplied by the activity of the vector in which it appears; it is then included in a *total positive contribution* or *total negative contribution* depending on the sign of the contribution. If the predicted effect of the coefficient is positive, it is tested against the total positive contribution; if the effect of the coefficient is negative, it is tested against the total negative contribution. If T_0 is the total positive or total negative contribution to the constraint (as appropriate)

and $\delta E < T_0 * XSLP_ITOL_R$

then the variable has passed the relative impact convergence criterion for this coefficient.

If a variable which has not converged on strict (closure or delta) criteria passes the (relative or absolute) impact or matrix criteria for all the coefficients in which it appears, then it is deemed to have converged. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target `XSLP_VALIDATIONTARGET_R`. Good values for the control are usually fall between $1e-3$ and $1e-6$.

Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ITOL_A, XSLP_MTOL_A, XSLP_MTOL_R, XSLP_STOL_A, XSLP_STOL_R

XSLP_MATRIXTOL, SLPMATRIXTOL

Description	Nonzero tolerance for dropping coefficients from the linearization.
Type	Double
Category	Control
Topic area	SLP
Note	Any value smaller than XSLP_MATRIXTOL in magnitude will not be loaded into the linearization. This only applies to the matrix coefficients; bounds, right hand sides and objectives are not affected.
Default value	0.0
Affects routines	XSLPconstruct, XSLPmaxim, XSLPminim

XSLP_MAXWEIGHT, SLPMAXWEIGHT

Description	Maximum penalty weight for delta or error vectors
Type	Double

Category	Control
Topic area	SLP
Note	When penalty vectors are created, or when their weight is increased by escalation, the maximum weight that will be used is given by <code>XSLP_MAXWEIGHT</code> .
Default value	100
Affects routines	<code>XSLPconstruct</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_ALGORITHM</code> , <code>XSLP_AUGMENTATION</code> , <code>XSLP_ESCALATION</code> , <code>XSLP_MINWEIGHT</code>

XSLP_MEMORYFACTOR

Description	Factor for expanding size of dynamic arrays in memory
Type	Double
Category	Control
Topic area	Memory
Note	When a dynamic array has to be increased in size, the new space allocated will be <code>XSLP_MEMORYFACTOR</code> times as big as the previous size. A larger value may result in improved performance because arrays need to be re-sized and moved less frequently; however, more memory may be required under such circumstances because not all of the previous memory area can be re-used efficiently.
Default value	1.6
See also	Memory control variables <code>XSLP_MEM*</code> Memory control variables <code>XSLP_MEM*</code>

XSLP_MERITLAMBDA, NLPMERITLAMBDA

Description	Factor by which the net objective is taken into account in the merit function
Type	Double
Category	Control
Topic areas	SLP, Solution
Note	The merit function is evaluated in the original, non-augmented / linearized space of the problem. A solution is deemed improved, if either feasibility improved, or if feasibility is not deteriorated but the net objective is improved, or if the combination of the two is improved, where the value of the <code>XSLP_MERITLAMBDA</code> control is used to combine the two measures. A nonpositive value indicates that the combined effect should not be checked.
Default value	0.0
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_FILTER</code> <code>XSLP_LSITERLIMIT</code> <code>XSLP_LSPATTERNLIMIT</code>

XSLP_MINSBFACTOR, SLPMINSBFACTOR

Description	Factor by which step bounds can be decreased beneath XSLP_ATOL_A
Type	Double
Category	Control
Topic area	SLP
Note	Normally, step bounds are not decreased beneath XSLP_ATOL_A, as such variables are treated as converged. However, it may be beneficial to decrease step bounds further, as individual variable value changes might affect the convergence of other variables in the model, even if the variable itself is deemed converged.
Default value	1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ATOL_A

XSLP_MINWEIGHT, SLPMINWEIGHT

Description	Minimum penalty weight for delta or error vectors
Type	Double
Category	Control
Topic area	SLP
Note	When penalty vectors are created, the minimum weight that will be used is given by XSLP_MINWEIGHT.
Default value	0.01
Affects routines	XSLPconstruct, XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_MAXWEIGHT

XSLP_MIPCUTOFF_A, SLPMIPCUTOFF_A

Description	Absolute objective function cutoff for MIP termination
Type	Double
Category	Control
Topic areas	MISLP, Tolerances
Note	<p>If the objective function is worse by a defined amount than the best integer solution obtained so far, then the SLP will be terminated (and the node will be cut off). The node will be cut off at the current SLP iteration if the objective function for the last XSLP_MIPCUTOFFCOUNT SLP iterations are all worse than the best obtained so far, and the difference is greater than XSLP_MIPCUTOFF_A and $OBJ * XSLP_MIPCUTOFF_R$ where OBJ is the best integer solution obtained so far.</p> <p>The MIP cutoff tests are only applied after XSLP_MIPCUTOFFLIMIT SLP iterations at the current node.</p>

Default value	0.00001
Affects routines	XSLPnlpoptimize
See also	XSLP_MIPCUTOFFCOUNT, XSLP_MIPCUTOFFLIMIT, XSLP_MIPCUTOFF_R

XSLP_MIPCUTOFF_R, SLPMIPCUTOFF_R

Description	Absolute objective function cutoff for MIP termination
Type	Double
Category	Control
Topic areas	MISLP, Tolerances
Note	<p>If the objective function is worse by a defined amount than the best integer solution obtained so far, then the SLP will be terminated (and the node will be cut off). The node will be cut off at the current SLP iteration if the objective function for the last XSLP_MIPCUTOFFCOUNT SLP iterations are all worse than the best obtained so far, and the difference is greater than XSLP_MIPCUTOFF_A and $OBJ * XSLP_MIPCUTOFF_R$ where OBJ is the best integer solution obtained so far.</p> <p>The MIP cutoff tests are only applied after XSLP_MIPCUTOFFLIMIT SLP iterations at the current node.</p>
Default value	0.00001
Affects routines	XSLPnlpoptimize
See also	XSLP_MIPCUTOFFCOUNT, XSLP_MIPCUTOFFLIMIT, XSLP_MIPCUTOFF_A

XSLP_MIPERRORTOL_A, SLPMIPERRORTOL_A

Description	Absolute penalty error cost tolerance for MIP cut-off
Type	Double
Category	Control
Topic areas	MISLP, Tolerances
Note	<p>The penalty error cost test is applied at each node where there are active penalties in the solution. If XSLP_MIPERRORTOL_A is nonzero and the absolute value of the penalty costs is greater than XSLP_MIPERRORTOL_A, the node will be declared infeasible. If XSLP_MIPERRORTOL_A is zero then no test is made and the node will not be declared infeasible on this criterion.</p>
Default value	0 (inactive)
Affects routines	XSLPnlpoptimize
See also	XSLP_MIPERRORTOL_R

XSLP_MIPERRORTOL_R, SLPMIPERRORTOL_R

Description	Relative penalty error cost tolerance for MIP cut-off
Type	Double
Category	Control
Topic areas	MISLP, Tolerances
Note	The penalty error cost test is applied at each node where there are active penalties in the solution. If XSLP_MIPERRORTOL_R is nonzero and the absolute value of the penalty costs is greater than $\text{XSLP_MIPERRORTOL_R} * \text{abs}(\text{Obj})$ where Obj is the value of the objective function, then the node will be declared infeasible. If XSLP_MIPERRORTOL_R is zero then no test is made and the node will not be declared infeasible on this criterion.
Default value	0 (inactive)
Affects routines	XSLPnlpoptimize
See also	XSLP_MIPERRORTOL_A

XSLP_MIPOTOL_A, SLPMIPOTOL_A

Description	Absolute objective function tolerance for MIP termination
Type	Double
Category	Control
Topic areas	MISLP, SLP-convergence, Tolerances
Note	The objective function test for MIP termination is applied only when step bounding has been applied (or XSLP_SBSTART SLP iterations have taken place if step bounding is not being used). The node will be terminated at the current SLP iteration if the range of the objective function values over the last XSLP_MIPOCOUNT SLP iterations is within XSLP_MIPOTOL_A or within $\text{OBJ} * \text{XSLP_MIPOTOL_R}$ where OBJ is the average value of the objective function over those iterations.
Default value	0.00001
Affects routines	XSLPnlpoptimize
See also	XSLP_MIPOCOUNT XSLP_MIPOTOL_R XSLP_SBSTART

XSLP_MIPOTOL_R, SLPMIPOTOL_R

Description	Relative objective function tolerance for MIP termination
Type	Double
Category	Control
Topic areas	MISLP, SLP-convergence, Tolerances

Note	The objective function test for MIP termination is applied only when step bounding has been applied (or XSLP_SBSTART SLP iterations have taken place if step bounding is not being used). The node will be terminated at the current SLP iteration if the range of the objective function values over the last XSLP_MIPOCOUNT SLP iterations is within XSLP_MIPOTOL_A or within $OBJ * XSLP_MIPOTOL_R$ where OBJ is the average value of the objective function over those iterations.
Default value	0.00001
Affects routines	XSLPnlpoptimize
See also	XSLP_MIPOCOUNT XSLP_MIPOTOL_A XSLP_SBSTART

XSLP_MSMBOUNDRANGE, MSMBOUNDRANGE

Description	Defines the maximum range inside which initial points are generated by multistart presets
Type	Double
Category	Control
Topic area	Multistart
Note	The is the maximum range in which initial points are generated; the actual range is expected to be smaller as bounds are domains are also considered.
Default value	1000
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART

XSLP_MTOL_A, SLPMTOL_A

Description	Absolute effective matrix element convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	The absolute effective matrix element convergence criterion assesses the change in the effect of a coefficient in a constraint. The <i>effect</i> of a coefficient is its value multiplied by the activity of the column in which it appears.

$$E = X * C$$

where X is the activity of the matrix column in which the coefficient appears, and C is the value of the coefficient. The linearization approximates the effect of the coefficient as

$$E = X * C_0 + \delta X * C'_0$$

where V is as before, C_0 is the value of the coefficient C calculated using the assumed values for the variables and C'_0 is the value of $\frac{\partial C}{\partial X}$ calculated using the assumed values for the variables.

If C_1 is the value of the coefficient C calculated using the actual values for the variables, then the error in the effect of the coefficient is given by

$$\delta E = X * C_1 - (X * C_0 + \delta X * C'_0)$$

If $\delta E < X * XSLP_MTOL_A$

then the variable has passed the absolute effective matrix element convergence criterion for this coefficient.

If a variable which has not converged on strict (closure or delta) criteria passes the (relative or absolute) impact or matrix criteria for all the coefficients in which it appears, then it is deemed to have converged. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target `XSLP_VALIDATIONTARGET_R`. Good values for the control are usually fall between $1e-3$ and $1e-6$.

Default value -1.0

Affects routines XSLPmaxim, XSLPminim

See also XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_R, XSLP_STOL_A, XSLP_STOL_R

XSLP_MTOL_R, SLPMTOL_R

Description Relative effective matrix element convergence tolerance

Type Double

Category Control

Topic areas SLP, SLP-convergence, Tolerances

Note The relative effective matrix element convergence criterion assesses the change in the effect of a coefficient in a constraint relative to the magnitude of the coefficient. The *effect* of a coefficient is its value multiplied by the activity of the column in which it appears.

$$E = X * C$$

where X is the activity of the matrix column in which the coefficient appears, and C is the value of the coefficient. The linearization approximates the effect of the coefficient as

$$E_1 = X * C_0 + \delta X * C'_0$$

where V is as before, C_0 is the value of the coefficient C calculated using the assumed values for the variables and C'_0 is the value of $\frac{\partial C}{\partial X}$ calculated using the assumed values for the variables.

If C_1 is the value of the coefficient C calculated using the actual values for the variables, then the error in the effect of the coefficient is given by

$$\delta E = X * C_1 - (X * C_0 + \delta X * C'_0)$$

If $\delta E < E_1 * XSLP_MTOL_R$

then the variable has passed the relative effective matrix element convergence criterion for this coefficient.

If a variable which has not converged on strict (closure or delta) criteria passes the (relative or absolute) impact or matrix criteria for all the coefficients in which it appears, then it is deemed to have converged. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target `XSLP_VALIDATIONTARGET_R`. Good values for the control are usually fall between $1e-3$ and $1e-6$.

Default value -1.0

Affects routines XSLPmaxim, XSLPminim

See also XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_A, XSLP_STOL_A, XSLP_STOL_R

XSLP_MVTOL, SLPMVTOL

Description	Marginal value tolerance for determining if a constraint is slack
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	If the absolute value of the marginal value of a constraint is less than XSLP_MVTOL, then (1) the constraint is regarded as not constraining for the purposes of the slack tolerance convergence criteria; (2) the constraint is not regarded as an <i>active constraint</i> when identifying unconverged variables in active constraints. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target XSLP_VALIDATIONTARGET_R. Good values for the control are usually fall between 1e-3 and 1e-6.
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_STOL_A, XSLP_STOL_R

XSLP_OBJSENSE

Description	<i>This parameter is deprecated and will be removed in a future release. Please use XPRSchgobjsense instead.</i> Objective function sense.
Type	Double
Category	Control
Note	XSLP_OBJSENSE is set to +1 for minimization and to -1 for maximization. It is automatically set by XSLPmaxim and XSLPminim; it must be set by the user before calling XSLPnloptimize.
Set by routines	XSLPmaxim, XSLPminim
Default value	+1
Affects routines	XSLPmaxim, XSLPminim, XSLPnloptimize

XSLP_OBJTHRESHOLD, SLPOBJTHRESHOLD

Description	Assumed maximum value of the objective function in absolute value.
Type	Double
Category	Control
Topic area	SLP
Default value	1.0e+15
Affects routines	XSLPmaxim, XSLPminim

XSLP_OBJTOPENALTYCOST, SLPOBJTOPENALTYCOST

Description	Factor to estimate initial penalty costs from objective function
Type	Double
Category	Control
Topic area	SLP
Notes	<p>The setting of initial penalty error costs can affect the path of the optimization and, indeed, whether a solution is achieved at all. If the penalty costs are too low, then unbounded solutions may result although Xpress-SLP will increase the costs in an attempt to recover. If the penalty costs are too high, then the requirement to achieve feasibility of the linearized constraints may be too strong to allow the system to explore the nonlinear feasible region. Low penalty costs can result in many SLP iterations, as feasibility of the nonlinear constraints is not achieved until the penalty costs become high enough; high penalty costs force feasibility of the linearizations, and so tend to find local optima close to an initial feasible point. Xpress-SLP can analyze the problem to estimate the size of penalty costs required to avoid an initial unbounded solution. XSLP_OBJTOPENALTYCOST can be used in conjunction with this procedure to scale the costs and give an appropriate initial value for balancing the requirements of feasibility and optimality.</p> <p>Not all models are amenable to the Xpress-SLP analysis. As the analysis is initially concerned with establishing a cost level to avoid unboundedness, a model which is sufficiently constrained will never show unboundedness regardless of the cost. Also, as the analysis is done at the start of the optimization to establish a penalty cost, significant changes in the coefficients, or a high degree of nonlinearity, may invalidate the initial analysis.</p> <p>A setting for XSLP_OBJTOPENALTYCOST of zero disables the analysis. A setting of 3 or 4 has proved successful for many models. If XSLP_OBJTOPENALTYCOST cannot be used because of the problem structure, its effect can still be emulated by some initial experiments to establish the cost required to avoid unboundedness, and then manually applying a suitable factor. If the problem is initially unbounded, then the penalty cost will be increased until either it reaches its maximum or the problem becomes bounded.</p>
Default value	0
Affects routines	XSLPmaxim, XSLPminim

XSLP_OPTIMALITYTOLTARGET, SLPOPTIMALITYTOLTARGET

Description	When set, this defines a target optimality tolerance to which the linearizations are solved to
Type	Double
Category	Control
Topic areas	SLP, Linearizations, Tolerances
Note	This is a soft version of XPRS_OPTIMALITYTOL, and will dynamically revert back to XPRS_OPTIMALITYTOL if the desired accuracy could not be achieved.
Default value	0 (ignored, not set)
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_FEASTOLTARGET,

XSLP_OTOL_A, SLPOTOL_A

Description	Absolute static objective (2) convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	<p>The static objective (2) convergence criterion does not measure convergence of individual variables. Instead, it measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables interacting with active constraints (those that have a marginal value of at least XSLP_MVTOL) have converged. The rationale is that if the remaining unconverged variables are not involved in active constraints and if the objective function is not changing significantly between iterations, then the solution is more-or-less practical.</p> <p>The variation in the objective function is defined as</p> $\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$ <p>where Iter is the XSLP_OCOUNT most recent SLP iterations and Obj is the corresponding objective function value.</p> <p>If $\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_OTOL_A}$ then the problem has converged on the absolute static objective (2) convergence criterion. The static objective function (2) test is applied only if XSLP_OCOUNT is at least 2. When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target XSLP_VALIDATIONTARGET_K. Good values for the control are usually fall between $1\text{e-}3$ and $1\text{e-}6$.</p>
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_OCOUNT, XSLP_OTOL_R

XSLP_OTOL_R, SLPOTOL_R

Description	Relative static objective (2) convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	<p>The static objective (2) convergence criterion does not measure convergence of individual variables. Instead, it measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables interacting with active constraints (those that have a marginal value of at least XSLP_MVTOL) have converged. The rationale is that if the remaining unconverged variables are not involved in active constraints and if the objective function is not changing significantly between iterations, then the solution is more-or-less practical.</p> <p>The variation in the objective function is defined as</p>

$$\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$$

where $Iter$ is the `XSLP_OCOUNT` most recent SLP iterations and Obj is the corresponding objective function value.

If $ABS(\delta Obj) \leq AVG_{Iter}(Obj) * XSLP_OTOL_R$

then the problem has converged on the relative static objective (2) convergence criterion.

The static objective function (2) test is applied only if `XSLP_OCOUNT` is at least 2. When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target `XSLP_VALIDATIONTARGET_K`. Good values for the control are usually fall between $1e-3$ and $1e-6$.

Default value -1.0

Affects routines `XSLPmaxim`, `XSLPminim`

XSLP_PRESOLVE_ELIMTOL, NLPPRESOLVE_ELIMTOL

Description Tolerance for nonlinear eliminations during SLP presolve

Type Double

Category Control

Topic areas Presolve, Tolerances

Note Any eliminations on smaller coefficients will be rejected.

Default value 0.001

Affects routines `XSLPpresolve`

XSLP_PRESOLVEZERO, NLPPRESOLVEZERO

Description Minimum absolute value for a variable which is identified as nonzero during SLP presolve

Type Double

Category Control

Topic areas Presolve, Tolerances

Note During the SLP (nonlinear)presolve, a variable may be identified as being nonzero (for example, because it is used as a divisor). A bound of plus or minus `XSLP_PRESOLVEZERO` will be applied to the variable if it is identified as non-negative or non-positive.

Default value 1.0E-09

Affects routines `XSLPpresolve`

XSLP_PRIMALINTEGRALALPHA, NLPPRIMALINTEGRALALPHA

Description Decay term for primal integral computation

Type Double

Category Control

Topic area Logging

Note This control represents the exponential decay term for computing the XSLP_PRIMALINTEGRAL. 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

XSLP_PRIMALINTEGRALREF, NLPPRIMALINTEGRALREF

Description Reference solution value to take into account when calculating the primal integral

Type Double

Category Control

Topic area Logging

Note When a global optimum is known, this can used to calculate a globally valid primal integral. It can also be used to indicate the target objective value still to be taken into account in the integral.

Default value XPRS_PLUSINFINITY

Affects routines XSLPminim, XSLPmaxim

XSLP_SHRINK, SLPSHRINK

Description Multiplier to reduce a step bound

Type Double

Category Control

Topic area SLP

Note If step bounding is enabled, the step bound for a variable will be decreased if successive changes are in opposite directions. The step bound (B) for the variable will be reset to $B * XSLP_SHRINK$.
If the step bound is already below the strict (delta or closure) tolerances, it will not be reduced further.

Default value 0.5

Affects routines XSLPmaxim, XSLPminim

See also XSLP_EXPAND, XSLP_SHRINKBIAS, XSLP_SAMECOUNT

XSLP_SHRINKBIAS, SLPSHRINKBIAS

Description Defines an overwrite / adjustment of step bounds for improving iterations

Type Double

Category Control

Topic area SLP

Note	Positive values overwrite XSLP_SHRINK only if the objective is improving. A negative value is used to scale all step bounds in improving iterations.
Default value	0 (ignored, not set)
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_SHRINK, XSLP_EXPAND, XSLP_SAMECOUNT

XSLP_STOL_A, SLPSTOL_A

Description	Absolute slack convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	The slack convergence criterion is identical to the impact convergence criterion, except that the tolerances used are XSLP_STOL_A (instead of XSLP_ITOL_A) and XSLP_STOL_R (instead of XSLP_ITOL_R). See XSLP_ITOL_A for a description of the test. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target XSLP_VALIDATIONTARGET_R. Good values for the control are usually fall between 1e-3 and 1e-6.
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_A, XSLP_MTOL_R, XSLP_STOL_R

XSLP_STOL_R, SLPSTOL_R

Description	Relative slack convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	The slack convergence criterion is identical to the impact convergence criterion, except that the tolerances used are XSLP_STOL_A (instead of XSLP_ITOL_A) and XSLP_STOL_R (instead of XSLP_ITOL_R). See XSLP_ITOL_R for a description of the test. When the value is set to be negative, the value is adjusted automatically by SLP, based on the feasibility target XSLP_VALIDATIONTARGET_R. Good values for the control are usually fall between 1e-3 and 1e-6.
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_A, XSLP_MTOL_R, XSLP_STOL_A

XSLP_VALIDATIONFACTOR, NLPVALIDATIONFACTOR

Description	Minimum improvement in validation targets to continue iterating
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Default value	0.001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_VALIDATIONTARGET_K XSLP_VALIDATIONTARGET_R

XSLP_VALIDATIONTARGET_R, NLPVALIDATIONTARGET_R

Description	Feasibility target tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	Primary feasibility control for SLP. When the relevant feasibility based convergence controls are left at their default values, SLP will adjust their value to match the target. The control defines a target value, that may not necessarily be attainable.
Default value	1e-6
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_VALIDATIONTARGET_K

XSLP_VALIDATIONTARGET_K, NLPVALIDATIONTARGET_K

Description	Optimality target tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	Primary optimality control for SLP. When the relevant optimality based convergence controls are left at their default values, SLP will adjust their value to match the target. The control defines a target value, that may not necessarily be attainable for problem with no strong constraint qualifications.
Default value	1e-6
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_VALIDATIONTARGET_R

XSLP_VALIDATIONTOL_A, NLPVALIDATIONTOL_A

Description	Absolute tolerance for the XSLPvalidate procedure
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Note	<p>XSLPvalidate checks the feasibility of a converged solution against relative and absolute tolerances for each constraint. The left hand side and the right hand side of the constraint are calculated using the converged solution values. If the calculated values imply that the constraint is infeasible, then the difference (D) is tested against the absolute and relative validation tolerances.</p> <p>If $D < \text{XSLP_VALIDATIONTOL_A}$ then the constraint is within the absolute validation tolerance. The total positive (TPos) and negative contributions (TNeg) to the left hand side are also calculated.</p> <p>If $D < \text{MAX}(\text{ABS}(\text{TPos}), \text{ABS}(\text{TNeg})) * \text{XSLP_VALIDATIONTOL_A}$ then the constraint is within the relative validation tolerance. For each constraint which is outside both the absolute and relative validation tolerances, validation factors are calculated which are the factors by which the infeasibility exceeds the corresponding validation tolerance; the smaller factor is printed in the validation report.</p> <p>The validation index XSLP_VALIDATIONINDEX_A is the largest of these factors which is an absolute validation factor multiplied by the absolute validation tolerance; the validation index XSLP_VALIDATIONINDEX_R is the largest of these factors which is a relative validation factor multiplied by the relative validation tolerance.</p>
Default value	0.00001
Affects routines	XSLPvalidate
See also	XSLP_VALIDATIONINDEX_A, XSLP_VALIDATIONINDEX_R, XSLP_VALIDATIONTOL_R

XSLP_VALIDATIONTOL_K, NLPVALIDATIONTOL_K

Description	Relative tolerance for the XSLPvalidatekkt procedure
Type	Double
Category	Control
Topic areas	SLP, Tolerances
Default value	0.00001

XSLP_VALIDATIONTOL_R, NLPVALIDATIONTOL_R

Description	Relative tolerance for the XSLPvalidate procedure
Type	Double
Category	Control
Topic area	Tolerances

Note	<p>XSLPvalidate checks the feasibility of a converged solution against relative and absolute tolerances for each constraint. The left hand side and the right hand side of the constraint are calculated using the converged solution values. If the calculated values imply that the constraint is infeasible, then the difference (D) is tested against the absolute and relative validation tolerances.</p> <p>If $D < \text{XSLP_VALIDATIONTOL_A}$</p> <p>then the constraint is within the absolute validation tolerance. The total positive (TPos) and negative contributions (TNeg) to the left hand side are also calculated.</p> <p>If $D < \text{MAX}(\text{ABS}(\text{TPos}), \text{ABS}(\text{TNeg})) * \text{XSLP_VALIDATIONTOL_R}$</p> <p>then the constraint is within the relative validation tolerance. For each constraint which is outside both the absolute and relative validation tolerances, validation factors are calculated which are the factors by which the infeasibility exceeds the corresponding validation tolerance; the smaller factor is printed in the validation report.</p> <p>The validation index $\text{XSLP_VALIDATIONINDEX_A}$ is the largest of these factors which is an absolute validation factor multiplied by the absolute validation tolerance; the validation index $\text{XSLP_VALIDATIONINDEX_R}$ is the largest of these factors which is a relative validation factor multiplied by the relative validation tolerance.</p>
Default value	0.00001
Affects routines	XSLPvalidate
See also	XSLP_VALIDATIONINDEX_A, XSLP_VALIDATIONINDEX_R, XSLP_VALIDATIONTOL_A

XSLP_VTOL_A, SLPVTOL_A

Description	Absolute static objective (3) convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	<p>The static objective (3) convergence criterion does not measure convergence of individual variables, and in fact does not in any way imply that the solution has converged. However, it is sometimes useful to be able to terminate an optimization once the objective function appears to have stabilized. One example is where a set of possible schedules are being evaluated and initially only a good estimate of the likely objective function value is required, to eliminate the worst candidates.</p> <p>The variation in the objective function is defined as</p> $\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$ <p>where Iter is the XSLP_VCOUNT most recent SLP iterations and Obj is the corresponding objective function value.</p> <p>If $\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_VTOL_A}$</p> <p>then the problem has converged on the absolute static objective function (3) criterion.</p> <p>The static objective function (3) test is applied only if after at least $\text{XSLP_VLIMIT} + \text{XSLP_SBSTART}$ SLP iterations have taken place and only if XSLP_VCOUNT is at least 2. Where step bounding is being used, this ensures that the test is not applied until after step bounding has been introduced. When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target $\text{XSLP_VALIDATIONTARGET_K}$. Good values for the control are usually fall between $1\text{e-}3$ and $1\text{e-}6$.</p>
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_SBSTART, XSLP_VCOUNT, XSLP_VLIMIT, XSLP_VTOL_R

XSLP_VTOL_R, SLPVTOL_R

Description	Relative static objective (3) convergence tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	<p>The static objective (3) convergence criterion does not measure convergence of individual variables, and in fact does not in any way imply that the solution has converged. However, it is sometimes useful to be able to terminate an optimization once the objective function appears to have stabilized. One example is where a set of possible schedules are being evaluated and initially only a good estimate of the likely objective function value is required, to eliminate the worst candidates.</p> <p>The variation in the objective function is defined as</p> $\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$ <p>where <i>Iter</i> is the XSLP_VCOUNT most recent SLP iterations and <i>Obj</i> is the corresponding objective function value.</p> <p>If $\text{ABS}(\delta\text{Obj}) \leq \text{AVG}_{\text{Iter}}(\text{Obj}) * \text{XSLP_VTOL_R}$ then the problem has converged on the absolute static objective function (3) criterion. The static objective function (3) test is applied only if after at least XSLP_VLIMIT + XSLP_SBSTART SLP iterations have taken place and only if XSLP_VCOUNT is at least 2. Where step bounding is being used, this ensures that the test is not applied until after step bounding has been introduced. When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target XSLP_VALIDATIONTARGET_K. Good values for the control are usually fall between 1e-3 and 1e-6.</p>
Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_SBSTART, XSLP_VCOUNT, XSLP_VLIMIT, XSLP_VTOL_A

XSLP_WTOL_A, SLPWTOL_A

Description	Absolute extended convergence continuation tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	<p>It may happen that all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. This means that, at least in the linearization, if the variable were to be allowed to move further the objective function would improve. This does not necessarily imply that the same is true of the original problem, but it is still possible that an improved result could be obtained by taking another SLP iteration.</p> <p>The extended convergence continuation criterion is applied after a converged solution has been found where at least one variable has converged on extended criteria and is at its step bound limit. The extended convergence continuation test measures whether any improvement is being achieved when additional SLP iterations are carried out. If not, then the last converged solution will be restored and the optimization will stop.</p>

For a maximization problem, the improvement in the objective function at the current iteration compared to the objective function at the last converged solution is given by:

$$\delta\text{Obj} = \text{Obj} - \text{LastConvergedObj}$$

For a minimization problem, the sign is reversed.

If $\delta\text{Obj} > \text{XSLP_WTOL_A}$ and

$\delta\text{Obj} > \text{ABS}(\text{ConvergedObj}) * \text{XSLP_WTOL_R}$ then the solution is deemed to have a significantly better objective function value than the converged solution.

When a solution is found which converges on extended criteria and with active step bounds, the solution is saved and SLP optimization continues until one of the following:

- (1) a new solution is found which converges on some other criterion, in which case the SLP optimization stops with this new solution;
- (2) a new solution is found which converges on extended criteria and with active step bounds, and which has a significantly better objective function, in which case this is taken as the new saved solution;
- (3) none of the XSLP_WCOUNT most recent SLP iterations has a significantly better objective function than the saved solution, in which case the saved solution is restored and the SLP optimization stops.

When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target $\text{XSLP_VALIDATIONTARGET_K}$. Good values for the control are usually fall between $1\text{e-}3$ and $1\text{e-}6$.

Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_WCOUNT, XSLP_WTOL_R

XSLP_WTOL_R, SLPWTOL_R

Description Relative extended convergence continuation tolerance

Type Double

Category Control

Topic areas SLP, SLP-convergence, Tolerances

Note It may happen that all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. This means that, at least in the linearization, if the variable were to be allowed to move further the objective function would improve. This does not necessarily imply that the same is true of the original problem, but it is still possible that an improved result could be obtained by taking another SLP iteration.

The extended convergence continuation criterion is applied after a converged solution has been found where at least one variable has converged on extended criteria and is at its step bound limit. The extended convergence continuation test measures whether any improvement is being achieved when additional SLP iterations are carried out. If not, then the last converged solution will be restored and the optimization will stop.

For a maximization problem, the improvement in the objective function at the current iteration compared to the objective function at the last converged solution is given by:

$$\delta\text{Obj} = \text{Obj} - \text{LastConvergedObj}$$

For a minimization problem, the sign is reversed.

If $\delta\text{Obj} > \text{XSLP_WTOL_A}$ and

$$\delta\text{Obj} > \text{ABS}(\text{ConvergedObj}) * \text{XSLP_WTOL_R}$$

then the solution is deemed to have a significantly better objective function value than the converged solution.

If XSLP_WCOUNT is greater than zero, and a solution is found which converges on extended criteria and with active step bounds, the solution is saved and SLP optimization continues until one of the following:

- (1) a new solution is found which converges on some other criterion, in which case the SLP optimization stops with this new solution;
- (2) a new solution is found which converges on extended criteria and with active step bounds, and which has a significantly better objective function, in which case this is taken as the new saved solution;
- (3) none of the XSLP_WCOUNT most recent SLP iterations has a significantly better objective function than the saved solution, in which case the saved solution is restored and the SLP optimization stops.

When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target XSLP_VALIDATIONTARGET_K. Good values for the control are usually fall between 1e-4 and 1e-6.

Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_WCOUNT, XSLP_WTOL_A

XSLP_XTOL_A, SLPXTOL_A

Description	Absolute static objective function (1) tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances

Note It may happen that all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. This means that, at least in the linearization, if the variable were to be allowed to move further the objective function would improve. This does not necessarily imply that the same is true of the original problem, but it is still possible that an improved result could be obtained by taking another SLP iteration. However, if the objective function has already been stable for several SLP iterations, then there is less likelihood of an improved result, and the converged solution can be accepted.

The static objective function (1) test measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. Because all the variables have converged, the solution is already converged but the fact that some variables are at their step bound limit suggests that the objective function could be improved by going further.

The variation in the objective function is defined as

$$\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$$

where Iter is the XSLP_XCOUNT most recent SLP iterations and Obj is the corresponding objective function value.

If $\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_XTOL_A}$

then the objective function is deemed to be static according to the absolute static objective function (1) criterion.

If $\text{ABS}(\delta\text{Obj}) \leq \text{AVG}_{\text{Iter}}(\text{Obj}) * \text{XSLP_XTOL_R}$

then the objective function is deemed to be static according to the relative static objective function (1) criterion.

The static objective function (1) test is applied only until `XSLP_XLIMIT` SLP iterations have taken place. After that, if all the variables have converged on strict or extended criteria, the solution is deemed to have converged.

If the objective function passes the relative or absolute static objective function (1) test then the solution is deemed to have converged.

When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target `XSLP_VALIDATIONTARGET_K`. Good values for the control are usually fall between $1e-3$ and $1e-6$.

Default value	-1.0
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_XCOUNT</code> , <code>XSLP_XLIMIT</code> , <code>XSLP_XTOL_R</code>

XSLP_XTOL_R, SLPXTOL_R

Description	Relative static objective function (1) tolerance
Type	Double
Category	Control
Topic areas	SLP, SLP-convergence, Tolerances
Note	It may happen that all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. This means that, at least in the linearization, if the variable were to be allowed to move further the objective function would improve. This does not necessarily imply that the same is true of the original problem, but it is still possible that an improved result could be obtained by taking another SLP iteration. However, if the objective function has already been stable for several SLP iterations, then there is less likelihood of an improved result, and the converged solution can be accepted.

The static objective function (1) test measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. Because all the variables have converged, the solution is already converged but the fact that some variables are at their step bound limit suggests that the objective function could be improved by going further.

The variation in the objective function is defined as

$$\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$$

where `Iter` is the `XSLP_XCOUNT` most recent SLP iterations and `Obj` is the corresponding objective function value.

If $\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_XTOL_A}$

then the objective function is deemed to be static according to the absolute static objective function (1) criterion.

If $\text{ABS}(\delta\text{Obj}) \leq \text{AVG}_{\text{Iter}}(\text{Obj}) * \text{XSLP_XTOL_R}$

then the objective function is deemed to be static according to the relative static objective function (1) criterion.

The static objective function (1) test is applied only until `XSLP_XLIMIT` SLP iterations have taken place. After that, if all the variables have converged on strict or extended criteria, the solution is deemed to have converged.

If the objective function passes the relative or absolute static objective function (1) test then the solution is deemed to have converged.

When the value is set to be negative, the value is adjusted automatically by SLP, based on the optimality target `XSLP_VALIDATIONTARGET_K`. Good values for the control are usually fall between $1e-4$ and $1e-6$.

Default value	-1.0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_XCOUNT, XSLP_XLIMIT, XSLP_XTOL_A

XSLP_ZERO, NLPZERO

Description	Absolute tolerance
Type	Double
Category	Control
Topic area	Tolerances
Note	<p>If a value is below <code>XSLP_ZERO</code> in magnitude, then it will be regarded as zero in certain formula calculations:</p> <ul style="list-style-type: none"> an attempt to divide by such a value will give a "divide by zero" error; an exponent of a negative number will produce a "negative number, fractional exponent" error if the exponent differs from an integer by more than <code>XSLP_ZERO</code>.
Default value	1.0E-15
Affects routines	XSLPevaluatecoef, XSLPevaluateformula

20.2 Integer control parameters

XSLP_ALGORITHM, SLPALGORITHM

Description	Bit map describing the SLP algorithm(s) to be used	
Type	Integer	
Category	Control	
Topic areas	SLP, Bit-vector	
Values	Bit	Meaning
	0	Do not apply step bounds.
	1	Apply step bounds to SLP delta vectors only when required.
	2	Estimate step bounds from early SLP iterations.
	3	Use dynamic damping.
	4	Do not update values which are converged within strict tolerance.
	5	Retain previous value when cascading if determining row is zero.
	6	Reset XSLP_DELTA_Z to zero when converged and continue SLP.
	7	Quick convergence check.
	8	Escalate penalties.
	9	Use the primal simplex algorithm when all error vectors become inactive.
	11	Continue optimizing after penalty cost reaches maximum.
	12	Accept a solution which has converged even if there are still significant active penalty error vectors.
	13	Skip the solution polishing step if the LP postsolve returns a slightly infeasible, but claimed optimal solution.
	14	Step bounds are updated to accomodate cascaded values (otherwise cascaded values are pushed to respect step bounds).
	15	Apply clamping when converged on extended criteria only with some variables having active step bounds.
	16	Apply clamping when converged on extended criteria only.
Notes	<p>Bit 0: Do not apply step bounds. The default algorithm uses step bounds to force convergence. Step bounds may not be appropriate if dynamic damping is used.</p> <p>Bit 1: Apply step bounds to SLP delta vectors only when required. Step bounds can be applied to all vectors simultaneously, or applied only when oscillation of the delta vector (change in sign between successive SLP iterations) is detected.</p> <p>Bit 2: Estimate step bounds from early SLP iterations. If initial step bounds are not being explicitly provided, this gives a good method of calculating reasonable values. Values will tend to be larger rather than smaller, to reduce the risk of infeasibility caused by excessive tightness of the step bounds.</p> <p>Bit 3: Use dynamic damping. Dynamic damping is sometimes an alternative to step bounding as a means of encouraging convergence, but it does not have the same power to force convergence as do step bounds.</p> <p>Bit 4: Do not update values which are converged within strict tolerance. Models which are numerically unstable may benefit from this setting, which does not update values which have effectively hardly changed. If a variable subsequently does move outside its strict convergence tolerance, it will be updated as usual.</p>	

Bit 5: Retain previous value when cascading if determining row is zero. If the determining row is zero (that is, all the coefficients interacting with it are either zero or in columns with a zero activity), then it is impossible to calculate a new value for the vector being cascaded. The choice is to use the solution value as it is, or to revert to the assumed value

Bit 6: Reset XSLP_DELTA_Z to zero when converged and continue SLP. One of the mechanisms to avoid local optima is to retain small non-zero coefficients between delta vectors and constraints, even when the coefficient should strictly be zero. If this option is set, then a converged solution will be continued with zero coefficients as appropriate.

Bit 7: Quick convergence check. Normally, each variable is checked against all convergence criteria until either a criterion is found which it passes, or it is declared "not converged". Later (extended convergence) criteria are more expensive to test and, once an unconverged variable has been found, the overall convergence status of the solution has been established. The quick convergence check carries out checks on the strict criteria, but omits checks on the extended criteria when an unconverged variable has been found.

Bit 8: Escalate penalties. Constraint penalties are increased after each SLP iteration where penalty vectors are present in the solution. Escalation applies an additional scaling factor to the penalty costs for active errors. This helps to prevent successive solutions becoming "stuck" because of a particular constraint, because its cost will be raised so that other constraints may become more attractive to violate instead and thus open up a new region to explore.

Bit 9: Use the primal simplex algorithm when all error vectors become inactive. The primal simplex algorithm often performs better than dual during the final stages of SLP optimization when there are relatively few basis changes between successive solutions. As it is impossible to establish in advance when the final stages are being reached, the disappearance of error vectors from the solution is used as a proxy.

Bit 11: Continue optimizing after penalty cost reaches maximum. Normally if the penalty cost reaches its maximum (by default the value of XPRS_PLUSINFINITY), the optimization will terminate with an unconverged solution. If the maximum value is set to a smaller value, then it may make sense to continue, using other means to determine when to stop.

Bit 12: Accept a solution which has converged even if there are still significant active penalty error vectors. Normally, the optimization will continue if there are active penalty vectors in the solution. However, it may be that there is no feasible solution (and so active penalties will always be present). Setting bit 12 means that, if other convergence criteria are met, then the solution will be accepted as converged and the optimization will stop.

Bit 13: Due to the nature of the SLP linearizations, and in particular because of the large differences in the objective function (model objective against penalty costs) some dual reductions in the linear presolver might introduce numerically instable reductions that cause slight infeasibilities to appear in postsolve. It is typically more efficient to remove these infeasibilities with an extra call to the linear optimizer; compared to switching these reductions off, which usually has a significant cost in performance. This bit is provided for numerically very hard problems, when the polishing step proves to be too expensive (XSLP will report these if any in the final log summary).

Bit 14: Normally, cascading will respect the step bounds of the SLP variable being cascaded. However, allowing the cascaded value to fall outside the step bounds (i.e. expanding the step bounds) can lead to better linearizations, as cascading will set better values for the SLP variables regarding their determining rows; note, that this later strategy might interfere with convergence of the cascaded variables.

Bit 15: When clamping is applied, then in any iteration when the solution would normally be deemed converged on extended criteria only, an extra step bound shrinking step is applied to help imposing strict convergence. In this variant, clamping is only applied on variables that have converged on extended criteria only and have active step bounds.

Bit 16: When clamping is applied, then in any iteration when the solution would normally be deemed converged on extended criteria only, an extra step bound shrinking step is applied to

help imposing strict convergence. In this variant, clamping is applied on all variables that have converged on extended criteria only.

The following constants are provided for setting these bits:

Setting bit 0	XSLP_NOSTEPBOUNDS
Setting bit 1	XSLP_STEPBOUNDSASREQUIRED
Setting bit 2	XSLP_ESTIMATESTEPBOUNDS
Setting bit 3	XSLP_DYNAMICDAMPING
Setting bit 4	XSLP_HOLDVALUES
Setting bit 5	XSLP_RETAINPREVIOUSVALUE
Setting bit 6	XSLP_RESETDELTAZ
Setting bit 7	XSLP_QUICKCONVERGENCECHECK
Setting bit 8	XSLP_ESCALATEPENALTIES
Setting bit 9	XSLP_SWITCHTOPRIMAL
Setting bit 11	XSLP_MAXCOSTOPTION
Setting bit 12	XSLP_RESIDUALERRORS
Setting bit 13	XSLP_NOLPPOLISHING
Setting bit 14	XSLP_CASCADEDBOUNDS
Setting bit 15	XSLP_CLAMPEXTENDEDACTIVESB
Setting bit 16	XSLP_CLAMPEXTENDEDALL

Recommended setting: Bits 1, 2, 5, 7 and usually bits 8 and 9.

Default value	166 (sets bits 1, 2, 5, 7)
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_DELTA_Z, XSLP_ERRORMAXCOST, XSLP_ESCALATION, XSLP_CLAMP SHRINK

XSLP_ANALYZE, SLPANALYZE

Description	Bit map activating additional options supporting model / solution path analysis												
Type	Integer												
Category	Control												
Topic areas	SLP, Bit-vector, Logging												
Values	<table> <tr> <th>Bit</th><th>Meaning</th></tr> <tr> <td>3</td><td>Include an extended iteration summary.</td></tr> <tr> <td>4</td><td>Run infeasibility analysis on infeasible iterations.</td></tr> <tr> <td>6</td><td>Write the linearizations to disk at every XSLP_AUTOSAVE iterations.</td></tr> <tr> <td>7</td><td>Write the initial basis of the linearizations to disk at every XSLP_AUTOSAVE iterations.</td></tr> <tr> <td>8</td><td>Create an XSLP save file at every XSLP_AUTOSAVE iterations.</td></tr> </table>	Bit	Meaning	3	Include an extended iteration summary.	4	Run infeasibility analysis on infeasible iterations.	6	Write the linearizations to disk at every XSLP_AUTOSAVE iterations.	7	Write the initial basis of the linearizations to disk at every XSLP_AUTOSAVE iterations.	8	Create an XSLP save file at every XSLP_AUTOSAVE iterations.
Bit	Meaning												
3	Include an extended iteration summary.												
4	Run infeasibility analysis on infeasible iterations.												
6	Write the linearizations to disk at every XSLP_AUTOSAVE iterations.												
7	Write the initial basis of the linearizations to disk at every XSLP_AUTOSAVE iterations.												
8	Create an XSLP save file at every XSLP_AUTOSAVE iterations.												

Note In most cases, the value of this control does not affect the solution process itself. However, bit 3 (extended summary) will cause SLP to do more function evaluations, and the presence of non-deterministic user functions might cause changes in the solution process. These options are off by default due to performance considerations. The following constants are provided for setting these bits:

Setting bit 3	XSLP_ANALYZE_EXTENDEDFINALSUMMARY
Setting bit 4	XSLP_ANALYZE_INFEASIBLE_ITERATION
Setting bit 6	XSLP_ANALYZE_SAVELINEARIZATIONS
Setting bit 7	XSLP_ANALYZE_SAVEITERBASIS
Setting bit 8	XSLP_ANALYZE_SAVEFILE

Default value	0
See also	XSLP_AUTOSAVE

XSLP_AUGMENTATION, SLPAUGMENTATION

Description	Bit map describing the SLP augmentation method(s) to be used	
Type	Integer	
Category	Control	
Topic areas	SLP, Bit-vector	
Values	Bit	Meaning
	0	Minimum augmentation.
	1	Even handed augmentation.
	2	Penalty error vectors on all non-linear equality constraints.
	3	Penalty error vectors on all non-linear inequality constraints.
	4	Penalty vectors to exceed step bounds.
	5	Use arithmetic means to estimate penalty weights.
	6	Estimate step bounds from values of row coefficients.
	7	Estimate step bounds from absolute values of row coefficients.
	8	Row-based step bounds.
	9	Penalty error vectors on all constraints.
	10	Initial values do not imply an SLP variable.
	12	Avoid running an LP around fixed initial values trying to get feasible.

Notes	<p>Bit 0: Minimum augmentation. Standard augmentation includes delta vectors for all variables involved in nonlinear terms (in non-constant coefficients or as vectors containing non-constant coefficients). Minimum augmentation includes delta vectors only for variables in non-constant coefficients. This produces a smaller linearization, but there is less control on convergence, because convergence control (for example, step bounding) cannot be applied to variables without deltas.</p> <p>Bit 1: Even handed augmentation. Standard augmentation treats variables which appear in non-constant coefficients in a different way from those which contain non-constant coefficients. Even-handed augmentation treats them all in the same way by replacing each non-constant coefficient C in a vector V by a new coefficient $C * V$ in the "equals" column (which has a fixed activity of 1) and creating delta vectors for all types of variable in the same way.</p> <p>Bit 2: Penalty error vectors on all non-linear equality constraints. The linearization of a nonlinear equality constraint is inevitably an approximation and so will not generally be feasible except at the point of linearization. Adding penalty error vectors allows the linear approximation to be violated at a cost and so ensures that the linearized constraint is feasible.</p> <p>Bit 3: Penalty error vectors on all non-linear inequality constraints. The linearization of a nonlinear constraint is inevitably an approximation and so may not be feasible except at the point of linearization. Adding penalty error vectors allows the linear approximation to be violated at a cost and so ensures that the linearized constraint is feasible.</p> <p>Bit 4: Penalty vectors to exceed step bounds. Although it has rarely been found necessary or desirable in practice, Xpress-SLP allows step bounds to be violated at a cost. This may help with feasibility but it generally slows down or prevents convergence, so it should be used only if found absolutely necessary.</p>
--------------	---

Bit 5: Use arithmetic means to estimate penalty weights. Penalty weights are estimated from the magnitude of the elements in the constraint or interacting rows. Geometric means are normally used, so that a few excessively large or small values do not distort the weights significantly. Arithmetic means will value the coefficients more equally.

Bit 6: Estimate step bounds from values of row coefficients. If step bounds are to be imposed from the start, the best approach is to provide explicit values for the bounds. Alternatively, Xpress-SLP can estimate the values from the range of estimated coefficient sizes in the relevant rows.

Bit 7: Estimate step bounds from absolute values of row coefficients. If step bounds are to be imposed from the start, the best approach is to provide explicit values for the bounds. Alternatively, Xpress-SLP can estimate the values from the largest estimated magnitude of the coefficients in the relevant rows.

Bit 8: Row-based step bounds. Step bounds are normally applied as bounds on the delta variables. Some applications may find that using explicit rows to bound the delta vectors gives better results.

Bit 9: Penalty error vectors on all constraints. If the linear portion of the underlying model may actually be infeasible, then applying penalty vectors to all rows may allow identification of the infeasibility and may also allow a useful solution to be found.

Bit 10: Having an initial value will not cause the augmentation to include the corresponding delta variable; i.e. treat the variable as an SLP variable. Useful to provide initial values necessary in the first linearization in case of a minimal augmentation, or as a convenience option when it's easiest to set an initial value for all variables for some reason.

Bit 12: Unless this bit is set, if almost all variables have initial values, an LP can be solved with all variables with initial values fixed to those, to try to extend the partial initial solution to a feasible solution.

The following constants are provided for setting these bits:

Setting bit 0	XSLP_MINIMUMAUGMENTATION
Setting bit 1	XSLP_EVENHANDEDAUGMENTATION
Setting bit 2	XSLP_EQUALITYERRORVECTORS
Setting bit 3	XSLP_ALLERRORVECTORS
Setting bit 4	XSLP_PENALTYDELTAVECTORS
Setting bit 5	XSLP_AMEANWEIGHT
Setting bit 6	XSLP_SBFROMVALUES
Setting bit 7	XSLP_SBFROMABSVALUES
Setting bit 8	XSLP_STEPBOUNDROWS
Setting bit 9	XSLP_ALLROWERRORVECTORS
Setting bit 10	XSLP_NOUPDATEIFONLYIV
Setting bit 12	XSLP_SKIPIVLPHEURISTICS

The recommended setting is bits 2 and 3 (penalty vectors on all nonlinear constraints).

Default value 12 (sets bits 2 and 3)

Affects routines XSLPconstruct

XSLP_AUTOSAVE, SLPAUTOSAVE

Description	Frequency with which to save the model
Type	Integer
Category	Control
Topic areas	SLP, Logging

Note	A value of zero means that the model will not automatically be saved. A positive value of <i>n</i> will save model information at every <i>n</i> th SLP iteration as requested by XSLP_ANALYZE.
Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ANALYZE

XSLP_BARCROSSOVERSTART, SLPBARCROSSOVERSTART

Description	Default crossover activation behaviour for barrier start
Type	Integer
Category	Control
Topic areas	SLP, Linearizations
Note	When XSLP_BARLIMIT is set, XSLP_BARCROSSOVERSTART offers an overwrite control on when crossover is applied. A positive value indicates that crossover should be disabled in iterations smaller than XSLP_BARCROSSOVERSTART and should be enabled afterwards, or when stalling is detected as described in XSLP_BARSTARTOPS. A value of 0 indicates to respect the value of XPRS_CROSSOVER and only overwrite its value when stalling is detected. A value of -1 indicates to always rely on the value of XPRS_CROSSOVER.
Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_BARLIMIT, XSLP_BARSTARTOPS, XSLP_BARSTALLINGLIMIT, XSLP_BARSTALLINGOBLIMIT, XSLP_BARSTALLINGTOL

XSLP_BARLIMIT, SLPBARLIMIT

Description	Number of initial SLP iterations using the barrier method
Type	Integer
Category	Control
Topic areas	SLP, Linearizations, Limits
Note	Particularly for larger models, using the Newton barrier method is faster in the earlier SLP iterations. Later on, when the basis information becomes more useful, a simplex method generally performs better. XSLP_BARLIMIT sets the number of SLP iterations which will be performed using the Newton barrier method.
Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_BARCROSSOVERSTART, XSLP_BARSTARTOPS, XSLP_BARSTALLINGLIMIT, XSLP_BARSTALLINGOBLIMIT, XSLP_BARSTALLINGTOL

XSLP_BARSTALLINGLIMIT, SLPBARSTALLINGLIMIT

Description	Number of iterations to allow numerical failures in barrier before switching to dual
Type	Integer
Category	Control
Topic areas	SLP, Linearizations, Limits
Note	On large problems, it may be beneficial to warm start progress by running a number of iterations with the barrier solver as specified by XSLP_BARLIMIT. On some numerically difficult problems, the barrier may stop prematurely due to numerical issues. Such solves can sometimes be finished if crossover is applied. After XSLP_BARSTALLINGLIMIT such attempts, SLP will automatically switch to use the dual simplex.
Default value	3
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_BARCROSSOVERSTART, XSLP_BARLIMIT, XSLP_BARSTARTOPS, XSLP_BARSTALLINGOBJLIMIT, XSLP_BARSTALLINGTOL

XSLP_BARSTALLINGOBJLIMIT, SLPBARSTALLINGOBJLIMIT

Description	Number of iterations over which to measure the objective change for barrier iterations with no crossover
Type	Integer
Category	Control
Topic areas	SLP, Linearizations, Limits
Note	On large problems, it may be beneficial to warm start progress by running a number of iterations with the barrier solver without crossover by setting XSLP_BARLIMIT to a positive value and setting XPRS_CROSSOVER to 0. A potential drawback is slower convergence due to the interior point provided by the barrier solve keeping a higher number of variables active. This may lead to stalling in progress, negating the benefit of using the barrier. When in the last XSLP_BARSTALLINGOBJLIMIT iterations no significant progress has been made, crossover is automatically enabled.
Default value	3
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_BARCROSSOVERSTART, XSLP_BARLIMIT, XSLP_BARSTARTOPS, XSLP_BARSTALLINGLIMIT, XSLP_BARSTALLINGTOL

XSLP_BARSTARTOPS, SLPBARSTARTOPS

Description	Controls behaviour when the barrier is used to solve the linearizations
Type	Integer
Category	Control
Topic areas	SLP, Linearizations

Values	Bit	Meaning
	0	Check objective progress when no crossover is applied.
	1	Fall back to dual simplex if too many numerical problems are reported by the barrier.
	2	If a non-vertex converged solution found by barrier without crossover can be returned as a final solution.
Note	The following constants are provided for setting these bits:	
	Setting bit 0	BARSTARTOPS_STALLING_OBJECTIVE
	Setting bit 1	BARSTARTOPS_STALLING_NUMERICAL
	Setting bit 2	BARSTARTOPS_ALLOWINTERIORSOLUTION
Default value	-1	
Affects routines	XSLPmaxim, XSLPminim	
See also	XSLP_BARCROSSOVERSTART, XSLP_BARLIMIT, XSLP_BARSTALLINGLIMIT, XSLP_BARSTALLINGOBJLIMIT, XSLP_BARSTALLINGTOL	

XSLP_CALCTHREADS, NLPCALCTHREADS

Description	Number of threads used for formula and derivatives evaluations
Type	Integer
Category	Control
Topic area	Parallel
Note	When beneficial, SLP can calculate formula values and partial derivative information in parallel.
Default value	-1 (determined by XSLP_THREADS)
Affects routines	XSLPmaxim, XSLPmaxim
See also	XSLP_THREADS,

XSLP_CASCADE, SLPCASCADE

Description	Bit map describing the cascading to be used
Type	Integer
Category	Control
Topic areas	SLP, Cascading

Values	Bit	Meaning
	0	Apply cascading to all variables with determining rows.
	1	Apply cascading to SLP variables which appear in coefficients and which would change by more than XPRS_FEASTOL.
	2	Apply cascading to all SLP variables which appear in coefficients.
	3	Apply cascading to SLP variables which are structural and which would change by more than XPRS_FEASTOL.
	4	Apply cascading to all SLP variables which are structural.
	5	Create secondary order grouping DR rows with instantiated user functions together in the order.
	6	In cases where the determining column is below XSLP_DRCOLTOL, fix at the previous rather than current value.
	7	In cases where the determining column is below XSLP_DRCOLTOL, fix within a range XSLP_DRFIXRANGE of previous value.
	8	Automatically determine whether to apply cascading.
Note	<p>Normal cascading (bit 0) uses determining rows to recalculate the values of variables to be consistent with values already available or already recalculated. Other bit settings are normally required only in quadratic programming where some of the SLP variables are in the objective function. The values of such variables may need to be corrected if the corresponding update row is slightly infeasible. The following constants are provided for setting these bits:</p> <p>Setting bit 0 XSLP_CASCADE_ALL Setting bit 1 XSLP_CASCADE_COEF_VAR Setting bit 2 XSLP_CASCADE_ALL_COEF_VAR Setting bit 3 XSLP_CASCADE_STRUCT_VAR Setting bit 4 XSLP_CASCADE_ALL_STRUCT_VAR Setting bit 5 XSLP_CASCADE_SECONDARY_GROUPS Setting bit 6 XSLP_CASCADE_DRCOL_PREVIOUSVALUE Setting bit 7 XSLP_CASCADE_DRCOL_PVRANGE Setting bit 8 XSLP_CASCADE_AUTOAPPLY</p>	
Default value	257	
Affects routines	XSLPcascade	
See also	XSLP_DRCOLTOL, XSLP_DRFIXRANGE	

XSLP_CASCADENLIMIT, SLPCASCADENLIMIT

Description	Maximum number of iterations for cascading with non-linear determining rows
Type	Integer
Category	Control
Topic areas	SLP, Cascading, Limits
Note	Re-calculation of the value of a variable uses a modification of the Newton-Raphson method. The maximum number of steps in the method is set by XSLP_CASCADENLIMIT. If the maximum number of steps is taken without reaching a converged value, the best value found will be used.
Default value	10
Affects routines	XSLPcascade

See also XSLP_CASCADE

XSLP_CONTROL

Description	Bit map describing which Xpress NonLinear functions also activate the corresponding Optimizer Library function	
Type	Integer	
Category	Control	
Topic area	Bit-vector	
Values	Bit	Meaning
	0	Xpress NonLinear problem management functions do NOT invoke the corresponding Optimizer Library function for the underlying linear problem.
	1	XSLPcopycontrols does NOT invoke XPRScopycontrols.
	2	XSLPcopycallbacks does NOT invoke XPRScopycallbacks.
	3	XSLPcopyprob does NOT invoke XPRScopyprob.
	4	XSLPsetdefaults does NOT invoke XPRSsetdefaults.
	5	XSLPsave does NOT invoke XPRSsave.
	6	XSLPrestore does NOT invoke XPRSrestore.
Note	<p>The problem management functions are:</p> <p>XSLPcopyprob to copy from an existing problem;</p> <p>XSLPcopycontrols and XSLPcopycallbacks to copy the current controls and callbacks from an existing problem;</p> <p>XSLPsetdefaults to reset the controls to their default values;</p> <p>XSLPsave and XSLPrestore for saving and restoring a problem.</p>	
Default value	0 (no bits set)	
Affects routines	XSLPcopycontrols, XSLPcopycallbacks, XSLPcopyprob, XSLPrestore, XSLPsave, XSLPsetdefaults	

XSLP_CONVERGENCEOPS, SLPCONVERGENCEOPS

Description	Bit map describing which convergence tests should be carried out
Type	Integer
Category	Control
Topic areas	SLP, SLP-convergence, Bit-vector

Values	Bit	Meaning
	0	Execute the closure tolerance checks.
	1	Execute the delta tolerance checks.
	2	Execute the matrix tolerance checks.
	3	Execute the impact tolerance checks.
	4	Execute the slack impact tolerance checks.
	5	Check for user provided convergence.
	6	Execute the objective range checks.
	7	Execute the objective range + constraint activity check.
	8	Execute the objective range + active step bound check.
	9	Execute the convergence continuation check.
	10	Take scaling of individual variables / rows into account.
	11	Execute the validation target convergence checks.
	12	Execute the first order optimality target convergence checks.
	13	Allow convex quadratic problems to converge on extended criteria.
	15	Do not declare converged if still sufficient improvement in objective.

Note Provides fine tuned control (over setting the related convergence tolerances) of which convergence checks are carried out.

The following constants are provided for setting these bits:

Setting bit 0	XSLP_CONVERGEBIT_CTOL
Setting bit 1	XSLP_CONVERGEBIT_ATOL
Setting bit 2	XSLP_CONVERGEBIT_MTOL
Setting bit 3	XSLP_CONVERGEBIT_ITOL
Setting bit 4	XSLP_CONVERGEBIT_STOL
Setting bit 5	XSLP_CONVERGEBIT_USER
Setting bit 6	XSLP_CONVERGEBIT_VTOL
Setting bit 7	XSLP_CONVERGEBIT_XTOL
Setting bit 8	XSLP_CONVERGEBIT_OTOL
Setting bit 9	XSLP_CONVERGEBIT_WTOL
Setting bit 10	XSLP_CONVERGEBIT_EXTENDEDSCALING
Setting bit 11	XSLP_CONVERGEBIT_VALIDATION
Setting bit 12	XSLP_CONVERGEBIT_VALIDATION_K
Setting bit 13	XSLP_CONVERGEBIT_NOQUADCHECK
Setting bit 15	XSLP_CONVERGEBIT_REQUIRE_OTOL_R

Default value 39935 (bits 0-9, 11-12, and 15 are set)

Affects routines XSLPmaxim, XSLPminim

XSLP_CUTSTRATEGY, SLPCUTSTRATEGY

Description	Determines which cuts to apply in the MISLP search when the default SLP-in-MIP strategy is used.
Type	Integer
Category	Control
Topic areas	MISLP, Cuts

Note	Cuts are derived from the linearizations and are local cuts in that they are valid in the linearization and not necessarily valid for the full problem. The values mirror that of XPRS_CUTSTRATEGY.
Default value	0
Affects routines	XSLPminim, XSLPmaxim

XSLP_DAMPSTART, SLPDAMPSTART

Description	SLP iteration at which damping is activated
Type	Integer
Category	Control
Topic area	SLP
Note	If damping is used as part of the SLP algorithm, it can be delayed until a specified SLP iteration. This may be appropriate when damping is used to encourage convergence after an un-damped algorithm has failed to converge.
Default value	0
Affects routines	XSLPmaxim, XSLPmaxim
See also	XSLP_ALGORITHM, XSLP_DAMPEXPAND, XSLP_DAMPMAX, XSLP_DAMPMIN, XSLP_DAMP SHRINK

XSLP_DELAYUPDATEROWS, SLPDELAYUPDATEROWS

Description	Number of SLP iterations before update rows are fully activated
Type	Integer
Category	Control
Topic area	SLP
Note	During augmentation, one or more delta vectors are created for each SLP variable. The values of these are linked to that of the variable through an <i>update row</i> which is created as part of the augmentation procedure.
Default value	2
Affects routines	XSLPmaxim, XSLPmaxim

XSLP_DELTAOFFSET, SLPDELTAOFFSET

Description	Position of first character of SLP variable name used to create name of delta vector
Type	Integer
Category	Control
Topic areas	SLP, Logging

Note	During augmentation, a delta vector, and possibly penalty delta vectors, are created for each SLP variable. They are created with names derived from the corresponding SLP variable. Customized naming is possible using <code>XSLP_DELTAFORMAT</code> etc to define a format and <code>XSLP_DELTAOFFSET</code> to define the first character (counting from zero) of the variable name to be used.
Default value	0
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_DELTAFORMAT</code> , <code>XSLP_MINUSDELTAFORMAT</code> , <code>XSLP_PLUSDELTAFORMAT</code>

XSLP_DELTAZLIMIT, SLPDELTAZLIMIT

Description	Number of SLP iterations during which to apply <code>XSLP_DELTA_Z</code>
Type	Integer
Category	Control
Topic area	SLP
Note	<code>XSLP_DELTA_Z</code> is used to retain small derivatives which would otherwise be regarded as zero. This is helpful in avoiding local optima, but may make the linearized problem more difficult to solve because of the number of small nonzero elements in the resulting matrix. <code>XSLP_DELTAZLIMIT</code> can be set to a nonzero value, which is then the number of iterations for which <code>XSLP_DELTA_Z</code> will be used. After that, small derivatives will be set to zero. A negative value indicates no automatic perturbations to the derivatives in any situation.
Default value	0
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_DELTA_Z</code>

XSLP_DERIVATIVES, NLPDERIVATIVES

Description	Bitmap describing the method of calculating derivatives						
Type	Integer						
Category	Control						
Topic area	Derivatives						
Values	<table> <tr> <th>Bit</th><th>Meaning</th></tr> <tr> <td>0</td><td>analytic derivatives where possible</td></tr> <tr> <td>1</td><td>avoid embedding numerical derivatives of instantiated functions into analytic derivatives</td></tr> </table>	Bit	Meaning	0	analytic derivatives where possible	1	avoid embedding numerical derivatives of instantiated functions into analytic derivatives
Bit	Meaning						
0	analytic derivatives where possible						
1	avoid embedding numerical derivatives of instantiated functions into analytic derivatives						
Notes	<p>If no bits are set then numerical derivatives are calculated using finite differences. Analytic derivatives cannot be used for formulae involving discontinuous functions. They may not work well with functions which are not smooth (such as <code>MAX</code>), or where the derivative changes very quickly with the value of the variable (such as <code>LOG10</code> of small values). Both first and second order analytic derivatives can either be calculated as symbolic formulas, or by the means of auto-differentiation, with the exception that the second order symbolic derivatives require that the first order derivatives are also calculated using the symbolic method.</p>						

Default value	1
Affects routines	XSLPconstruct, XSLPmaxim, XSLPminim
See also	XSLP_JACOBIAN, XSLP_HESSIAN

XSLP_DETERMINISTIC, NLPDETERMINISTIC

Description	Determines if the parallel features of SLP should be guaranteed to be deterministic
Type	Integer
Category	Control
Topic areas	SLP, MISLP, Parallel
Note	Determinism can only be guaranteed if no callbacks are used, or if in the presence of callbacks the effect of the callbacks only depend on local information provided by SLP.
Default value	1
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART_POOLSIZE,

XSLP_ECFCHECK, SLPECFCHECK

Description	Check feasibility at the point of linearization for extended convergence criteria
Type	Integer
Category	Control
Topic areas	SLP, SLP-convergence
Values	0 no check (extended criteria are always used); 1 check until one infeasible constraint is found; 2 check all constraints.
Notes	<p>The extended convergence criteria measure the accuracy of the solution of the linear approximation compared to the solution of the original nonlinear problem. For this to work, the linear approximation needs to be reasonably good at the point of linearization. In particular, it needs to be reasonably close to feasibility.</p> <p>XSLP_ECFCHECK is used to determine what checking of feasibility is carried out at the point of linearization. If the point of linearization at the start of an SLP iteration is deemed to be infeasible, then the extended convergence criteria are not used to decide convergence at the end of that SLP iteration.</p> <p>If all that is required is to decide that the point of linearization is not feasible, then the search can stop after the first infeasible constraint is found (parameter is set to 1). If the actual number of infeasible constraints is required, then XSLP_ECFCHECK should be set to 2, and all constraints will be checked.</p> <p>The number of infeasible constraints found at the point of linearization is returned in XSLP_ECFCOUNT.</p>
Default value	1
Affects routines	Convergence criteria, XSLPmaxim, XSLPminim
See also	XSLP_ECFCOUNT, XSLP_ECFTOL_A, XSLP_ECFTOL_R

XSLP_ECHOXPRSMESSAGES

Description	Controls if the XSLP message callback should relay messages from the XPRS library.	
Type	Integer	
Category	Control	
Topic area	Logging	
Values	-1	automatic: if an XSLP message callback is not set, then messages from the nonlinear solver are sent to the XPRS message callback; if an XSLP message callback is set, then messages are not echoed.
	0	the XPRS and XSLP message callbacks are treated as independent.
	1	messages from the XPRS message callback are sent to the XSLP message callback.
	2	messages from the nonlinear solver are sent to the XPRS message callback.
Default value	-1	

XSLP_ERROROFFSET, SLPERROROFFSET

Description	Position of first character of constraint name used to create name of penalty error vectors	
Type	Integer	
Category	Control	
Topic areas	SLP, Logging	
Note	During augmentation, penalty error vectors may be created for some or all of the constraints. The vectors are created with names derived from the corresponding constraint name. Customized naming is possible using XSLP_MINUSERERRORFORMAT and XSLP_PLUSERERRORFORMAT to define a format and XSLP_ERROROFFSET to define the first character (counting from zero) of the constraint name to be used.	
Default value	0	
Affects routines	XSLPconstruct	
See also	XSLP_MINUSERERRORFORMAT, XSLP_PLUSERERRORFORMAT	

XSLP_EVALUATE, NLPEVALUATE

Description	Evaluation strategy for user functions	
Type	Integer	
Category	Control	
Topic area	User Functions	
Values	0	use derivatives where possible;
	1	always re-evaluate.
Note	If a user function returns derivatives or returns more than one value, then it is possible for Xpress NonLinear to estimate the value of the function from its derivatives if the new point of evaluation is sufficiently close to the original. Setting XSLP_EVALUATE to 1 will force re-evaluation of all functions regardless of how much or little the point of evaluation has changed.	

Default value	0
Affects routines	XSLPevaluatecoef, XSLPevaluateformula
See also	XSLP_FUNCEVAL

XSLP_FILTER, SLPFILTER

Description	Bit map for controlling solution updates	
Type	Integer	
Category	Control	
Topic areas	SLP, Bit-vector, Solution	
Values	Bit	Meaning
	0	retain best solution according to the merit function.
	1	check cascaded solutions against improvements in the merit function.
	2	force minimum step sizes in line search.
	3	accept the trust region step if the line search returns a zero step size.
Notes	<p>Bit 0 determines whether XSLPgetslopsol should return the final converged solution (if the bit is off), or the solution which had the best value according to the merit function (if the bit is on, which is the default).</p> <p>If bit 1 is set, a cascaded solution which does not improve the merit function will be rejected (XSLP will revert to the solution of the linearization).</p> <p>Bits 2-3 determine the strategy for when the step direction is not improving according to the merit function.</p> <p>The following constants are provided for setting these bits:</p> <p>Setting bit 0 XSLP_FILTER_KEEPPBEST</p> <p>Setting bit 1 XSLP_FILTER_CASCADE</p> <p>Setting bit 2 XSLP_FILTER_ZEROLINESEARCH</p> <p>Setting bit 3 XSLP_FILTER_ZEROLINESEARCHTR</p>	
Default value	3 (bit 0,1)	
Affects routines	XSLPmaxim, XSLPminim, XSLPcascade	
See also	XSLP_MERITLAMBDA, XSLP_CASCADE, XSLP_LSSTART, XSLP_LSITERLIMIT, XSLP_LSPATTERNLIMIT	

XSLP_FINDIV, NLPFINDIV

Description	Option for running a heuristic to find a feasible initial point	
Type	Integer	
Category	Control	
Topic area	Heuristics	
Values	-1	Automatic (default).
	0	Disable the heuristic.
	1	Enable the heuristic.

Notes The procedure uses bound reduction (and, up to an extent, probing) to obtain a point in the initial bounding box that is feasible for the bound reduction techniques. If an initial point is already specified and is found not to violate bound reduction, then the heuristic is not run and the given point is used as the initial solution.

Default value -1

Affects routines XSLPmaxim, XSLPminim

XSLP_FUNCEVAL, NLPFUNCEVAL

Description Bit map for determining the method of evaluating user functions and their derivatives

Type Integer

Category Control

Topic areas User Functions, Derivatives

Values	Bit	Meaning
	3	evaluate function whenever independent variables change.
	4	evaluate function when independent variables change outside tolerances.
	5	application of bits 3-4: 0 = functions which do not have a defined re-evaluation mode; 1 = all functions.
	6	tangential derivatives.
	7	forward derivatives
	8	application of bits 6-7: 0 = functions which do not have a defined derivative mode; 1 = all functions.

Notes Bits 3-4 determine the type of function re-evaluation. If both bits are zero, then the settings for each individual function are used. If bit 3 or bit 4 is set, then bit 5 defines which functions the setting applies to. If it is set to 1, then it applies to all functions. Otherwise, it applies only to functions which do not have an explicit setting of their own. Bits 6-7 determine the type of calculation for numerical derivatives. If both bits are zero, then the settings for each individual function are used. If bit 6 or bit 7 is set, then bit 8 defines which functions the setting applies to. If it is set to 1, then it applies to all functions. Otherwise, it applies only to functions which do not have an explicit setting of their own.

The following constants are provided for setting these bits:

Setting bit 3	XSLP_RECALC
Setting bit 4	XSLP_TOLCALC
Setting bit 5	XSLP_ALLCALCS
Setting bit 6	XSLP_2DERIVATIVE
Setting bit 7	XSLP_1DERIVATIVE
Setting bit 8	XSLP_ALLDERIVATIVES

Default value 0

Affects routines XSLPevaluatecoef, XSLPevaluateformula

See also XSLP_EVALUATE

XSLP_GRIDHEURSELECT, SLPGRIDHEURSELECT

Description	Bit map selectin which heuristics to run if the problem has variable with an integer delta	
Type	Integer	
Category	Control	
Topic areas	SLP, Heuristics	
Values	Bit	Meaning
	0	Enumeration: try all combinations.
	1	Simple search heuristics.
	2	Simulated annealing.
Note	A value of 0 indicates that integer deltas are only taken into consideration during the SLP iterations.	
Note	The enumeration option can be useful for cases where the number of possible values of the variables with an integer delta is small.	
Default value	6	
Affects routines	XSLPmaxim, XSLPminim	

XSLP_HEURSTRATEGY, SLPHEURSTRATEGY

Description	Branch and Bound: This specifies the MINLP heuristic strategy. On some problems it is worth trying more comprehensive heuristic strategies by setting HEURSTRATEGY to 2 or 3.	
Type	Integer	
Category	Control	
Topic areas	SLP, Heuristics	
Values	-1	Automatic selection of heuristic strategy.
	0	No heuristics.
	1	Basic heuristic strategy.
	2	Enhanced heuristic strategy.
	3	Extensive heuristic strategy.
	4	Run all heuristics without effort limits.
Default value	-1	
Affects routines	XSLPminim, XSLPmaxim.	

XSLP_HESSIAN, NLPHESSIAN

Description	Second order differentiation mode when using analytical derivatives
Type	Integer
Category	Control
Topic area	Derivatives

Values	-1, 0 automatic selection 1 numerical derivatives (finite difference) 2 symbolic differentiation 3 automatic differentiation
Note	Symbolic mode differentiation for the second order derivatives is only available when XSLP_JACOBIAN is also set to symbolic mode.
Default value	-1
See also	XSLP_DERIVATIVES, XSLP_JACOBIAN

XSLP_INFEASLIMIT, SLPINFEASLIMIT

Description	The maximum number of consecutive infeasible SLP iterations which can occur before Xpress-SLP terminates
Type	Integer
Category	Control
Topic areas	SLP, Limits
Note	An infeasible solution to an SLP iteration means that is likely that Xpress-SLP will create a poor linear approximation for the next SLP iteration. Sometimes, small infeasibilities arise because of numerical difficulties and do not seriously affect the solution process. However, if successive solutions remain infeasible, it is unlikely that Xpress-SLP will be able to find a feasible converged solution. XSLP_INFEASLIMIT sets the number of successive SLP iterations which must take place before Xpress-SLP terminates with a status of "infeasible solution".
Default value	3
Affects routines	XSLPmaxim, XSLPminim

XSLP_ITERLIMIT, SLPITERLIMIT

Description	The maximum number of SLP iterations
Type	Integer
Category	Control
Topic areas	SLP, Limits
Note	If Xpress-SLP reaches XSLP_ITERLIMIT without finding a converged solution, it will stop. For MISLP, the limit is on the number of SLP iterations at each node.
Default value	1000
Affects routines	XSLPnlpoptimize, XSLPmaxim, XSLPminim

XSLP_JACOBIAN, NLPJACOBIAN

Description	First order differentiation mode when using analytical derivatives
Type	Integer
Category	Control

Topic area	Derivatives
Values	-1, 0 automatic selection 1 numerical derivatives (finite difference) 2 symbolic differentiation 3 automatic differentiation
Note	Symbolic mode differentiation for the second order derivatives is only available when XSLP_JACOBIAN is set to symbolic mode.
Default value	-1
See also	XSLP_DERIVATIVES, XSLP_HESSIAN

XSLP_KEEPEQUALSCOLUMN, NLPKEEPEQUALSCOLUMN

Description	When set to a nonzero value, the MPS reader will keep the equals column in the problem
Type	Integer
Category	Control
Topic area	Misc
Note	This control is provided mainly for backward compatibility.
Default value	0
Affects routines	XSLPreadprob

XSLP_LINQUADBR, NLPLINQUADBR

Description	Use linear and quadratic constraints and objective function to further reduce bounds on all variables
Type	Integer
Category	Control
Topic area	Presolve
Values	-1 automatic selection 0 disable 1 enable
Note	While bound reduction is effective when performed on nonlinear, nonquadratic constraints and objective function, it can be useful to obtain tightened bounds from linear and quadratic constraints, as the corresponding variables may appear in other nonlinear constraints. This option then allows for a slightly more expensive bound reduction procedure, at the benefit of further reduction in the problem's bounds.
Default value	-1
See also	XSLP_PRESOLVEOPS, XSLP_PROBING

XSLP_LOG, NLPLOG

Description	Level of printing during SLP iterations
Type	Integer
Category	Control
Topic areas	SLP, Logging
Values	-1 none 0 minimal 1 normal: iteration, penalty vectors 2 omit from convergence log any variables which have converged 3 omit from convergence log any variables which have already converged (except variables on step bounds) 4 include all variables in convergence log 5 include user function call communications in the log
Default value	0
Affects routines	XSLPmaxim, XSLPminim

XSLP_LSITERLIMIT, SLPLSITERLIMIT

Description	Number of iterations in the line search
Type	Integer
Category	Control
Topic areas	SLP, Limits
Notes	The line search attempts to refine the step size suggested by the trust region step bounds. The line search is a local method; the control sets a maximum on the number of model evaluations during the line search.
Default value	0
See also	XSLP_LSPATTERNLIMIT, XSLP_LSSTART, XSLP_LSZEROLIMIT, XSLP_FILTER
Affects routines	XSLPmaxim, XSLPminim

XSLP_LSPATTERNLIMIT, SLPLSPATTERNLIMIT

Description	Number of iterations in the pattern search preceding the line search
Type	Integer
Category	Control
Topic areas	SLP, Limits
Notes	When positive, defines the number of samples taken along the step size suggested by the trust region step bounds before initiating the line search. Useful for highly non-convex problems.
Default value	0
See also	XSLP_LSITERLIMIT, XSLP_LSSTART, XSLP_LSZEROLIMIT, XSLP_FILTER
Affects routines	XSLPmaxim, XSLPminim

XSLP_LSSTART, SLPLSSTART

Description	Iteration in which to active the line search
Type	Integer
Category	Control
Topic area	SLP
Notes	
Default value	8
See also	XSLP_LSITERLIMIT, XSLP_LSPATTERNLIMIT, XSLP_LSZEROLIMIT, XSLP_FILTER
Affects routines	XSLPmaxim, XSLPminim

XSLP_LSZEROLIMIT, SLPLSZEROLIMIT

Description	Maximum number of zero length line search steps before line search is deactivated
Type	Integer
Category	Control
Topic areas	SLP, Limits
Notes	When the line search repeatedly returns a zero step size, counteracted by bits set on XSLP_FILTER, the effort spent in line search is redundant, and line search will be deactivated after XSLP_LSZEROLIMIT consecutive such iteration.
Default value	5
See also	XSLP_LSITERLIMIT, XSLP_LSPATTERNLIMIT, XSLP_LSSTART, XSLP_FILTER
Affects routines	XSLPmaxim, XSLPminim

XSLP_MAXTIME, NLPMAXTIME

Description	The maximum time in seconds that the SLP optimization will run before it terminates
Type	Integer
Category	Control
Topic area	Limits
Notes	The (elapsed) time is measured from the beginning of the first SLP optimization. If XSLP_MAXTIME is negative, Xpress NonLinear will terminate after (-XSLP_MAXTIME) seconds. If it is positive, Xpress NonLinear will terminate in MISLP after XSLP_MAXTIME seconds or as soon as an integer solution has been found thereafter.
Default value	0
Affects routines	XSLPnloptimize, XSLPmaxim, XSLPminim

XSLP_MIPALGORITHM, SLP_MIPALGORITHM

Description	Bitmap describing the MISLP algorithms to be used	
Type	Integer	
Category	Control	
Topic areas	MISLP, Bit-vector	
Values	Bit	Meaning
	0	Solve initial SLP to convergence.
	2	Relax step bounds according to XSLP_MIPRELAXSTEPBOUNDS after initial node.
	3	Fix step bounds according to XSLP_MIPFIXSTEPBOUNDS after initial node.
	4	Relax step bounds according to XSLP_MIPRELAXSTEPBOUNDS at each node.
	5	Fix step bounds according to XSLP_MIPFIXSTEPBOUNDS at each node.
	6	Limit iterations at each node to XSLP_MIPITERLIMIT.
	7	Relax step bounds according to XSLP_MIPRELAXSTEPBOUNDS after MIP solution is found.
	8	Fix step bounds according to XSLP_MIPFIXSTEPBOUNDS after MIP solution is found.
	9	Use MIP at each SLP iteration instead of SLP at each node.
	10	Use MIP on converged SLP solution and then SLP on the resulting MIP solution.
Notes	<p>XSLP_MIPALGORITHM determines the strategy of XSLPnloptimize for solving MINLP problems. The recommended approach is to solve the problem first without reference to the discrete variables. This can be handled automatically by setting bit 0 of XSLP_MIPALGORITHM; if done manually, then optimize using the "I" option to prevent the Optimizer presolve from changing the problem.</p> <p>Some versions of the optimizer re-run the initial node as part of the tree search; it is possible to initiate a new SLP optimization at this point by relaxing or fixing step bounds (use bits 2 and 3). If step bounds are fixed for a class of variable, then the variables in that class will not change their value in any child node.</p> <p>At each node, it is possible to relax or fix step bounds. It is recommended that step bounds are relaxed, so that the new problem can be solved starting from its parent, but without undue restrictions caused by step bounding (use bit 4). Exceptionally, it may be preferable to restrict the freedom of child nodes by relaxing fewer types of step bound or fixing the values of some classes of variable (use bit 5).</p> <p>When the optimal node has been found, it is possible to fix the discrete variables and then re-optimize with SLP. Step bounds can be relaxed or fixed for this optimization as well (use bits 7 and 8).</p> <p>Although it is ultimately necessary to solve the optimal node to convergence, individual nodes can be truncated after XSLP_MIPITERLIMIT SLP iterations. Set bit 6 to activate this feature. The normal MISLP algorithm uses SLP at each node. One alternative strategy is to use the MIP optimizer for solving each SLP iteration. Set bit 9 to implement this strategy ("MIP within SLP").</p> <p>Another strategy is to solve the problem to convergence ignoring the nature of the integer variables. Then, fixing the linearization, use MIP to find the optimal setting of the discrete variables. Then, fixing the discrete variables, but varying the linearization, solve to convergence. Set bit 10 to implement this strategy ("SLP then MIP").</p> <p>For more details about MISLP algorithms and strategies, see the separate section.</p> <p>The following constants are provided for setting these bits:</p>	

Setting bit 0	XSLP_MIPINITIALSLP
Setting bit 2	XSLP_MIPINITIALRELAXSLP
Setting bit 3	XSLP_MIPINITIALFIXSLP
Setting bit 4	XSLP_MIPNODERELAXSLP
Setting bit 5	XSLP_MIPNODEFIXSLP
Setting bit 6	XSLP_MIPNODELIMITSLP
Setting bit 7	XSLP_MIPFINALRELAXSLP
Setting bit 8	XSLP_MIPFINALFIXSLP
Setting bit 9	XSLP_MIPWITHINSLP
Setting bit 10	XSLP_SLPTHENMIP
Setting bit 11	XSLP_NOFINALROUNDING

Default value 17 (bits 0 and 4 are set)

Affects routines XSLPnlpoptimize

See also XSLP_ALGORITHM, XSLP_MIPFIXSTEPBOUNDS, XSLP_MIPITERLIMIT, XSLP_MIPRELAXSTEPBOUNDS

XSLP_MIPCUTOFFCOUNT, SLPMIPCUTOFFCOUNT

Description Number of SLP iterations to check when considering a node for cutting off

Type Integer

Category Control

Topic areas MISLP, Limits

Notes If the objective function is worse by a defined amount than the best integer solution obtained so far, then the SLP will be terminated (and the node will be cut off). The node will be cut off at the current SLP iteration if the objective function for the last XSLP_MIPCUTOFFCOUNT SLP iterations are all worse than the best obtained so far, and the difference is greater than $XSLP_MIPCUTOFF_A$ and $OBJ * XSLP_MIPCUTOFF_R$ where OBJ is the best integer solution obtained so far.
The test is not applied until at least XSLP_MIPCUTOFFLIMIT SLP iterations have been carried out at the current node.

Default value 5

Affects routines XSLPnlpoptimize

See also XSLP_MIPCUTOFF_A, XSLP_MIPCUTOFF_R, XSLP_MIPCUTOFFLIMIT

XSLP_MIPCUTOFFLIMIT, SLPMIPCUTOFFLIMIT

Description Number of SLP iterations to check when considering a node for cutting off

Type Integer

Category Control

Topic areas MISLP, Limits

Notes If the objective function is worse by a defined amount than the best integer solution obtained so far, then the SLP will be terminated (and the node will be cut off). The node will be cut off at the current SLP iteration if the objective function for the last XSLP_MIPCUTOFFCOUNT SLP iterations are all worse than the best obtained so far, and the difference is greater than

$XSLP_MIPCUTOFF_A$ and $OBJ * XSLP_MIPCUTOFF_R$ where OBJ is the best integer solution obtained so far.

The test is not applied until at least $XSLP_MIPCUTOFFLIMIT$ SLP iterations have been carried out at the current node.

Default value	10
Affects routines	<code>XSLPnlpoptimize</code>
See also	<code>XSLP_MIPCUTOFF_A</code> , <code>XSLP_MIPCUTOFF_R</code> , <code>XSLP_MIPCUTOFFCOUNT</code>

`XSLP_MIPDEFAULTALGORITHM`, `SLPMIPDEFAULTALGORITHM`

Description	Default algorithm to be used during the tree search in MISLP
Type	Integer
Category	Control
Topic areas	MISLP, Linearizations
Note	The default algorithm used within SLP during the MISLP optimization can be set using <code>XSLP_MIPDEFAULTALGORITHM</code> . It will not necessarily be the same as the one best suited to the initial SLP optimization.
Default value	3
Affects routines	<code>XSLPnlpoptimize</code>
See also	<code>XPRS_DEFAULTALG</code> , <code>XSLP_MIPALGORITHM</code>

`XSLP_MIPFIXSTEPBOUNDS`, `SLPMIPFIXSTEPBOUNDS`

Description	Bitmap describing the step-bound fixing strategy during MISLP	
Type	Integer	
Category	Control	
Topic area	MISLP	
Values	Bit	Meaning
	0	Fix step bounds on structural SLP variables which are not in coefficients.
	1	Fix step bounds on all structural SLP variables.
	2	Fix step bounds on SLP variables appearing only in coefficients.
	3	Fix step bounds on SLP variables appearing in coefficients.
Note	At any node (including the initial and optimal nodes) it is possible to fix the step bounds of classes of variables so that the variables themselves will not change. This may help with convergence, but it does increase the chance of a local optimum because of excessive artificial restrictions on the variables.	
Default value	0	
Affects routines	<code>XSLPnlpoptimize</code>	
See also	<code>XSLP_MIPALGORITHM</code> , <code>XSLP_MIPRELAXSTEPBOUNDS</code>	

XSLP_MIPITERLIMIT, SLPMIPITERLIMIT

Description	Maximum number of SLP iterations at each node
Type	Integer
Category	Control
Topic areas	MISLP, Limits
Note	If bit 6 of XSLP_MIPALGORITHM is set, then the number of iterations at each node will be limited to XSLP_MIPITERLIMIT.
Default value	0
Affects routines	XSLPnlpoptimize
See also	XSLP_ITERLIMIT, XSLP_MIPALGORITHM

XSLP_MIPLOG, SLPMIPLOG

Description	Frequency with which MIP status is printed
Type	Integer
Category	Control
Topic areas	MISLP, Logging
Note	By default (zero or negative value) the MIP status is printed after synchronization points. If XSLP_MIPLOG is set to a positive integer, then the current MIP status (node number, best value, best bound) is printed every XSLP_MIPLOG nodes.
Default value	0 (deterministic logging)
Affects routines	XSLPnlpoptimize
See also	XSLP_LOG, XSLP_SLPMIPLOG

XSLP_MIPOCOUNT, SLPMIPOCOUNT

Description	Number of SLP iterations at each node over which to measure objective function variation
Type	Integer
Category	Control
Topic areas	MISLP, SLP-convergence
Note	The objective function test for MIP termination is applied only when step bounding has been applied (or XSLP_SBSTART SLP iterations have taken place if step bounding is not being used). The node will be terminated at the current SLP iteration if the range of the objective function values over the last XSLP_MIPOCOUNT SLP iterations is within XSLP_MIPOTOL_A or within $OBJ * XSLP_MIPOTOL_R$ where OBJ is the average value of the objective function over those iterations.
Default value	5
Affects routines	XSLPnlpoptimize
See also	XSLP_MIPOTOL_A XSLP_MIPOTOL_R XSLP_SBSTART

XSLP_MIPRELAXSTEPBOUNDS, SLP_MIPRELAXSTEPBOUNDS

Description	Bitmap describing the step-bound relaxation strategy during MISLP	
Type	Integer	
Category	Control	
Topic area	MISLP	
Values	Bit	Meaning
	0	Relax step bounds on structural SLP variables which are not in coefficients.
	1	Relax step bounds on all structural SLP variables.
	2	Relax step bounds on SLP variables appearing only in coefficients.
	3	Relax step bounds on SLP variables appearing in coefficients.
Note	At any node (including the initial and optimal nodes) it is possible to relax the step bounds of classes of variables so that the variables themselves are completely free to change. This may help with finding a global optimum, but it may also increase the solution time, because more SLP iterations are necessary at each node to obtain a converged solution.	
Default value	15 (relax all types)	
Affects routines	XSLPnlpoptimize	
See also	XSLP_MIPALGORITHM, XSLP_MIPFIXSTEPBOUNDS	

XSLP_MULTISTART, MULTISTART

Description	The multistart master control. Defines if the multistart search is to be initiated, or if only the baseline model is to be solved.	
Type	Integer	
Category	Control	
Topic area	Multistart	
Values	-1	Depends on if any multistart jobs have been added.
	0	Multistart is off.
	1	Multistart is on.
Note	By default, the multistart search will always be initiated if multistart jobs have been added to the problem. The (original) base problem is not part of the multisearch job pool. To make it so, add an job with no extra settings (template job). It might be useful to load multiple template jobs, and customize them from callbacks.	
Default value	-1	
Note	FICO Xpress Global supports the solution of nonconvex problems to global optimality.	
Affects routines	XSLPminim, XSLPmaxim	
See also	XSLP_MULTISTART_MAXSOLVES, XSLP_MULTISTART_MAXTIME	

XSLP_MULTISTART_LOG, MULTISTART_LOG

Description	The level of logging during the multistart run.
Type	Integer
Category	Control
Topic areas	Multistart, Logging
Default value	0
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART,

XSLP_MULTISTART_MAXSOLVES, MULTISTART_MAXSOLVES

Description	The maximum number of jobs to create during the multistart search.
Type	Integer
Category	Control
Topic areas	Multistart, Limits
Note	This control can be increased on the fly during the multistart search: for example, if a job gets refused by a user callback, the callback may increase this limit to account for the rejected job.
Default value	-1 (no upper limit)
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART, XSLP_MULTISTART_MAXTIME

XSLP_MULTISTART_MAXTIME, MULTISTART_MAXTIME

Description	The maximum total time to be spent in the multistart search.
Type	Integer
Category	Control
Topic areas	Multistart, Limits
Note	XSLP_MAXTIME applies on a per job instance basis. There will be some time spent even after XSLP_MULTISTART_MAXTIME has elapsed, while the running jobs get terminated and their results collected.
Default value	0 (no upper limit)
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART, XSLP_MULTISTART_MAXSOLVES

XSLP_MULTISTART_POOLSIZE, MULTISTART_POOLSIZE

Description	The maximum number of problem objects allowed to pool up before synchronization in the deterministic multistart.
Type	Integer
Category	Control
Topic areas	Multistart, Limits
Default value	2
Note	Deterministic multistart is ensured by guaranteeing that the multistart solve results are evaluated in the same order every time. Solves that finish too soon can be pooled until all earlier started solves finish, allowing the system to start solving other multistart instances in the meantime on idle threads. Larger pool sizes will provide better speedups, but will require larger amounts of memory. Positive values are interpreted as a multiplier on the maximum number of active threads used, while negative values are interpreted as an absolute limit (and the absolute value is used). A value of zero will mean no result pooling.
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART, XSLP_DETERMINISTIC

XSLP_MULTISTART_SEED, MULTISTART_SEED

Description	Random seed used for the automatic generation of initial point when loading multistart presets
Type	Integer
Category	Control
Topic area	Multistart
Default value	0
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART

XSLP_NLPSOLVER, NLPSOLVER

Description	Controls whether to call FICO Xpress Global or one of the local solvers
Type	Integer
Category	Control
Topic area	Solution Process
Values	<div> <div>-1</div> <div>If the license allows and there are no user functions or multistart jobs, FICO Xpress Global will be called, otherwise a local solver.</div> </div> <div> <div>1</div> <div>The algorithm selected by XSLP_SOLVER will be used to find a locally optimal solution</div> </div> <div> <div>2</div> <div>FICO Xpress Global will be used to find a globally optimal solution</div> </div>
Default value	-1

Note	<p>The following constants are provided for setting this control:</p> <ul style="list-style-type: none"> -1 NLPSOLVER_AUTOMATIC 1 NLPSOLVER_LOCAL 2 NLPSOLVER_GLOBAL <p>Note that in cases where the problem can be reformulated to a convex MIQCQP, the Optimizer may be called independent of this control.</p>
See also	XSLP_SOLVER

XSLP_MULTISTART_THREADS, MULTISTART_THREADS

Description	The maximum number of threads to be used in multistart
Type	Integer
Category	Control
Topic areas	Multistart, Parallel
Default value	-1 (determined by XSLP_THREADS)
Note	The current hard upper limit on the number of threads to be used in multistart is 64.
Affects routines	XSLPminim, XSLPmaxim
See also	XSLP_MULTISTART XSLP_THREADS

XSLP_OCOUNT, SLPOCOUNT

Description	Number of SLP iterations over which to measure objective function variation for static objective (2) convergence criterion
Type	Integer
Category	Control
Topic areas	SLP, SLP-convergence
Note	<p>The static objective (2) convergence criterion does not measure convergence of individual variables. Instead, it measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables interacting with active constraints (those that have a marginal value of at least XSLP_MVTOL) have converged. The rationale is that if the remaining unconverged variables are not involved in active constraints and if the objective function is not changing significantly between iterations, then the solution is more-or-less practical.</p> <p>The variation in the objective function is defined as</p> $\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$ <p>where Iter is the XSLP_OCOUNT most recent SLP iterations and Obj is the corresponding objective function value.</p> <p>If $\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_OTOL_A}$ then the problem has converged on the absolute static objective (2) convergence criterion. The static objective function (2) test is applied only if XSLP_OCOUNT is at least 2.</p>
Default value	5
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_OTOL_A XSLP_OTOL_R

XSLP_PENALTYINFOSTART, SLPPENALTYINFOSTART

Description	Iteration from which to record row penalty information
Type	Integer
Category	Control
Topic areas	SLP, Logging
Note	Information about the size (current and total) of active penalties of each row and the number of times a penalty vector has been active is recorded starting at the SLP iteration number given by XSLP_PENALTYINFOSTART.
Default value	3

XSLP_POSTSOLVE, NLPPOSTSOLVE

Description	This control determines whether postsolving should be performed automatically
Type	Integer
Category	Control
Topic area	Presolve
Values	-1 Postsolve if the problem could be solved to optimality/infeasibility. 0 Do not automatically postsolve. 1 Postsolve automatically.
Default value	-1
See also	XSLP_PRESOLVE

XSLP_PRESOLVE, NLPPRESOLVE

Description	This control determines whether presolving should be performed on the nonlinear problem prior to starting the main algorithm
Type	Integer
Category	Control
Topic area	Presolve
Values	-1 Disable nonlinear presolve if and only if Optimizer presolve is disabled. 0 Disable nonlinear presolve. 1 Activate nonlinear presolve. 2 Low memory presolve. Original problem is not restored by postsolve and dual solution may not be completely postsolved.
Note	The Xpress NonLinear nonlinear presolve (which is carried out once, before augmentation in SLP or introducing auxiliaries in global) is independent of the Optimizer presolve (which is carried out during each SLP iteration or after introducing auxiliaries for a global solve).
Default value	-1
Affects routines	XSLPconstruct, XSLPpresolve
See also	XSLP_PRESOLVELEVEL, XSLP_PRESOLVEOPS, XSLP_REFORMULATE

XSLP_PRESOLVELEVEL, NLPPRESOLVELEVEL

Description	This control determines the level of changes presolve may carry out on the problem and whether column/row indices may change	
Type	Integer	
Category	Control	
Topic area	Presolve	
Values	1	Individual rows only presolve, no dropped columns/rows or index changes, no nonlinear transformations (XSLP_PRESOLVELEVEL_LOCALIZED).
	2	All linear presolve that does not drop columns/rows, no index changes, no nonlinear transformations (XSLP_PRESOLVELEVEL_BASIC).
	3	Full linear presolve including dropping columns/rows and index changes, no nonlinear transformations (XSLP_PRESOLVELEVEL_LINEAR).
	4	Full presolve (XSLP_PRESOLVELEVEL_FULL).
Note	XSLP_PRESOLVEOPS and XSLP_REFORMULATE controls the operations carried out in presolve. XSLP_PRESOLVELEVEL controls how those operations may change the problem. XSLP_PRESOLVE controls whether presolve is performed at all.	
Default value	XSLP_PRESOLVELEVEL_FULL	
Affects routines	XSLPconstruct, XSLPpresolve	
See also	XSLP_PRESOLVE, XSLP_PRESOLVEOPS, XSLP_REFORMULATE	

XSLP_PRESOLVEOPS, NLPPRESOLVEOPS

Description	Bitmap indicating the SLP presolve actions to be taken	
Type	Integer	
Category	Control	
Topic areas	Presolve, Bit-vector	
Values	Bit	Meaning
	0	Generic SLP presolve.
	1	Explicitly fix columns identified as fixed to zero.
	2	Explicitly fix all columns identified as fixed.
	3	SLP bound tightening.
	4	MISLP bound tightening.
	5	Bound tightening based on function domains.
	8	Do not presolve coefficients.
	9	Do not remove delta variables.
	10	Avoid reductions that can not be dual postsolved.
	11	Allow eliminations on determined variables.
	12	Avoid performing linear reductions at the nlp level.
	13	Avoid simplifying nonlinear expressions.

Note	The Xpress NonLinear nonlinear presolve (which is carried out once, before augmentation) is independent of the Optimizer presolve (which is carried out during each SLP iteration). Linear reductions are performed according to XPRS_PRESOLVEOPS if bit 12 is not set.
Default value	2104
Affects routines	XSLPconstruct, XSLPpresolve
See also	XSLP_PRESOLVELEVEL, XSLP_PRESOLVE, XSLP_PRESOLVEOPS, XSLP_REFORMULATE

XSLP_PROBING, NLPPROBING

Description	This control determines whether probing on a subset of variables should be performed prior to starting the main algorithm. Probing runs multiple times bound reduction in order to further tighten the bounding box.
Type	Integer
Category	Control
Topic area	Presolve
Values	<ul style="list-style-type: none"> -1 Automatic. 0 Disable SLP probing. 1 Activate SLP probing only on binary variables. 2 Activate SLP probing only on binary or unbounded integer variables. 3 Activate SLP probing only on binary or integer variables. 4 Activate SLP probing only on binary, integer variables, and unbounded continuous variables. 5 Activate SLP probing on any variable.
Default value	-1
Note	The Xpress NonLinear nonlinear probing, which is carried out once, is independent of the Optimizer presolve (which is carried out during each SLP iteration). The probing level allows for probing on an expanding set of variables, allowing for probing on all variables (level 5) or only those for which probing is more likely to be useful (binary variables).
Affects routines	XSLPpresolve
See also	XSLP_PRESOLVEOPS,

XSLP_REFORMULATE, NLPREFORMULATE

Description	Controls the problem reformulations carried out before augmentation. This allows SLP to take advantage of dedicated algorithms for special problem classes.
Type	Integer
Category	Control
Topic area	Presolve

Values	Bit	Meaning
	0	Solve convex quadratic objectives using the XPRS library .
	1	Convert non-convex quadratic objectives to SLP constructs .
	2	Solve convex quadratic constraints using the XPRS library.
	3	Convert non-convex QCQP constraints to SLP constructs.
	4	Keep second order cones in the XPRS problem to keep them in the linearizations.
	5	Convexity of a quadratic only problem may be checked by calling the optimizer to solve the instance.
	6	Convert pieewise linear functions to MIP constructs.
	7	Convert ABS functions to MIP constraints if the full problem can be made not nonlinear.
	8	Convert MIN and MAX functions to MIP expressions if the full problem can be made not nonlinear.
	9	Always convert ABS expressions.
	10	Always convert MIN and MAX expressions.

Default value 511 (bits 0 — 8 incl. are set)

Note The reformulation is part of XSLP presolve, and is only carried out if XSLP_PRESOLVE is nonzero. The following constants are provided for setting these bits:

Setting bit 0	XSLP_REFORMULATE_SLP2QP
Setting bit 1	XSLP_REFORMULATE_QP2SLP
Setting bit 2	XSLP_REFORMULATE_SLP2QCQP
Setting bit 3	XSLP_REFORMULATE_QCQP2SLP
Setting bit 4	XSLP_REFORMULATE_SOCP2SLP
Setting bit 5	XSLP_REFORMULATE_QPSOLVE
Setting bit 6	XSLP_REFORMULATE_PWL
Setting bit 7	XSLP_REFORMULATE_ABS
Setting bit 8	XSLP_REFORMULATE_MINMAX
Setting bit 9	XSLP_REFORMULATE_ALLABS
Setting bit 10	XSLP_REFORMULATE_ALLMINMAX

Affects routines XSLPconstruct, XSLPminim, XSLPmaxim, XSLPminim, XSLPmaxim, XSLPnloptimize

XSLP_SAMECOUNT, SLPSAMECOUNT

Description	Number of steps reaching the step bound in the same direction before step bounds are increased
Type	Integer
Category	Control
Topic area	SLP
Note	If step bounding is enabled, the step bound for a variable will be increased if successive changes are in the same direction. More precisely, if there are XSLP_SAMECOUNT successive changes reaching the step bound and in the same direction for a variable, then the step bound (B) for the variable will be reset to $B * XSLP_EXPAND$.
Default value	3
Affects routines	XSLPmaxim, XSLPminim

See also XSLP_EXPAND

XSLP_SAMEDAMP, SLPSAMEDAMP

Description	Number of steps in same direction before damping factor is increased
Type	Integer
Category	Control
Topic area	SLP
Note	If dynamic damping is enabled, the damping factor for a variable will be increased if successive changes are in the same direction. More precisely, if there are XSLP_SAMEDAMP successive changes in the same direction for a variable, then the damping factor (D) for the variable will be reset to $D * XSLP_DAMPEXPAND + XSLP_DAMPMAX * (1 - XSLP_DAMPEXPAND)$
Default value	3
See also	Xpress-SLP Solution Process, XSLP_ALGORITHM, XSLP_DAMP, XSLP_DAMPMAX
Affects routines	XSLPmaxim, XSLPminim

XSLP_SBROWOFFSET, SLPSBROWOFFSET

Description	Position of first character of SLP variable name used to create name of SLP lower and upper step bound rows
Type	Integer
Category	Control
Topic areas	SLP, Logging
Note	During augmentation, a delta vector is created for each SLP variable. Step bounds are provided for each delta variable, either using explicit bounds, or by using rows to provide lower and upper bounds. If such rows are used, they are created with names derived from the corresponding SLP variable. Customized naming is possible using XSLP_SBLOROWFORMAT and XSLP_SBUPROWFORMAT to define a format and XSLP_SBROWOFFSET to define the first character (counting from zero) of the variable name to be used.
Default value	0
Affects routines	XSLPconstruct
See also	XSLP_SBLOROWFORMAT, XSLP_SBUPROWFORMAT

XSLP_SBSTART, SLPSBSTART

Description	SLP iteration after which step bounds are first applied
Type	Integer
Category	Control
Topic area	SLP

Note If step bounds are used, they can be applied for the whole of the SLP optimization process, or started after a number of SLP iterations. In general, it is better not to apply step bounds from the start unless one of the following applies:

- (1) the initial estimates are known to be good, and explicit values can be provided for initial step bounds on all variables; or
- (2) the problem is unbounded unless all variables are step-bounded.

Default value 8

Affects routines XSLPmaxim, XSLPminim

XSLP_SCALE

Description When to re-scale the SLP problem

Type Integer

Category Control

Topic areas SLP, Numerics

Values

0	No re-scaling.
1	Re-scale every SLP iteration up to XSLP_SCALECOUNT iterations after the end of barrier optimization.
2	Re-scale every SLP iteration up to XSLP_SCALECOUNT iterations in total.
3	Re-scale every SLP iteration until primal simplex is automatically invoked.
4	Re-scale every SLP iteration.
5	Re-scale every XSLP_SCALECOUNT SLP iterations.
6	Re-scale every XSLP_SCALECOUNT SLP iterations after the end of barrier optimization.

Note During the SLP optimization, matrix entries can change considerably in magnitude, even when the formulae in the coefficients are not very nonlinear. Re-scaling of the matrix can reduce numerical errors, but may increase the time taken to achieve convergence.

Default value 1

Affects routines XSLPmaxim, XSLPminim

See also XSLP_SCALECOUNT

XSLP_SCALECOUNT

Description Iteration limit used in determining when to re-scale the SLP matrix

Type Integer

Category Control

Topic areas SLP, Numerics

Notes If XSLP_SCALE is set to 1 or 2, then XSLP_SCALECOUNT determines the number of iterations (after the end of barrier optimization or in total) in which the matrix is automatically re-scaled.

Default value 0

Affects routines XSLPmaxim, XSLPminim

See also XSLP_SCALE

XSLP_SOLVER, LOCALSOLVER

Description	Selects the library to use for local solves
Type	Integer
Category	Control
Topic area	Solution Process
Values	-1 automatic selection, based on model characteristics and solver availability 0 use Xpress-SLP (always available) 1 use Knitro if available 2 use Xpress-Optimizer if possible (convex quadratic problems only)
Note	<p>The presence of Knitro is detected automatically. Knitro can be used to solve any problem loaded into XSLP, independently from how the problem was loaded. XSLP_SOLVER is set to automatic, XSLP will be selected if any SLP specific construct has been loaded (these are ignored if Knitro is selected manually).</p> <p>When solving problems to global optimality, the XSLP_SOLVER control is used to decide which local solver to call for reoptimizing NLP-infeasible solutions heuristically.</p>
Default value	-1
See also	XSLP_NLPSOLVER

XSLP_SLPLOG

Description	Frequency with which SLP status is printed
Type	Integer
Category	Control
Topic areas	SLP, Logging
Note	If XSLP_LOG is set to zero (minimal logging) then a nonzero value for XSLP_SLPLOG defines the frequency (in SLP iterations) when summary information is printed out.
Default value	1
Affects routines	XSLPnlpoptimize, XSLPmaxim, XSLPminim
See also	XSLP_LOG, XSLP_MIPLOG

XSLP_STOPOUTOFRANGE, NLPSTOPOUTOFRANGE

Description	Stop optimization and return error code if internal function argument is out of range
Type	Integer
Category	Control
Topic area	SLP
Note	If XSLP_STOPOUTOFRANGE is set to 1, then if an internal function receives an argument which is out of its allowable range (for example, LOG of a negative number), an error code is set and the optimization is terminated.

Default value	0
Affects routines	XSLPevaluatecoef, XSLPevaluateformula XSLPmaxim, XSLPminim

XSLP_THREADS, NLPTHREADS

Description	Default number of threads to be used
Type	Integer
Category	Control
Topic area	Parallel
Note	Overall thread control value, used to determine the number of threads used where parallel calculations are possible.
Default value	-1 (use XPRS_THREADS value)
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_CALCTHREADS, XSLP_MULTISTART_THREADS,

XSLP_THREADSAFEUSERFUNC, NLPTHREADSAFEUSERFUNC

Description	Defines if user functions are allowed to be called in parallel
Type	Integer
Category	Control
Topic areas	User Functions, Parallel
Note	Date and time printing can be useful for identifying slow procedures during the SLP optimization. Setting XSLP_TIMEPRINT to 1 prints times at additional points during the optimization.
Values	0 user function are not thread safe, and will not be called in parallel 1 user functions are thread safe, and may be called in parallel
Default value	0 (no parallel user function calls)
Affects routines	XSLPmaxim, XSLPminim

XSLP_TRACEMASKOPS, SLPTRACEMASKOPS

Description	Controls the information printed for XSLP_TRACEMASK. The order in which the information is printed is determined by the order of bits in XSLP_TRACEMASKOPS.
Type	Integer
Category	Control
Topic areas	SLP, Logging, Bit-vector

Values	Bit	Meaning
	0	The variable name is used as a mask, not as an exact fit.
	1	Use mask to trace rows.
	2	Use mask to trace columns.
	3	Use mask to trace cascaded SLP variables.
	4	Show row / column category.
	5	Trace slack values.
	6	Trace dual values.
	7	Trace row penalty multiplier.
	8	Trace variable values (as returned by the linearization).
	9	Trace reduced costs.
	10	Trace slp value (value used in linearization and cascaded).
	11	Trace step bounds.
	12	Trace convergence status.
	13	Trace line search.
Default value	-1	

Note The following constants are provided for setting these bits:

Setting bit 0	XSLP_TRACEMASK_GENERALFIT
Setting bit 1	XSLP_TRACEMASK_ROWS
Setting bit 2	XSLP_TRACEMASK_COLS
Setting bit 3	XSLP_TRACEMASK_CASCADE
Setting bit 4	XSLP_TRACEMASK_TYPE
Setting bit 5	XSLP_TRACEMASK_SLACK
Setting bit 6	XSLP_TRACEMASK_DUAL
Setting bit 7	XSLP_TRACEMASK_WEIGHT
Setting bit 8	XSLP_TRACEMASK_SOLUTION
Setting bit 9	XSLP_TRACEMASK_REDUCEDCOST
Setting bit 10	XSLP_TRACEMASK_SLPVALUE
Setting bit 11	XSLP_TRACEMASK_STEPBOUND
Setting bit 12	XSLP_TRACEMASK_CONVERGE
Setting bit 13	XSLP_TRACEMASK_LINESEARCH

Affects routines XSLPminim, XSLPmaxim, XSLPminim, XSLPmaxim, XSLPnloptimize

XSLP_UNFINISHEDLIMIT, SLPUNFINISHEDLIMIT

Description	The number of consecutive SLP iterations that may have an unfinished status before the solve is terminated.
Type	Integer
Category	Control
Topic areas	SLP, Linearizations
Note	If the optimization of the current linear approximation terminates with an "unfinished" status, then first a number of strategies are applied to attempt a successful solve of the same linearization. If this fails, then a new iteration is started to change the linearization itself. This control limits the number of such repeated attempts.
Default value	3
Affects routines	XSLPnloptimize, XSLPmaxim, XSLPminim

XSLP_UPDATEOFFSET, SLPUPDATEOFFSET

Description	Position of first character of SLP variable name used to create name of SLP update row
Type	Integer
Category	Control
Topic areas	SLP, Logging
Note	During augmentation, one or more delta vectors are created for each SLP variable. The values of these are linked to that of the variable through an <i>update row</i> which is created as part of the augmentation procedure. Update rows are created with names derived from the corresponding SLP variable. Customized naming is possible using XSLP_UPDATEFORMAT to define a format and XSLP_UPDATEOFFSET to define the first character (counting from zero) of the variable name to be used.
Default value	0
Affects routines	XSLPconstruct
See also	XSLP_UPDATEFORMAT

XSLP_VCOUNT, SLPVCOUNT

Description	Number of SLP iterations over which to measure static objective (3) convergence
Type	Integer
Category	Control
Topic areas	SLP, SLP-convergence
Note	<p>The static objective (3) convergence criterion does not measure convergence of individual variables, and in fact does not in any way imply that the solution has converged. However, it is sometimes useful to be able to terminate an optimization once the objective function appears to have stabilized. One example is where a set of possible schedules are being evaluated and initially only a good estimate of the likely objective function value is required, to eliminate the worst candidates.</p> <p>The variation in the objective function is defined as</p> $\delta Obj = MAX_{Iter}(Obj) - MIN_{Iter}(Obj)$ <p>where <i>Iter</i> is the XSLP_VCOUNT most recent SLP iterations and <i>Obj</i> is the corresponding objective function value.</p> <p>If $ABS(\delta Obj) \leq XSLP_VTOL_A$ then the problem has converged on the absolute static objective function (3) criterion.</p> <p>The static objective function (3) test is applied only if after at least XSLP_VLIMIT + XSLP_SBSTART SLP iterations have taken place and only if XSLP_VCOUNT is at least 2. Where step bounding is being used, this ensures that the test is not applied until after step bounding has been introduced.</p>
Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_SBSTART, XSLP_VLIMIT, XSLP_VTOL_A, XSLP_VTOL_R

XSLP_VLIMIT, SLPVLIMIT

Description	Number of SLP iterations after which static objective (3) convergence testing starts
Type	Integer
Category	Control
Topic areas	SLP, SLP-convergence
Note	<p>The static objective (3) convergence criterion does not measure convergence of individual variables, and in fact does not in any way imply that the solution has converged. However, it is sometimes useful to be able to terminate an optimization once the objective function appears to have stabilized. One example is where a set of possible schedules are being evaluated and initially only a good estimate of the likely objective function value is required, to eliminate the worst candidates.</p> <p>The variation in the objective function is defined as</p> $\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$ <p>where Iter is the XSLP_VCOUNT most recent SLP iterations and Obj is the corresponding objective function value.</p> <p>If $\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_VTOL_A}$ then the problem has converged on the absolute static objective function (3) criterion. The static objective function (3) test is applied only after at least XSLP_VLIMIT + XSLP_SBSTART SLP iterations have taken place and only if XSLP_VCOUNT is at least 2. Where step bounding is being used, this ensures that the test is not applied until after step bounding has been introduced.</p>
Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_SBSTART, XSLP_VCOUNT, XSLP_VTOL_A, XSLP_VTOL_R

XSLP_WCOUNT, SLPWCOUNT

Description	Number of SLP iterations over which to measure the objective for the extended convergence continuation criterion
Type	Integer
Category	Control
Topic areas	SLP, SLP-convergence
Note	<p>It may happen that all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. This means that, at least in the linearization, if the variable were to be allowed to move further the objective function would improve. This does not necessarily imply that the same is true of the original problem, but it is still possible that an improved result could be obtained by taking another SLP iteration.</p> <p>The extended convergence continuation criterion is applied after a converged solution has been found where at least one variable has converged on extended criteria and is at its step bound limit. The extended convergence continuation test measures whether any improvement is being achieved when additional SLP iterations are carried out. If not, then the last converged solution will be restored and the optimization will stop.</p>

For a maximization problem, the improvement in the objective function at the current iteration compared to the objective function at the last converged solution is given by:

$$\delta\text{Obj} = \text{Obj} - \text{LastConvergedObj}$$

For a minimization problem, the sign is reversed.

If $\delta\text{Obj} > \text{XSLP_WTOL_A}$ and

$\delta\text{Obj} > \text{ABS}(\text{ConvergedObj}) * \text{XSLP_WTOL_R}$ then the solution is deemed to have a significantly better objective function value than the converged solution.

When a solution is found which converges on extended criteria and with active step bounds, the solution is saved and SLP optimization continues until one of the following:

- (1) a new solution is found which converges on some other criterion, in which case the SLP optimization stops with this new solution;
- (2) a new solution is found which converges on extended criteria and with active step bounds, and which has a significantly better objective function, in which case this is taken as the new saved solution;
- (3) none of the XSLP_WCOUNT most recent SLP iterations has a significantly better objective function than the saved solution, in which case the saved solution is restored and the SLP optimization stops.

If XSLP_WCOUNT is zero, then the extended convergence continuation criterion is disabled.

Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_WTOL_A, XSLP_WTOL_R

XSLP_XCOUNT, SLPXCOUNT

Description Number of SLP iterations over which to measure static objective (1) convergence

Type Integer

Category Control

Topic areas SLP, SLP-convergence

Note It may happen that all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. This means that, at least in the linearization, if the variable were to be allowed to move further the objective function would improve. This does not necessarily imply that the same is true of the original problem, but it is still possible that an improved result could be obtained by taking another SLP iteration. However, if the objective function has already been stable for several SLP iterations, then there is less likelihood of an improved result, and the converged solution can be accepted.

The static objective function (1) test measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. Because all the variables have converged, the solution is already converged but the fact that some variables are at their step bound limit suggests that the objective function could be improved by going further.

The variation in the objective function is defined as

$$\delta\text{Obj} = \text{MAX}_{\text{Iter}}(\text{Obj}) - \text{MIN}_{\text{Iter}}(\text{Obj})$$

where Iter is the XSLP_XCOUNT most recent SLP iterations and Obj is the corresponding objective function value.

If $\text{ABS}(\delta\text{Obj}) \leq \text{XSLP_XTOL_A}$

then the objective function is deemed to be static according to the absolute static objective function (1) criterion.

If $ABS(\delta Obj) \leq AVG_{Iter}(Obj) * XSLP_XTOL_R$

then the objective function is deemed to be static according to the relative static objective function (1) criterion.

The static objective function (1) test is applied only until `XSLP_XLIMIT` SLP iterations have taken place. After that, if all the variables have converged on strict or extended criteria, the solution is deemed to have converged.

If the objective function passes the relative or absolute static objective function (1) test then the solution is deemed to have converged.

Default value	5
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_XLIMIT, XSLP_XTOL_A, XSLP_XTOL_R

XSLP_XLIMIT, SLPXLIMIT

Description	Number of SLP iterations up to which static objective (1) convergence testing is performed
Type	Integer
Category	Control
Topic areas	SLP, SLP-convergence, Limits

Note It may happen that all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. This means that, at least in the linearization, if the variable were to be allowed to move further the objective function would improve. This does not necessarily imply that the same is true of the original problem, but it is still possible that an improved result could be obtained by taking another SLP iteration. However, if the objective function has already been stable for several SLP iterations, then there is less likelihood of an improved result, and the converged solution can be accepted.

The static objective function (1) test measures the significance of the changes in the objective function over recent SLP iterations. It is applied when all the variables have converged, but some have converged on extended criteria and at least one of these variables is at its step bound. Because all the variables have converged, the solution is already converged but the fact that some variables are at their step bound limit suggests that the objective function could be improved by going further.

The variation in the objective function is defined as

$$\delta Obj = MAX_{Iter}(Obj) - MIN_{Iter}(Obj)$$

where `Iter` is the `XSLP_XCOUNT` most recent SLP iterations and `Obj` is the corresponding objective function value.

If $ABS(\delta Obj) \leq XSLP_XTOL_A$

then the objective function is deemed to be static according to the absolute static objective function (1) criterion.

If $ABS(\delta Obj) \leq AVG_{Iter}(Obj) * XSLP_XTOL_R$

then the objective function is deemed to be static according to the relative static objective function (1) criterion.

The static objective function (1) test is applied only until `XSLP_XLIMIT` SLP iterations have taken place. After that, if all the variables have converged on strict or extended criteria, the solution is deemed to have converged.

If the objective function passes the relative or absolute static objective function (1) test then the solution is deemed to have converged.

Default value	100
----------------------	-----

Affects routines XSLPmaxim, XSLPminim

See also XSLP_XCOUNT, XSLP_XTOL_A, XSLP_XTOL_R

XSLP_ZEROCRITERION, SLPZEROCRITERION

Description Bitmap determining the behavior of the placeholder deletion procedure

Type Integer

Category Control

Topic areas SLP, Bit-vector

Values	Bit	Meaning
	0	(=1) Remove placeholders in nonbasic SLP variables
	1	(=2) Remove placeholders in nonbasic delta variables
	2	(=4) Remove placeholders in a basic SLP variable if its update row is nonbasic
	3	(=8) Remove placeholders in a basic delta variable if its update row is nonbasic and the corresponding SLP variable is nonbasic
	4	(=16) Remove placeholders in a basic delta variable if the determining row for the corresponding SLP variable is nonbasic
	5	(=32) Print information about zero placeholders

Note For an explanation of deletion of placeholder entries in the matrix see *Management of zero placeholder entries*.

The following constants are provided for setting these bits:

Setting bit 0	XSLP_ZEROCRITERION_NBSLPVAR
Setting bit 1	XSLP_ZEROCRITERION_NBDELTA
Setting bit 2	XSLP_ZEROCRITERION_SLPVARNBUPDATEROW
Setting bit 3	XSLP_ZEROCRITERION_DELTANBUPSATEROW
Setting bit 4	XSLP_ZEROCRITERION_DELTANBDRROW
Setting bit 5	XSLP_ZEROCRITERION_PRINT

Default value 0

Affects routines XSLPmaxim, XSLPminim

See also XSLP_ZEROCRITERIONCOUNT, XSLP_ZEROCRITERIONSTART, *Management of zero placeholder entries*

XSLP_ZEROCRITERIONCOUNT, SLPZEROCRITERIONCOUNT

Description Number of consecutive times a placeholder entry is zero before being considered for deletion

Type Integer

Category Control

Topic areas SLP, Limits

Note For an explanation of deletion of placeholder entries in the matrix see *Management of zero placeholder entries*.

Default value 0

Affects routines XSLPmaxim, XSLPminim

See also XSLP_ZEROCRITERION, XSLP_ZEROCRITERIONSTART, *Management of zero placeholder entries*

XSLP_ZEROCRITERIONSTART, SLPZEROCRITERIONSTART

Description	SLP iteration at which criteria for deletion of placeholder entries are first activated.
Type	Integer
Category	Control
Topic area	SLP
Note	For an explanation of deletion of placeholder entries in the matrix see <i>Management of zero placeholder entries</i> .
Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ZEROCRITERION, XSLP_ZEROCRITERIONCOUNT, <i>Management of zero placeholder entries</i>

20.3 String control parameters

XSLP_DELTAFORMAT, SLPDELTAFORMAT

Description	Formatting string for creation of names for SLP delta vectors
Type	String
Category	Control
Topic areas	SLP, Logging
Note	This control can be used to create a specific naming structure for delta vectors. The structure follows the normal C-style printf form, and can contain printing characters plus one %s string. This will be replaced by sequential characters from the name of the variable starting at position XSLP_DELTAOFFSET.
Default value	pD_%s where p is a unique prefix for names in the current problem
Affects routines	XSLPconstruct
See also	XSLP_DELTAOFFSET

XSLP_ITERFALLBACKOPS, SLPITERFALLBACKOPS

Description	Alternative LP level control values for numerically challenging problems
Type	String
Category	Control
Topic areas	SLP, Linearizations, Numerics
Notes	<p>When set, this control provides alternative ways of solving a linearization called adaptive iteration solves. This can be useful for numerically challenging problems that either solve to a non-satisfactory accuracy (relative to XSLP_FEASTOLTARGET) with the default solves, or that can incorrectly report infeasibility or unboundedness. In such cases, the solve will try the controls listed by XSLP_ITERFALLBACKOPS until a satisfactory solution is found or all options are exhausted.</p> <p>The individual controls for each solve are separated by a comma (','), while the set of controls for an attempt by a colon (':'). Example: 'XPRS_DEFAULTALG=3 : XPRS_BARORDER = 2, XPRS_PRESOVLE = 0' will try primal in one solve, and the homogenous barrier with presolve turned off in another. Optimizer flags are not inherited by the solve, so use XPRS_DEFAULTALG for selecting an LP solver to use.</p> <p>The resulting LP solves are carried out in a parallel manner, using XSLP_MULTISTART_THREADS number of threads.</p> <p>The result of the adaptive solves are always deterministic.</p> <p>Once a satisfactory solution is found, remaining solves are progressed only as far as necessary to guarantee determinism, so it is beneficial to list more promising control sets first.</p>
Default value	none
Affects routines	XSLPminim, XSLPmaxim, XSLPnlpoptimize
See also	XSLP_FEASTOLTARGET,

XSLP_IVNAME, NLPIVNAME

Description	Name of the set of initial values to be used
Type	String
Category	Control
Topic area	File IO
Notes	<p>This variable may be required for input from a file using <code>XSLPreadprob</code> if there is more than one set of initial values in the file. If no name is set, then the first set of initial values will be used, and the name will be set accordingly.</p> <p>This variable may also be required for output using <code>XSLPwriteprob</code> where initial values are included in the problem. If it is not set, then a default name will be used.</p>
Set by routines	<code>XSLPreadprob</code>
Default value	none
Affects routines	<code>XSLPreadprob</code> , <code>XSLPwriteprob</code>
See also	<code>XSLP_SBNAME</code> , <code>XSLP_TOLNAME</code>

XSLP_MINUSDELTAFORMAT, SLPMINUSDELTAFORMAT

Description	Formatting string for creation of names for SLP negative penalty delta vectors
Type	String
Category	Control
Topic areas	SLP, Logging
Note	<p>This control can be used to create a specific naming structure for negative penalty delta vectors. The structure follows the normal C-style <code>printf</code> form, and can contain printing characters plus one <code>%s</code> string. This will be replaced by sequential characters from the name of the variable starting at position <code>XSLP_DELTAOFFSET</code>.</p>
Default value	<p><code>pD-%s</code> where <code>p</code> is a unique prefix for names in the current problem</p>
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_DELTAOFFSET</code>

XSLP_MINUSERERRORFORMAT, SLPMINUSERERRORFORMAT

Description	Formatting string for creation of names for SLP negative penalty error vectors
Type	String
Category	Control
Topic areas	SLP, Logging
Note	<p>This control can be used to create a specific naming structure for negative penalty error vectors. The structure follows the normal C-style <code>printf</code> form, and can contain printing characters plus one <code>%s</code> string. This will be replaced by sequential characters from the name of the variable starting at position <code>XSLP_ERROROFFSET</code>.</p>

Default value	pE-%s where p is a unique prefix for names in the current problem
Affects routines	XSLPconstruct
See also	XSLP_ERROROFFSET

XSLP_PENALTYCOLFORMAT, SLPPENALTYCOLFORMAT

Description	Formatting string for creation of the names of the SLP penalty transfer vectors
Type	String
Category	Control
Topic areas	SLP, Logging
Note	This control can be used to create a specific naming structure for the penalty transfer vectors which transfer penalty costs into the objective. The structure follows the normal C-style printf form, and can contain printing characters plus one %s string. This will be replaced by "DELTA" for the penalty delta transfer vector and "ERR" for the penalty error transfer vector.
Default value	pPC-%s where p is a unique prefix for names in the current problem
Affects routines	XSLPconstruct

XSLP_PENALTYROWFORMAT, SLPPENALTYROWFORMAT

Description	Formatting string for creation of the names of the SLP penalty rows
Type	String
Category	Control
Topic areas	SLP, Logging
Note	This control can be used to create a specific naming structure for the penalty rows which total the penalty costs for the objective. The structure follows the normal C-style printf form, and can contain printing characters plus one %s string. This will be replaced by "DELTA" for the penalty delta row and "ERR" for the penalty error row.
Default value	pPR-%s where p is a unique prefix for names in the current problem
Affects routines	XSLPconstruct

XSLP_PLUSDELTAFORMAT, SLPPLUSDELTAFORMAT

Description	Formatting string for creation of names for SLP positive penalty delta vectors
Type	String
Category	Control
Topic areas	SLP, Logging

Note	This control can be used to create a specific naming structure for positive penalty delta vectors. The structure follows the normal C-style printf form, and can contain printing characters plus one %s string. This will be replaced by sequential characters from the name of the variable starting at position XSLP_DELTAOFFSET.
Default value	pD+%s where p is a unique prefix for names in the current problem
Affects routines	XSLPconstruct
See also	XSLP_DELTAOFFSET

XSLP_PLUSERERRORFORMAT, SLPPLUSERERRORFORMAT

Description	Formatting string for creation of names for SLP positive penalty error vectors
Type	String
Category	Control
Topic areas	SLP, Logging
Note	This control can be used to create a specific naming structure for positive penalty error vectors. The structure follows the normal C-style printf form, and can contain printing characters plus one %s string. This will be replaced by sequential characters from the name of the variable starting at position XSLP_ERROROFFSET.
Default value	pE+%s where p is a unique prefix for names in the current problem
Affects routines	XSLPconstruct
See also	XSLP_ERROROFFSET

XSLP_SBLOROWFORMAT, SLPSBLOROWFORMAT

Description	Formatting string for creation of names for SLP lower step bound rows
Type	String
Category	Control
Topic areas	SLP, Logging
Note	This control can be used to create a specific naming structure for lower limits on step bounds modeled as rows. The structure follows the normal C-style printf form, and can contain printing characters plus one %s string. This will be replaced by sequential characters from the name of the variable starting at position XSLP_SBROWOFFSET.
Default value	pSB-%s where p is a unique prefix for names in the current problem
Affects routines	XSLPconstruct
See also	XSLP_SBROWOFFSET

XSLP_SBNAME, SLPSBNAME

Description	Name of the set of initial step bounds to be used
Type	String
Category	Control
Topic areas	SLP, Logging
Notes	<p>This variable may be required for input from a file using <code>XSLPreadprob</code> if there is more than one set of initial step bounds in the file. If no name is set, then the first set of initial step bounds will be used, and the name will be set accordingly.</p> <p>This variable may also be required for output using <code>XSLPwriteprob</code> where initial step bounds are included in the problem. If it is not set, then a default name will be used.</p>
Set by routines	<code>XSLPreadprob</code>
Default value	none
Affects routines	<code>XSLPreadprob</code> , <code>XSLPwriteprob</code>
See also	<code>XSLP_IVNAME</code> , <code>XSLP_TOLNAME</code>

XSLP_SBUPROWFORMAT, SLPSBUPROWFORMAT

Description	Formatting string for creation of names for SLP upper step bound rows
Type	String
Category	Control
Topic areas	SLP, Logging
Note	<p>This control can be used to create a specific naming structure for upper limits on step bounds modeled as rows. The structure follows the normal C-style <code>printf</code> form, and can contain printing characters plus one <code>%s</code> string. This will be replaced by sequential characters from the name of the variable starting at position <code>XSLP_SBROWOFFSET</code>.</p>
Default value	<p><code>pSB+%s</code> where <code>p</code> is a unique prefix for names in the current problem</p>
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_SBROWOFFSET</code>

XSLP_TOLNAME, SLPTOLNAME

Description	Name of the set of tolerance sets to be used
Type	String
Category	Control
Topic areas	SLP, File IO

Notes	This variable may be required for input from a file using <code>XSLPreadprob</code> if there is more than one set of tolerance sets in the file. If no name is set, then the first set of tolerance sets will be used, and the name will be set accordingly. This variable may also be required for output using <code>XSLPwriteprob</code> where tolerance sets are included in the problem. If it is not set, then a default name will be used.
Set by routines	<code>XSLPreadprob</code>
Default value	none
Affects routines	<code>XSLPreadprob</code> , <code>XSLPwriteprob</code>
See also	<code>XSLP_IVNAME</code> , <code>XSLP_SBNAME</code>

XSLP_TRACEMASK, SLPTRACEMASK

Description	Mask of variable or row names that are to be traced through the SLP iterates
Type	String
Category	Control
Topic areas	SLP, Logging
Notes	If the mask is nonempty, variables and rows matching the mask are listed after each SLP iteration and each cascade, allowing for a convenient means to observe how certain variables change through the iterates. This feature is provided for tuning and model debugging purposes. The actual information printed is controlled by <code>XSLP_TRACEMASKOPS</code> . The string in the tracemask may contain several variable or row names, separated by a whitespace. Wildcards may also be used.
Default value	none: no tracing
Affects routines	<code>XSLPminim</code> , <code>XSLPmaxim</code> , <code>XSLPpreminim</code> , <code>XSLPpremaxim</code> , <code>XSLPnlpoptimize</code> ,
See also	<code>XSLP_TRACEMASKOPS</code>

XSLP_UPDATEFORMAT, SLPUPDATEFORMAT

Description	Formatting string for creation of names for SLP update rows
Type	String
Category	Control
Topic areas	SLP, Logging
Note	This control can be used to create a specific naming structure for update rows. The structure follows the normal C-style printf form, and can contain printing characters plus one %s string. This will be replaced by sequential characters from the name of the variable starting at position <code>XSLP_UPDATEOFFSET</code> .
Default value	<code>pU_%s</code> where p is a unique prefix for names in the current problem
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_UPDATEOFFSET</code>

20.4 Knitro controls

All Knitro controls are available with an 'X' pre-tag. For example the Knitro integer control 'KTR_PARAM_ALGORITHM' can be set using XSLPsetintcontrol using the control ID defined as 'XKTR_PARAM_ALGORITHM'. Please refer to the Xpress Knitro manual for the description of the Knitro controls.

CHAPTER 21

Library functions and the programming interface

21.1 Counting

All Xpress NonLinear entities are numbered from 1. The 0th item is defined, and is an empty entity of the appropriate type. Therefore, whenever an Xpress NonLinear function returns a zero value, it means that there is no data of that type.

In parsed and unparsed function arrays, types XSLP_COL and XSLP_ROW use indices counted from zero, the same as the Xpress Optimizer library.

21.2 The Xpress NonLinear problem pointer

Xpress NonLinear uses the same concept as the Optimizer library, with a "pointer to a problem". The optimizer problem must be initialized first in the normal way. Then the corresponding Xpress NonLinear problem must be initialized, including a pointer to the underlying optimizer problem. For example:

```
{
...
XPRSprob prob=NULL;
XSLPprob SLPprob=NULL;

XPRSinit("");
XSLPinit();
XPRScreateprob(&prob);
XSLPcreateprob(&SLPprob, &prob);
...
}
```

At the end of the program, the Xpress NonLinear problem should be destroyed. You are responsible for destroying the underlying XPRSprob linear problem afterwards. For example:

```
{
...
XSLPdestroyprob(SLPprob);
XPRSdestroyprob(prob);
XSLPfree();
XPRSfree();
...
}
```

The following functions are provided to manage Xpress NonLinear problems. See the documentation below on the individual functions for more details.

XSLPcopycontrols(XSLPprob prob1, XSLPprob prob2)

Copy the settings of control variables

```
XSLPcopycallbacks(XSLPprob probl, XSLPprob prob2)
```

Copy the callback settings

```
XSLPcopyprob(XSLPprob probl, XSLPprob prob2, char *ProbName)
```

Copy a problem completely

```
XSLPcreateprob(XSLPprob *probl, XPRSprob *prob2)
```

Create an Xpress NonLinear problem

```
XSLPdestroyprob(XSLPprob probl)
```

Delete an Xpress NonLinear problem from memory

```
XSLPrestore(XSLPprob probl)
```

Restore Xpress NonLinear data structures from file

```
XSLPsave(XSLPprob probl)
```

Save Xpress NonLinear data structures to file

21.3 The load functions

The load functions can be used to load an Xpress NonLinear problem directly into the Xpress data structures. Because there are so many additional items which can be loaded apart from the basic (linear) matrix, the loading process is divided into several functions.

The best practice is to load the linear part of the problem first, using the normal Optimizer Library functions `XPRSloadlp` or `XPRSloadmip`. Then the appropriate parts of the Xpress NonLinear problem can be loaded. After all the load functions have been called, `XSLPconstruct` should be called to create the SLP matrix and data structures. If `XSLPconstruct` is not invoked before a call to one of the Xpress NonLinear optimization routines, then it will be called by the optimization routine itself.

All of these functions initialize their data areas. Therefore, if a second call is made to the same function for the same problem, the previous data will be deleted. If you want to include additional data of the same type, then use the corresponding add function.

It is possible to remove parts of the SLP structures with the various `XSLPdel` functions, and `XSLPunconstruct` can also be used to remove the augmentation.

Xpress NonLinear is compatible with the Xpress quadratic programming optimizer. `XPRSloadqp` and `XPRSloadmiqp` can be used to load quadratic problems (or quadratically constrained problems using `XPRSloadqcqp` and `XPRSloadmiqcqp`). The quadratic objective will be optimized using the Xpress quadratic optimizer; the nonlinear constraints will be handled with the normal SLP procedures. Please note, that this separation is only useful for a convex quadratic objective and convex quadratic inequality constraints. All nonconvex quadratic matrices should be handled as SLP structures.

For a description on when it's more beneficial to use the XPRS library to solve QP or QCQP problems, please see *Selecting the right algorithm for a nonlinear problem - when to use the XPRS library instead of XSLP*.

21.4 Library functions

A large number of routines are available for Library users of Xpress NonLinear, 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. Library users have access to a set of functions providing advanced control over their program's interaction with the SLP module and catering for more complicated problem development. When called from the SLP library, these functions have an `XSLP` prefix and take an SLP problem as their first argument. They are also exported to the XPRS library, where they can be called on an `XPRSprob` without having to explicitly create an `XSLPprob` first.

<code>XPRSnlpchgformulastr</code>	Add or replace a single matrix formula using a character string for the formula. [Problem Modification]	p. 298
<code>XPRSnlpgetformulastr</code>	Retrieve a single matrix formula in a character string. [Problem Information]	p. 331
<code>XPRSslpchgcoefstr</code>	Add or change a single matrix coefficient using a character string for the formula. [Problem Modification, SLP]	p. 291
<code>XPRSslpgetcoefstr</code>	Retrieve a single matrix coefficient as a formula in a character string. [Problem Information, SLP]	p. 323
<code>XSLPaddcoefs, XPRSslpaddcoefs</code>	Add non-linear coefficients to the SLP problem. [Problem Modification, SLP]	p. 277
<code>XSLPaddformulas, XPRSnlpaddformulas</code>	Add non-linear formulas to the SLP problem. [Problem Modification]	p. 279
<code>XSLPadduserfunction, XPRSnlpadduserfunction</code>	Add user function definitions to an SLP problem. [User Functions]	p. 283
<code>XSLPcalclacks, XPRSnlpcalclacks</code>	Calculate the slack values for the provided solution in the non-linear problem [Solution]	p. 287
<code>XSLPcascade, XPRSslpcascade</code>	Re-calculate consistent values for SLP variables based on the current values of the remaining variables. [Cascading, SLP, Solution Process]	p. 288
<code>XSLPcascadeorder, XPRSslpcascadeorder</code>	Establish a re-calculation sequence for SLP variables with determining rows. [Data Input, SLP]	p. 289
<code>XSLPchgcascadenlimit, XPRSslpchgcascadenlimit</code>	Set a variable specific cascade iteration limit [Data Input, SLP]	p. 290
<code>XSLPchgcoef, XPRSslpchgcoef</code>	Add or change a single matrix coefficient using a parsed or unparsed formula. [Problem Modification, SLP]	p. 293
<code>XSLPchgdelatatype, XPRSslpchgdelatatype</code>	Changes the type of the delta assigned to a nonlinear variable [Problem Modification, SLP]	p. 295
<code>XSLPchgformula, XPRSnlpchgformula</code>	Add or replace a single matrix formula using a parsed or unparsed formula [Problem Modification]	p. 299
<code>XSLPchgrowstatus, XPRSslpchgrowstatus</code>	Change the status setting of a constraint [Bit-vector, Data Input, SLP]	p. 300
<code>XSLPchgrowwt, XPRSslpchgrowwt</code>	Set or change the initial penalty error weight for a row [Data Input, SLP]	p. 301
<code>XSLPconstruct, XPRSslpconstruct</code>	Create the full augmented SLP matrix and data structures, ready for optimization [SLP, Solution Process]	p. 306

XSLPcopycallbacks	Copy the user-defined callbacks from one SLP problem to another [Callback]	p. 307
XSLPcopycontrols	Copy the values of the control variables from one SLP problem to another [Controls and Attributes]	p. 308
XSLPcopyprob	Copy an existing SLP problem to another [Problem Creation]	p. 309
XSLPcreateprob	Create a new SLP problem [Problem Creation]	p. 310
XSLPdelcoefs, XPRSslpdelcoefs	Delete coefficients from the current problem. [Problem Modification, SLP]	p. 311
XSLPdelformulas, XPRSnlpdelformulas	Delete nonlinear formulas from the current problem [Problem Modification]	p. 312
XSLPdeluserfunction, XPRSnlpdeluserfunction	Delete a user function from the current problem [User Functions]	p. 314
XSLPdestroyprob	Delete an SLP problem and release all the associated memory [Problem Creation]	p. 316
XSLPevaluatecoef, XPRSslpevaluatecoef	Evaluate a coefficient using the current values of the variables [Solution]	p. 317
XSLPevaluateformula, XPRSnlpevaluateformula	Evaluate a formula using the current values of the variables [Solution]	p. 318
XSLPfixpenalties, XPRSslpfixpenalties	Fixe the values of the error vectors [Solution Process]	p. 319
XSLPfree	Free any memory allocated by Xpress NonLinear and close any open Xpress NonLinear files [Licensing]	p. 320
XSLPgetcoeffformula, XPRSslpgetcoeffformula	Retrieve a single matrix coefficient as a formula split into tokens. [Problem Information, SLP]	p. 324
XSLPgetcoefs, XPRSslpgetcoefs	Retrieve the list of positions of the nonlinear coefficients in the problem. [Problem Information, SLP]	p. 325
XSLPgetcolinfo, XPRSslpgetcolinfo	Get current column information. [SLP, Solution]	p. 326
XSLPgetdblattrib	Retrieve the value of a double precision problem attribute [Controls and Attributes]	p. 327
XSLPgetdblcontrol	Retrieve the value of a double precision problem control [Controls and Attributes]	p. 328
XSLPgetformula, XPRSnlpgetformula	Retrieve a single matrix formula as a formula split into tokens. [Problem Information]	p. 330

XSLPgetformularows, XPRSnlpgetformularows	Retrieve the list of positions of the nonlinear formulas in the problem [Problem Information]	p. 333
XSLPgetindex	Retrieve the index of an Xpress NonLinear entity with a given name [Problem Information, User Functions]	p. 334
XSLPgetintattrib	Retrieve the value of an integer problem attribute [Controls and Attributes]	p. 335
XSLPgetintcontrol	Retrieve the value of an integer problem control [Controls and Attributes]	p. 336
XSLPgetlasterror	Retrieve the error message corresponding to the last Xpress NonLinear error during an SLP run [Misc]	p. 337
XSLPgetptrattrib	Retrieve the value of a problem pointer attribute [Controls and Attributes]	p. 338
XSLPgetrowinfo, XPRSslpgetrowinfo	Get current row information. [SLP, Solution]	p. 339
XSLPgetrowstatus, XPRSslpgetrowstatus	Retrieve the status setting of a constraint [Bit-vector, SLP, Solution]	p. 340
XSLPgetrowwt, XPRSslpgetrowwt	Get the initial penalty error weight for a row [Data Information, SLP]	p. 341
XSLPgetstrattrib	Retrieve the value of a string problem attribute [Controls and Attributes]	p. 343
XSLPgetstrcontrol	Retrieve the value of a string problem control [Controls and Attributes]	p. 344
XSLPimportlibfunc, XPRSnlpimportlibfunc	Imports a function from a library file to be called as a user function [User Functions]	p. 348
XSLPinit	Initializes the Xpress NonLinear system [Licensing]	p. 349
XSLPinterrupt	Interrupts the current SLP optimization [Solution Process]	p. 350
XSLPitemname	Retrieves the name of an Xpress NonLinear entity or the value of a function token as a character string. [Names Manager]	p. 351
XSLPloadcoefs, XPRSslploadcoefs	Load non-linear coefficients into the SLP problem. [Problem Information, SLP]	p. 352
XSLPloadformulas, XPRSnlploadformulas	Load non-linear formulas into the SLP problem [Problem Information]	p. 355
XSLPmsaddcustompreset, XPRSmsaddcustompreset	A combined version of XSLPmsaddjob and XSLPmsaddpreset. [Data Input, Multistart]	p. 362
XSLPmsaddjob, XPRSmsaddjob	Adds a multistart job to the multistart pool [Data Input, Multistart]	p. 363

<code>XSLPmsaddpreset, XPRSmsaddpreset</code>	Loads a preset of jobs into the multistart job pool. [Data Input, Multistart]	p. 364
<code>XSLPmsclear, XPRSmsclear</code>	Removes all scheduled jobs from the multistart job pool [Multistart]	p. 365
<code>XSLPnlpoptimize, XPRSnlpoptimize</code>	Maximize or minimize an SLP problem [Solution Process]	p. 366
<code>XSLPpostsolve, XPRSnlppostsolve</code>	Restores the problem to its pre-solve state [Presolve]	p. 367
<code>XSLPpresolve</code>	Perform a nonlinear presolve on the problem [Presolve]	p. 368
<code>XSLPprntevalinfo, XPRSnlpprntevalinfo</code>	Print a summary of any evaluation errors that may have occurred during solving a problem [Logging]	p. 370
<code>XSLPprntmemory</code>	Print the dimensions and memory allocations for a problem [Logging]	p. 369
<code>XSLPreadprob</code>	Read an Xpress NonLinear extended MPS format matrix from a file into an SLP problem [File IO, Problem Creation]	p. 371
<code>XSLPreinitialize, XPRSslpreinitialize</code>	Reset the SLP problem to match a just augmented system [SLP, Solution Process]	p. 375
<code>XSLPrestore</code>	Restore the Xpress NonLinear problem from a file created by <code>XSLPsave</code> [File IO, Save Restore]	p. 374
<code>XSLPsave</code>	Save the Xpress NonLinear problem to file [File IO, Save Restore]	p. 376
<code>XSLPsaveas</code>	Save the Xpress NonLinear problem to a named file [File IO, Save Restore]	p. 377
<code>XSLPscaling</code>	Analyze the current matrix for largest/smallest coefficients and ratios [Numerics]	p. 378
<code>XSLPsetcbcascadeend, XPRSsetcbslpcascadeend</code>	Set a user callback to be called at the end of the cascading process, after the last variable has been cascaded [Callback, Cascading, SLP]	p. 379
<code>XSLPsetcbcascadestart, XPRSsetcbslpcascadestart</code>	Set a user callback to be called at the start of the cascading process, before any variables have been cascaded [Callback, Cascading, SLP]	p. 380
<code>XSLPsetcbcascadevar, XPRSsetcbslpcascadevar</code>	Set a user callback to be called after each column has been cascaded [Callback, Cascading, SLP]	p. 381
<code>XSLPsetcbcascadevarfail, XPRSsetcbslpcascadevarfail</code>	Set a user callback to be called after cascading a column was not successful [Callback, Cascading, SLP]	p. 382
<code>XSLPsetcbcofevalerror, XPRSsetcbnlpcfevalerror</code>	Set a user callback to be called when an evaluation of a coefficient fails during the solve [Callback, Numerics]	p. 383

<code>XSLPsetcbconstruct</code> , <code>XPRSsetcbslpconstruct</code>	Set a user callback to be called during the Xpress-SLP augmentation process [Callback, SLP]	p. 384
<code>XSLPsetcbdestroy</code>	Set a user callback to be called when an SLP problem is about to be destroyed [Callback]	p. 386
<code>XSLPsetcbdrcol</code> , <code>XPRSsetcbslpdrcol</code>	Set a user callback used to override the update of variables with small determining column [Callback, Cascading, SLP]	p. 387
<code>XSLPsetcbintsol</code> , <code>XPRSsetcbslpintsol</code>	Set a user callback to be called during MISLP when an integer solution is obtained [Callback, MISLP]	p. 388
<code>XSLPsetcbiterend</code> , <code>XPRSsetcbslpiterend</code>	Set a user callback to be called at the end of each SLP iteration [Callback, SLP]	p. 389
<code>XSLPsetcbiterstart</code> , <code>XPRSsetcbslpiterstart</code>	Set a user callback to be called at the start of each SLP iteration [Callback, SLP]	p. 390
<code>XSLPsetcbitervar</code> , <code>XPRSsetcbslpitervar</code>	Set a user callback to be called after each column has been tested for convergence [Callback, SLP, SLP-convergence]	p. 391
<code>XSLPsetcbmessage</code>	Set a user callback to be called whenever Xpress NonLinear outputs a line of text according to <code>XSLP_ECHOXPRSMESSAGES</code> . [Callback, Logging]	p. 392
<code>XSLPsetcbmsjobend</code> , <code>XPRSsetcbmsjobend</code>	Set a user callback to be called every time a new multistart job finishes. [Callback, Multistart]	p. 394
<code>XSLPsetcbmsjobstart</code> , <code>XPRSsetcbmsjobstart</code>	Set a user callback to be called every time a new multistart job is created, and the pre-loaded settings are applied [Callback, Multistart]	p. 395
<code>XSLPsetcbmswinner</code> , <code>XPRSsetcbmswinner</code>	Set a user callback to be called every time a multistart winner has been declared [Callback, Multistart]	p. 396
<code>XSLPsetcboptnode</code>	Set a user callback to be called during MISLP when an optimal SLP solution is obtained at a node [Callback, MISLP]	p. 397
<code>XSLPsetcbprenode</code>	Set a user callback to be called during MISLP after the set-up of the SLP problem to be solved at a node, but before SLP optimization [Callback, MISLP]	p. 398
<code>XSLPsetcbpresolved</code>	Set a user callback to be called after the nonlinear presolver has been applied. [Callback, SLP]	p. 401
<code>XSLPsetcbpreupdatelinearization</code> , <code>XPRSsetcbslppreupdatelinearization</code>	Set a user callback to be called before the linearization is updated [Callback, SLP]	p. 399

XSLPsetcbslpend	Set a user callback to be called at the end of the SLP optimization [Callback, SLP]	p. 402
XSLPsetcbslpnode	Set a user callback to be called during MISLP after the SLP optimization at each node. [Callback, MISLP]	p. 403
XSLPsetcbslpstart	Set a user callback to be called at the start of the SLP optimization [Callback, SLP]	p. 404
XSLPsetcurrentiv, XPRSnlpsetcurrentiv	Transfer the current solution to initial values [Data Input]	p. 405
XSLPsetdblcontrol	Set the value of a double precision problem control [Controls and Attributes]	p. 406
XSLPsetdefaultcontrol	Set the values of one SLP control to its default value [Controls and Attributes]	p. 407
XSLPsetdefaults	Set the values of all SLP controls to their default values [Controls and Attributes]	p. 408
XSLPsetdetrow, XPRSnlpsetdetrow	Set the determining row of a variable [Cascading, Data Input, SLP]	p. 409
XSLPsetfunctionerror, XPRSnlpsetfunctionerror	Set the function error flag for the problem [Misc]	p. 410
XSLPsetinitval, XPRSnlpsetinitval	Set the initial value of a variable [Data Input]	p. 411
XSLPsetintcontrol	Set the value of an integer problem control [Controls and Attributes]	p. 412
XSLPsetlogfile	Define an output file to be used to receive messages from Xpress NonLinear [File IO, Logging]	p. 413
XSLPsetparam	Set the value of a control parameter by name [Controls and Attributes]	p. 414
XSLPsetstrcontrol	Set the value of a string problem control [Controls and Attributes]	p. 415
XSLPunconstruct, XPRSnlpunconstruct	Removes the augmentation and returns the problem to its pre-linearization state [SLP, Solution Process]	p. 416
XSLPupdatelinearization, XPRSnlpupdatelinearization	Updates the current linearization [SLP, Solution Process]	p. 417
XSLPvalidate, XPRSnlpvalidate	Validate the feasibility of constraints in a converged solution [Solution]	p. 418
XSLPvalidatekkt, XPRSnlpvalidatekkt	Validates the first order optimality conditions also known as the Karush-Kuhn-Tucker (KKT) conditions versus the current solution [Solution]	p. 419

XSLPvalidateprob, XPRSnlpvalidateprob	Validates the current problem formulation and statement [Solution]	p. 420
XSLPvalidaterow, XPRSnlpvalidaterow	Prints an extensive analysis on a given constraint of the SLP problem [Solution]	p. 421
XSLPvalidatevector, XPRSnlpvalidatevector	Validate the feasibility of constraints for a given solution [Solution]	p. 422
XSLPwriteprob	Write the current problem to a file in extended MPS or text format [File IO, Problem Information]	p. 423
XSLPwriteslxsol	Write the current solution to an MPS like file format [File IO, Solution]	p. 424

XSLPaddcoefs, XPRSslpaddcoefs

Purpose

Add non-linear coefficients to the SLP problem. For a simpler version of this function see XSLPaddformulas.

Topic areas

SLP, Problem Modification

Synopsis

```
int XPRS_CC XSLPaddcoefs(XSLPprob prob, int ncoefs, const int[] rowind,
    const int[] colind, const double[] factor, const int[] formulastart,
    int parsed, const int[] type, const double[] value);
```

Arguments

prob	The current SLP problem.
ncoefs	Number of non-linear coefficients to be added.
rowind	Integer array holding index of row for the coefficient.
colind	Integer array holding index of column for the coefficient.
factor	Double array holding factor by which formula is scaled. If this is NULL, then a value of 1.0 will be used.
formulastart	Integer array of length ncoefs+1 holding the start position in the arrays type and value of the formula for the coefficients. formulastart[ncoefs] should be set to the next position after the end of the last formula.
parsed	Integer indicating whether the token arrays are formatted as internal unparsed (parsed=0) or internal parsed reverse Polish (parsed=1).
type	Array of token types providing the formula for each coefficient.
value	Array of values corresponding to the types in type.

Example

Assume that the rows and columns of prob are named Row1, Row2 ..., Col1, Col2 ... The following example adds coefficients representing:

Col2 * Col3 + Col6 * Col2^2 into Row1 and

Col2 ^ 2 into Row3.

```
int rowind[3], colind[3], formulastart[4], type[8];
int n, ncoefs;
double value[8];

rowind[0] = 1; colind[0] = 2;
rowind[1] = 1; colind[1] = 6;
rowind[2] = 3; colind[2] = 2;

n = ncoefs = 0;
formulastart[ncoefs++] = n;
type[n] = XSLP_COL; value[n++] = 3;
type[n++] = XSLP_EOF;

formulastart[ncoefs++] = n;
type[n] = XSLP_COL; value[n++] = 2;
type[n] = XSLP_CON; value[n++] = 2;
type[n] = XSLP_OP; value[n++] = XSLP_EXPONENT;
type[n++] = XSLP_EOF;

formulastart[ncoefs++] = n;
type[n] = XSLP_COL; value[n++] = 2;
type[n] = XSLP_CON; value[n++] = 2;
```

```

type[n] = XSLP_OP; value[n++] = XSLP_EXPONENT;
type[n++] = XSLP_EOF;

formulastart[ncoefs] = n;

XSLPaddcoefs(prob, ncoefs, rowind, colind,
             NULL, formulastart, 1 /* reversed Polish */, type, value);

```

The first coefficient in Row1 is in Col2 and has the formula Col3, so it represents Col2 * Col3.

The second coefficient in Row1 is in Col6 and has the formula Col2 * Col2 so it represents Col6 * Col2^2.

The formulae are described as *parsed* (parsed=1), so the formula is written as

Col2 Col2 *

rather than the unparsed form

Col2 * Col2

The last coefficient, in Row3, is in Col2 and has the formula Col2, so it represents Col2 * Col2.

Further information

The j^{th} coefficient is made up of two parts: *factor* and *Formula*. *factor* is a constant multiplier, which can be provided in the *factor* array. If Xpress NonLinear can identify a constant factor in *Formula*, then it will use that as well, to minimize the size of the formula which has to be calculated. *Formula* is made up of a list of tokens in *type* and *value* starting at *formulastart[j]*. The tokens follow the rules for *parsed* or *unparsed* formulae as indicated by the setting of *parsed*. The formula must be terminated with an *XSLP_EOF* token. If several coefficients share the same formula, they can have the same value in *formulastart*. For possible token types and values see the chapter on 13.

The add functions load additional items into the SLP problem. The corresponding load functions delete any existing items first.

The behaviour for existing coefficients is additive: the formula defined in the parameters are added to any existing formula coefficients. However, due to performance considerations, such duplications should be avoided when possible.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula,
XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPaddformulas, XPRSnlpaddformulas

Purpose

Add non-linear formulas to the SLP problem.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XSLPaddformulas(XSLPprob prob, int ncoefs, const int[] rowind,
    const int[] formulastart, int parsed, const int[] type, const
    double[] value);
```

Arguments

prob	The current SLP problem.
ncoefs	Number of non-linear coefficients to be added.
rowind	Integer array holding index of row for the coefficient.
formulastart	Integer array of length ncoefs+1 holding the start position in the arrays type and value of the formula for the coefficients. formulastart[ncoefs] should be set to the next position after the end of the last formula.
parsed	Integer indicating whether the token arrays are formatted as internal unparsed (parsed=0) or internal parsed reverse Polish (parsed=1).
type	Array of token types providing the formula for each coefficient.
value	Array of values corresponding to the types in type.

Example

Assume that the rows and columns of prob are named Row0, Row1 ..., Col0, Col1 ... The following example adds nonlinear formulas representing:

Col2 * Col3 * Col6^2 into Row1 and

Col2 * Col3 ^ 2 into Row3.

```
int rowind[3], formulastart[4], type[8];
int n, ncoefs;
double value[8];

rowind[0] = 1;
rowind[1] = 1;
rowind[2] = 3;

n = ncoefs = 0;
formulastart[ncoefs++] = n;
type[n] = XSLP_COL; value[n++] = 3;
type[n++] = XSLP_EOF;

formulastart[ncoefs++] = n;
type[n] = XSLP_COL; value[n++] = 2;
type[n] = XSLP_COL; value[n++] = 3;
type[n] = XSLP_OP; value[n++] = XSLP_MULTIPLY;
type[n] = XSLP_COL; value[n++] = 6;
type[n] = XSLP_CON; value[n++] = 2;
type[n] = XSLP_OP; value[n++] = XSLP_EXPONENT;
type[n] = XSLP_OP; value[n++] = XSLP_MULTIPLY;
type[n++] = XSLP_EOF;

formulastart[ncoefs++] = n;
type[n] = XSLP_COL; value[n++] = 2;
type[n] = XSLP_COL; value[n++] = 3;
type[n] = XSLP_CON; value[n++] = 2;
```

```

type[n] = XSLP_OP;  value[n++] = XSLP_EXPONENT;
type[n] = XSLP_OP;  value[n++] = XSLP_MULTIPLY;
type[n++] = XSLP_EOF;

formulastart[ncoefs] = n;

XSLPaddformulas(prob, ncoefs, rowind, formulastart, 1 /* reversed Polish */, type, value);

```

Further information

1. Formula is made up of a list of tokens in type and value starting at formulastart[j]. The tokens follow the rules for parsed or unparsed formulae as indicated by the setting of parsed. The formula must be terminated with an XSLP_EOF token. If several rows share the same nonlinear expression, they can have the same value in formulastart. For possible token types and values see the chapter on 13.
2. The add functions load additional items into the SLP problem. The corresponding load functions delete any existing items first.
3. The behaviour for existing formulas is additive: the formula defined in the parameters are added to any existing nonlinear expressions in the row. However, due to performance considerations, such duplications should be avoided when possible.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula,
XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPadddfs

Purpose

This subroutine is deprecated and will be removed in a future release.
Add a set of distribution factors.

Topic areas

SLP, Data Input

Synopsis

```
int XSLP_CC XSLPadddfs(XSLPprob prob, int ndf, const int *colind, const int
    *rowind, const double *value)
```

Arguments

prob	The current SLP problem.
ndf	The number of distribution factors.
colind	Array of indices of columns whose distribution factor is to be changed.
rowind	Array of indices of the rows where each distribution factor applies.
value	Array of double precision variables holding the new values of the distribution factors.

Example

The following example adds distribution factors as follows:

column 282 in row 134 = 0.1

column 282 in row 136 = 0.15

column 285 in row 133 = 1.0.

```
int colind[3], rowind[3];
double value[3];
colind[0] = 282; rowind[0] = 134; value[0] = 0.1;
colind[1] = 282; rowind[1] = 136; value[1] = 0.15;
colind[2] = 285; rowind[2] = 133; value[2] = 1.0;
XSLPadddfs(prob, 3, colind, rowind, value);
```

Further information

The *distribution factor* of a column in a row is the matrix coefficient of the corresponding delta vector in the row. Distribution factors are used in conventional recursion models, and are essentially normalized first-order derivatives. Xpress-SLP can accept distribution factors instead of initial values, provided that the values of the variables involved can all be calculated after optimization using determining rows, or by a callback.

The add functions load additional items into the SLP problem. The corresponding load functions delete any existing items first.

Related topics

XSLPchgdf, XSLPgetdf, XSLPloaddfs

XSLPaddtolsets

Purpose

This subroutine is deprecated and will be removed in a future release.
Add sets of standard tolerance values to an SLP problem.

Topic areas

SLP, SLP-convergence

Synopsis

```
int XPRS_CC XSLPaddtolsets(XSLPprob prob, int ntol, double *tol);
```

Arguments

prob	The current SLP problem.
ntol	The number of tolerance sets to be added.
tol	Double array of (ntol * 9) items containing the 9 tolerance values for each set in order.

Example

The following example creates two tolerance sets: the first has values of 0.005 for all tolerances; the second has values of 0.001 for relative tolerances (numbers 2,4,6,8), values of 0.01 for absolute tolerances (numbers 1,3,5,7) and zero for the closure tolerance (number 0).

```
double tol[18];
for (i=0;i<9;i++) tol[i] = 0.005;
tol[9] = 0;
for (i=10;i<18;i=i+2) tol[i] = 0.01;
for (i=11;i<18;i=i+2) tol[i] = 0.001;
XSLPaddtolsets(prob, 2, tol);
```

Further information

A tolerance set is an array of 9 values containing the following tolerances:

Entry / Bit	Tolerance	XSLP constant	XSLP bit constant
0	Closure tolerance (TC)	XSLP_TOLSET_TC	XSLP_TOLSETBIT_TC
1	Absolute delta tolerance (TA)	XSLP_TOLSET_TA	XSLP_TOLSETBIT_TA
2	Relative delta tolerance (RA)	XSLP_TOLSET_RA	XSLP_TOLSETBIT_RA
3	Absolute coefficient tolerance (TM)	XSLP_TOLSET_TM	XSLP_TOLSETBIT_TM
4	Relative coefficient tolerance (RM)	XSLP_TOLSET_RM	XSLP_TOLSETBIT_RM
5	Absolute impact tolerance (TI)	XSLP_TOLSET_TI	XSLP_TOLSETBIT_TI
6	Relative impact tolerance (RI)	XSLP_TOLSET_RI	XSLP_TOLSETBIT_RI
7	Absolute slack tolerance (TS)	XSLP_TOLSET_TS	XSLP_TOLSETBIT_TS
8	Relative slack tolerance (RS)	XSLP_TOLSET_RS	XSLP_TOLSETBIT_RS

The XSLP_TOLSET constants can be used to access the corresponding entry in the value arrays, while the XSLP_TOLSETBIT constants are used to set or retrieve which tolerance values are used for a given SLP variable.

Once created, a tolerance set can be used to set the tolerances for any SLP variable.

If a tolerance value is zero, then the default tolerance will be used instead. To force the use of a tolerance, use the XSLPchgtolset function and set the Status variable appropriately.

See the section Convergence criteria for a fuller description of tolerances and their uses.

The add functions load additional items into the SLP problem. The corresponding load functions delete any existing items first.

Related topics

XSLPchgtolset, XSLPdeltolsets, XSLPgettolset, XSLPloadtolsets

XSLPadduserfunction, XPRSnlpadduserfunction

Purpose

Add user function definitions to an SLP problem.

Topic area

User Functions

Synopsis

```
int XPRS_CC XSLPadduserfunction(XSLPprob prob, const char * funcname, int
    functype, int nin, int nout, int options, XPRSfunctionptr
    function, void * data, int * p_type);
```

Arguments

prob	The current SLP problem.
funcname	The name of the function as it appears in text formula expressions.
functype	The type of the user function, one of <ol style="list-style-type: none"> 1 (XSLP_USERFUNCTION_MAP) function takes double, returns double. 2 (XSLP_USERFUNCTION_VECMAP) function takes double array, returns double. 3 (XSLP_USERFUNCTION_MULTIMAP) function takes double array, returns double array. 4 (XSLP_USERFUNCTION_MAPDELTA) function takes double, returns double and delta. 5 (XSLP_USERFUNCTION_VECMAPDELTA) function takes double array, returns double and deltas. 6 (XSLP_USERFUNCTION_MULTIMAPDELTA) function takes double array, returns double array and deltas.
nin	Number of arguments the user function takes.
nout	Number of return arguments for the function.
options	options as a bitmap to the user function XSLP_INSTANCEFUNCTION always instantiate the function.
function	Pointer of the user function to call.
data	Context pointer to provide the user function with.
p_type	The token id of the user function added, to be used in the Value array when defining formulas and using with XSLP_FUN.

Further information

The type XPRSfunctionptr is a generic function pointer.

The function declarations expected for the user functions are defined by the functype argument.

The function of type XSLP_USERFUNCTION_MAP expects a function in the form of 'double XPRS_CC F(double Value, void *Context)'.

The function of type XSLP_USERFUNCTION_VECMAP expects a function in the form of 'double XPRS_CC F(double *Value, void *Context)'.

The function of type XSLP_USERFUNCTION_MULTIMAP expects a function in the form of 'int XPRS_CC F(double *Value, double *Out, void *Context)'.

The function of type XSLP_USERFUNCTION_MAPDELTA expects a function in the form of 'int XPRS_CC F(double Value, double Delta, double *Evaluation, double *Partial, void *Context)'.

The function of type XSLP_USERFUNCTION_VECMAPDELTA expects a function in the form of 'int XPRS_CC F(double *Value, double *Deltas, double *Evaluation, double *Partials, void *Context)'.

The function of type XSLP_USERFUNCTION_MULTIMAPDELTA expects a function in the form of 'int XPRS_CC F(double *Value, double *Deltas, double *Out, void *Context)'.

Related topics

User functions, `XSLPdeluserfunction`, `XSLPimportlibfunc`

XSLPaddvars

Purpose

This subroutine is deprecated and will be removed in a future release.
Add SLP variables defined as matrix columns to an SLP problem.

Topic area

Data Input

Synopsis

```
int XPRS_CC XSLPaddvars(XSLPprob prob, int nvars, int *colind, int
    *coltype, int *detrow, int *seqnum, int *tolind, double *initial,
    double *stepbound);
```

Arguments

prob	The current SLP problem.
nvars	The number of SLP variables to be added.
colind	Integer array holding the index of the matrix column corresponding to each SLP variable.
coltype	Bitmap giving information about the SLP variables, compare the variable status keys, in particular: Bit 2 (XSLP_HASIV) Variable has an initial value; May be NULL if not required.
detrow	Integer array holding the index of the determining row for each SLP variable (a negative value means there is no determining row) May be NULL if not required.
seqnum	Integer array holding the index sequence number for cascading for each SLP variable (a zero value means there is no pre-defined order for this variable) May be NULL if not required.
tolind	Integer array holding the index of the tolerance set for each SLP variable (a zero value means the default tolerances are used) May be NULL if not required.
initial	Double array holding the initial value for each SLP variable (use the coltype bit map to indicate if a value is being provided) May be NULL if not required.
stepbound	Double array holding the initial step bound size for each SLP variable (a zero value means that no initial step bound size has been specified). If a value of XPRS_PLUSINFINITY is used for a value in stepbound, the delta will never have step bounds applied, and will almost always be regarded as converged. May be NULL if not required.

Example

The following example loads two SLP variables into the problem. They correspond to columns 23 and 25 of the underlying LP problem. Column 25 has an initial value of 1.42; column 23 has no specific initial value

```
int colind[2], coltype[2];
double initial[2];

colind[0] = 23; coltype[0] = 0; initial[0] = 0.0;
colind[1] = 25; coltype[1] = XSLP_HASIV; initial[1] = 1.42;

XSLPaddvars(prob, 2, colind, coltype, NULL, NULL,
    NULL, initial, NULL);
```

Note that `initial` for the first variable will not actually be used, because the `XSLP_HASIV` flag is not set (`coltype` = 0). Setting the variable type to `XSLP_HASIV` for the second variable indicates that the initial value has been set.

The arrays for determining rows, sequence numbers, tolerance sets and step bounds are not used at all, and so have been passed to the function as `NULL`.

Further information

The add functions load additional items into the SLP problem. The corresponding load functions delete any existing items first.

Related topics

`XSLPchgvar`, `XSLPdelvars`, `XSLPgetvar`, `XSLPloadvars`

XSLPcalcslacks, XPRSnlpcalcslacks

Purpose

Calculate the slack values for the provided solution in the non-linear problem

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPcalcslacks(XSLPprob prob, const double[] solution, double[]  
    slack);
```

Arguments

prob	The current SLP problem.
solution	The solution for which the slacks are requested for.
slack	Vector of length NROWS to return the slack in.

Related topics

XSLPvalidate, XSLPvalidaterow

XSLPcascade, XPRSslpcascade

Purpose

Re-calculate consistent values for SLP variables based on the current values of the remaining variables.

Topic areas

SLP, Solution Process, Cascading

Synopsis

```
int XPRS_CC XSLPcascade(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

The following example changes the solution value for column 91, and then re-calculates the values of those dependent on it.

```
int ColNum;
double Value;

ColNum = 91;
XSLPgetvar(prob, ColNum, NULL, NULL, NULL, NULL,
           NULL, NULL, &Value, NULL, NULL, NULL,
           NULL, NULL, NULL, NULL, NULL);

Value = Value + 1.42;
XSLPchgvar(prob, ColNum, NULL, NULL, NULL, NULL,
           NULL, NULL, &Value, NULL, NULL, NULL,
           NULL);
XSLPcascade(prob);
```

XSLPgetvar and XSLPchgvar are being used to get and change the current value of a single variable.

Provided no other values have been changed since the last execution of XSLPcascade, values will be changed only for variables which depend on column 91.

Further information

See the section on cascading for an extended discussion of the types of cascading which can be performed.

XSLPcascade is called automatically during the SLP iteration process and so it is not normally necessary to perform an explicit cascade calculation.

The variables are re-calculated in accordance with the order generated by XSLPcascadeorder.

Related topics

XSLPcascadeorder, XSLP_CASCADE, XSLP_CASCADENLIMIT, XSLP_CASCADETOL_PA,
XSLP_CASCADETOL_PR

XSLPcascadeorder, XPRSslpcascadeorder

Purpose

Establish a re-calculation sequence for SLP variables with determining rows.

Topic areas

SLP, Data Input

Synopsis

```
int XPRS_CC XSLPcascadeorder(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

Assuming that all variables are SLP variables, the following example sets default values for the variables, creates the re-calculation order and then calls XSLPcascade to calculate consistent values for the dependent variables.

```
int ColNum;
for (ColNum=1; ColNum<=nCol; ColNum++)
    XSLPchgvar(prob, ColNum, NULL, NULL, NULL, NULL,
               NULL, NULL, &DefaultValue[ColNum], NULL, NULL, NULL,
               NULL);
XSLPcascadeorder(prob);
XSLPcascade(prob);
```

Further information

XSLPcascadeorder is called automatically at the start of the SLP iteration process and so it is not normally necessary to perform an explicit cascade ordering.

Related topics

XSLPcascade

XSLPchgascadenlimit, XPRSslpchgascadenlimit

Purpose

Set a variable specific cascade iteration limit

Topic areas

SLP, Data Input

Synopsis

```
int XPRS_CC XSLPchgascadenlimit(XSLPprob prob, int col, int limit);
```

Arguments

prob	The current SLP problem.
col	The index of the column corresponding to the SLP variable for which the cascading limit is to be imposed.
limit	The new cascading iteration limit.

Further information

A value set by this function will overwrite the value of XSLP_CASCADENLIMIT for this variable. To remove any previous value set by this function, use an iteration limit of 0.

Related topics

XSLPcascadeorder, XSLP_CASCADE, XSLP_CASCADENLIMIT, XSLP_CASCADETOL_PA,
XSLP_CASCADETOL_PR

XPRSslpchgcoefstr

Purpose

Add or change a single matrix coefficient using a character string for the formula. For a simpler version of this function see XPRSslpchgformulastr.

Topic areas

SLP, Problem Modification

Synopsis

```
int XPRS_CC XPRSslpchgcoefstr(XPRSprob prob, int row, int col, const double
    *factor, const char *formula);
```

Arguments

prob	The current problem.
row	The index of the matrix row for the coefficient.
col	The index of the matrix column for the coefficient.
factor	Address of a double precision variable holding the constant multiplier for the formula. If factor is NULL, a value of 1.0 will be used.
formula	Character string holding the formula with the tokens separated by spaces.

Example

Assuming that the columns of the matrix are named Col1, Col2, etc, the following example puts the formula `2.5*sin(Col1)` into the coefficient in row 1, column 3.

```
char *formula="sin ( Col1 )";
double factor;

factor = 2.5;
XPRSslpchgcoefstr(prob, 1, 3, &factor, formula);
```

Note that all the tokens in the formula (including mathematical operators and separators) are separated by one or more spaces.

Further information

If the coefficient already exists as a constant or formula, it will be changed into the new coefficient. If it does not exist, it will be added to the problem.

A coefficient is made up of two parts: `factor` and `formula`. `factor` is a constant multiplier which can be provided in the `factor` variable. If Xpress NonLinear can identify a constant factor in the formula, then it will use that as well, to minimize the size of the formula which has to be calculated.

This function can only be used if all the operands in the formula can be correctly identified as constants, existing columns, character variables or functions. Therefore, if a formula refers to a new column, that new item must be added to the Xpress NonLinear problem first.

Related topics

XPRSslpgetformulastr, XSLPaddformulas, XPRSslpchgformulastr, XSLPchgformula,
XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPchgcoef, XPRSslpchgcoef

Purpose

This subroutine is deprecated and will be removed in a future release. Use XPRSslpchgcoefstr instead.

Add or change a single matrix coefficient using a character string for the formula. For a simpler version of this function see XPRSslpchgformulastr.

Topic areas

SLP, Problem Modification

Synopsis

```
int XPRS_CC XSLPchgcoef(XSLPprob prob, int row, int col, const double
    *factor, const char *formula);
```

Arguments

prob	The current SLP problem.
row	The index of the matrix row for the coefficient.
col	The index of the matrix column for the coefficient.
factor	Address of a double precision variable holding the constant multiplier for the formula. If factor is NULL, a value of 1.0 will be used.
formula	Character string holding the formula with the tokens separated by spaces.

Example

Assuming that the columns of the matrix are named Col1, Col2, etc, the following example puts the formula `2.5*sin(Col1)` into the coefficient in row 1, column 3.

```
char *formula="sin ( Col1 )";
double factor;

factor = 2.5;
XSLPchgcoef(prob, 1, 3, &factor, formula);
```

Note that all the tokens in the formula (including mathematical operators and separators) are separated by one or more spaces.

Further information

If the coefficient already exists as a constant or formula, it will be changed into the new coefficient. If it does not exist, it will be added to the problem.

A coefficient is made up of two parts: `factor` and `formula`. `factor` is a constant multiplier which can be provided in the `factor` variable. If Xpress NonLinear can identify a constant factor in the formula, then it will use that as well, to minimize the size of the formula which has to be calculated.

This function can only be used if all the operands in the formula can be correctly identified as constants, existing columns, character variables or functions. Therefore, if a formula refers to a new column, that new item must be added to the Xpress NonLinear problem first.

Related topics

XPRSslpchgcoefstr, XPRSslpgetcoefstr, XPRSslpgetformulastr, XPRSslpchgformulastr, XSLPaddformulas, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPchgcoef, XPRSslpchgcoef

Purpose

Add or change a single matrix coefficient using a parsed or unparsed formula. For a simpler version of this function see XSLPchgformula.

Topic areas

SLP, Problem Modification

Synopsis

```
int XPRS_CC XSLPchgcoef(XSLPprob prob, int row, int col, const double
    *factor, int parsed, const int[] type, const double[] value);
```

Arguments

prob	The current SLP problem.
row	The index of the matrix row for the coefficient.
col	The index of the matrix column for the coefficient.
factor	Address of a double precision variable holding the constant multiplier for the formula. If factor is NULL, a value of 1.0 will be used.
parsed	Integer indicating the whether the token arrays are formatted as internal unparsed (parsed=0) or internal parsed reverse Polish (parsed=1).
type	Array of token types providing the description and formula for each item.
value	Array of values corresponding to the types in type.

Example

Assuming that the columns of the matrix are named Col1, Col2, etc, the following example puts the formula $2.5 * \sin(\text{Col1})$ into the coefficient in row 1, column 3.

```
int n, iSin, type[4];
double value[4];
double factor;

XSLPgetIndex(prob, XSLP_INTERNALFUNCNAMESNOCASE,
    "sin", &iSin);

n = 0;
type[n] = XSLP_IFUN; value[n++] = iSin;
type[n] = XSLP_COL; value[n++] = 0;
type[n++] = XSLP_RB;
type[n++] = XSLP_EOF;

factor = 2.5;
XSLPchgcoef(prob, 1, 3, &factor, 0, type, value);
```

XSLPgetIndex is used to retrieve the index for the internal function sin. The "nocase" version matches the function name regardless of the (upper or lower) case of the name.

Token type XSLP_COL always counts from 0, so Col1 is always 0.

The formula is written in unparsed form (parsed = 0) and so it is provided as tokens in the same order as they would appear if the formula were written in character form.

Further information

If the coefficient already exists as a constant or formula, it will be changed into the new coefficient. If it does not exist, it will be added to the problem.

A coefficient is made up of two parts: factor and Formula. factor is a constant multiplier which can be provided in the factor variable. If Xpress NonLinear can identify a constant factor in the Formula, then it will use that as well, to minimize the size of the formula which has to be calculated.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula,
XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPchgdelattype, XPRSslpchgdeltatype

Purpose

Changes the type of the delta assigned to a nonlinear variable

Topic areas

SLP, Problem Modification

Synopsis

```
int XPRS_CC XSLPchgdelattype(XSLPprob prob, int nvars, const int[] varind,
    const int[] deltatypes, const double[] values);
```

Arguments

prob	The current SLP problem.	
nvars	The number of SLP variables to change the delta type for.	
varind	Indices of the variables to change the deltas for.	
deltatypes	Type of the delta variable:	
	0 (XSLP_DELTA_CONT)	Differentiable variable, default.
	1 (XSLP_DELTA_SEMICON)	Variable where a minimum perturbation size given in values may be required before a significant change in the problem is achieved.
	2 (XSLP_DELTA_INTEGER)	Variable defined over the grid size given in values.
	3 (XSLP_DELTA_EXPLORE)	Variable where a meaningful step size should automatically be detected, with an upper limit given in values.
values	Grid or minimum step sizes for the variables.	

Further information

Changing the delta type of a variables makes the variable nonlinear.

Related topics

XSLP_SEMICONDELTA, XSLP_INTEGERDELTA, XSLP_EXPLOREDELTA

XSLPchgdf

Purpose

This subroutine is deprecated and will be removed in a future release.
Set or change a distribution factor.

Topic areas

SLP, Data Input

Synopsis

```
int XSLP_CC XSLPchgdf(XSLPprob prob, int col, int row, const double *value)
```

Arguments

prob	The current SLP problem.
col	The index of the column whose distribution factor is to be set or changed.
row	The index of the row where the distribution applies.
value	Address of a double precision variable holding the new value of the distribution factor. May be NULL if not required.

Example

The following example retrieves the value of the distribution factor for column 282 in row 134 and changes it to be twice as large.

```
double value;
XSLPgetdf(prob, 282, 134, &value);
value = value * 2;
XSLPchgdf(prob, 282, 134, &value);
```

Further information

The *distribution factor* of a column in a row is the matrix coefficient of the corresponding delta vector in the row. Distribution factors are used in conventional recursion models, and are essentially normalized first-order derivatives. Xpress NonLinear can accept distribution factors instead of initial values, provided that the values of the variables involved can all be calculated after optimization using determining rows, or by a callback.

Related topics

XSLPadddfs, XSLPgetdf, XSLPladdfs

XSLPchgformulastring, XPRSnlpchgformulastring

Purpose

This subroutine is deprecated and will be removed in a future release. Use XPRSnlpchgformulastr instead.

Add or replace a single matrix formula using a character string for the formula.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XSLPchgformulastring(XSLPprob prob, int row, const char
    *formula);
```

Arguments

prob	The current SLP problem.
row	The index of the matrix row for the coefficient.
formula	Character string holding the formula with the tokens separated by spaces.

Example

Assuming that the columns of the matrix are named Col1, Col2, etc, the following example puts the formula `sin (Col1)` into row 1.

```
char *formula="sin ( Col1 )";
XSLPchgformulastring(prob, 1, formula);
```

Note that all the tokens in the formula (including mathematical operators and separators) are separated by one or more spaces.

Further information

If the coefficient already exists as a constant or formula, it will be changed into the new coefficient. If it does not exist, it will be added to the problem.

This function can only be used if all the operands in the formula can be correctly identified as constants, existing columns, character variables or functions. Therefore, if a formula refers to a new column, that new item must be added to the Xpress NonLinear problem first.

Related topics

XPRSnlpchgformulastr, XPRSnlpgetformulastr, XSLPaddformulas, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XPRSnlpchgformulastr

Purpose

Add or replace a single matrix formula using a character string for the formula.

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XPRSnlpchgformulastr(XPRSprob prob, int row, const char
    *formula);
```

Arguments

prob	The current problem.
row	The index of the matrix row for the coefficient.
formula	Character string holding the formula with the tokens separated by spaces.

Example

Assuming that the columns of the matrix are named Col1, Col2, etc, the following example puts the formula `sin (Col1)` into row 1.

```
char *formula="sin ( Col1 )";
XPRSnlpchgformulastr(prob, 1, formula);
```

Note that all the tokens in the formula (including mathematical operators and separators) are separated by one or more spaces.

Further information

If the coefficient already exists as a constant or formula, it will be changed into the new coefficient. If it does not exist, it will be added to the problem.

This function can only be used if all the operands in the formula can be correctly identified as constants, existing columns, character variables or functions. Therefore, if a formula refers to a new column, that new item must be added to the Xpress NonLinear problem first.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPchgformula, XPRSnlpchgformula

Purpose

Add or replace a single matrix formula using a parsed or unparsed formula

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XSLPchgformula(XSLPprob prob, int row, int parsed, const int[]
    type, const double []value);
```

Arguments

prob	The current SLP problem.
row	The index of the matrix row for the coefficient.
parsed	Integer indicating the whether the token arrays are formatted as internal unparsed (parsed=0) or internal parsed reverse Polish (parsed=1).
type	Array of token types providing the description and formula for each item.
value	Array of values corresponding to the types in type.

Example

Assuming that the columns of the matrix are named Col1, Col2, etc, the following example puts the formula `sin(Col1)` into the coefficient in row 1.

```
int n, iSin, type[4];
double value[4];

XSLPgetindex(prob, XSLP_INTERNALFUNCNAMESNOCASE,
    "sin", &iSin);

n = 0;
type[n] = XSLP_IFUN; value[n++] = iSin;
type[n] = XSLP_COL; value[n++] = 0;
type[n++] = XSLP_RB;
type[n++] = XSLP_EOF;

Factor = 2.5;
XSLPchgformula(prob, 1, 0, type, value);
```

XSLPgetindex is used to retrieve the index for the internal function `sin`. The "nocase" version matches the function name regardless of the (upper or lower) case of the name.

Token type XSLP_COL always counts from 0, so Col1 is always 0.

The formula is written in unparsed form (parsed = 0) and so it is provided as tokens in the same order as they would appear if the formula were written in character form.

Further information

If the row already has a nonlinear expression in it, it will be changed into the new formula. If it does not exist, it will be added to the problem.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPchgrowstatus, XPRSslpchgrowstatus

Purpose

Change the status setting of a constraint

Topic areas

SLP, Data Input, Bit-vector

Synopsis

```
int XPRS_CC XSLPchgrowstatus(XSLPprob prob, int row, int *status);
```

Arguments

prob	The current SLP problem.
row	The index of the matrix row to be changed.
status	Address of an integer holding a bitmap with the new status settings. If the status is to be changed, always get the current status first (use XSLPgetrowstatus) and then change settings as required. The only settings likely to be changed are: Bit 11 Set if row must not have a penalty error vector. This is the equivalent of an enforced constraint (SLPDATA type EC).

Example

The following example changes the status of row 9 to be an enforced constraint.

```
int row, status;
row = 9;
XSLPgetrowstatus(prob, row, &status);
status = status | (1<<11);
XSLPchgrowstatus(prob, row, &status);
```

Further information

If status is NULL the current status will remain unchanged.

Related topics

XSLPgetrowstatus

XSLPchgrowwt, XPRSSlpchgrowwt

Purpose

Set or change the initial penalty error weight for a row

Topic areas

SLP, Data Input

Synopsis

```
int XSLP_CC XSLPchgrowwt(XSLPprob prob, int row, const double *weight);
```

Arguments

prob	The current SLP problem.
row	The index of the row whose weight is to be set or changed.
weight	Address of a double precision variable holding the new value of the weight. May be NULL if not required.

Example

The following example sets the initial weight of row number 2 to a fixed value of 3.6 and the initial weight of row 4 to a value twice the calculated default value.

```
double weight;
weight = -3.6;
XSLPchgrowwt(prob, 2, &weight);
weight = 2.0;
XSLPchgrowwt(prob, 4, &weight);
```

Further information

A positive value is interpreted as a multiplier of the default row weight calculated by Xpress-SLP.

A negative value is interpreted as a fixed value: the absolute value is used directly as the row weight.

The initial row weight is used only when the augmented structure is created.

Related topics

XSLPgetrowwt

XSLPchgtolset

Purpose

This subroutine is deprecated and will be removed in a future release.
Add or change a set of convergence tolerances used for SLP variables.

Topic areas

SLP, SLP-convergence

Synopsis

```
int XPRS_CC XSLPchgtolset(XSLPprob prob, int ntols, int *status, double
    *tols);
```

Arguments

prob	The current SLP problem.
tolset	Tolerance set for which values are to be changed. A zero value for <code>tolset</code> will create a new set.
status	Address of an integer holding a bitmap describing which tolerances are active in this set. See below for the settings.
tols	Array of 9 double precision values holding the values for the corresponding tolerances.

Example

The following example creates a new tolerance set with the default values for all tolerances except the relative delta tolerance, which is set to 0.005. It then changes the value of the absolute delta and absolute impact tolerances in tolerance set 6 to 0.015

```
int status;
double tols[9];

tolset[2] = 0.005;
status = 1<<2;
XSLPchgtolset(prob, 0, &status, tols);
tolset[1] = tolset[5] = 0.015;
status = 1<<1 | 1<<5;
XSLPchgtolset(prob, 6, &status, tols);
```

Further information

The bits in `status` are set to indicate that the corresponding tolerance is to be changed in the tolerance set. The meaning of the bits is as follows:

Entry / Bit	Tolerance	XSLP constant	XSLP bit constant
0	Closure tolerance (TC)	XSLP_TOLSET_TC	XSLP_TOLSETBIT_TC
1	Absolute delta tolerance (TA)	XSLP_TOLSET_TA	XSLP_TOLSETBIT_TA
2	Relative delta tolerance (RA)	XSLP_TOLSET_RA	XSLP_TOLSETBIT_RA
3	Absolute coefficient tolerance (TM)	XSLP_TOLSET_TM	XSLP_TOLSETBIT_TM
4	Relative coefficient tolerance (RM)	XSLP_TOLSET_RM	XSLP_TOLSETBIT_RM
5	Absolute impact tolerance (TI)	XSLP_TOLSET_TI	XSLP_TOLSETBIT_TI
6	Relative impact tolerance (RI)	XSLP_TOLSET_RI	XSLP_TOLSETBIT_RI
7	Absolute slack tolerance (TS)	XSLP_TOLSET_TS	XSLP_TOLSETBIT_TS
8	Relative slack tolerance (RS)	XSLP_TOLSET_RS	XSLP_TOLSETBIT_RS

The XSLP_TOLSET constants can be used to access the corresponding entry in the value arrays, while the XSLP_TOLSETBIT constants are used to set or retrieve which tolerance values are used for a given SLP variable. The members of the `tols` array corresponding to nonzero bit settings in `status` will be used to change the tolerance set. So, for example, if bit 3 is set in `status`, then `tols[3]` will replace the current value of the absolute coefficient tolerance. If a bit is not set in `status`, the value of the corresponding element of `tols` is unimportant.

Related topics

`XSLPaddtolsets`, `XSLPdeltolsets`, `XSLPgettolset`, `XSLPloadtolsets`

XSLPchgvar

Purpose

This subroutine is deprecated and will be removed in a future release.

Define a column as an SLP variable or change the characteristics and values of an existing SLP variable.

Topic area

Data Input

Synopsis

```
int XPRS_CC XSLPchgvar(XSLPprob prob, int col, int *detrow, double
    *initstepbound, double *stepbound, double *penalty, double *damp,
    double *initial, double *value, int *tolset, int *history, int
    *converged, int *vartype);
```

Arguments

prob	The current SLP problem.
col	The index of the matrix column.
detrow	Address of an integer holding the index of the determining row. Use -1 if there is no determining row. May be NULL if not required.
initstepbound	Address of a double precision variable holding the initial step bound size. May be NULL if not required.
stepbound	Address of a double precision variable holding the current step bound size. Use zero to disable the step bounds. May be NULL if not required.
penalty	Address of a double precision variable holding the weighting of the penalty cost for exceeding the step bounds. May be NULL if not required.
damp	Address of a double precision variable holding the damping factor for the variable. May be NULL if not required.
initial	Address of a double precision variable holding the initial value for the variable. May be NULL if not required.
value	Address of a double precision variable holding the current value for the variable. May be NULL if not required.
tolset	Address of an integer holding the index of the tolerance set for this variable. Use zero if there is no specific tolerance set. May be NULL if not required.
history	Address of an integer holding the history value for this variable. May be NULL if not required.
converged	Address of an integer holding the convergence status for this variable. May be NULL if not required.
vartype	Address of an integer holding a bitmap defining the existence of certain properties for this variable: Bit 1: Variable has a delta vector Bit 2: Variable has an initial value Bit 14: Variable is the reserved "=" column May be NULL if not required.

Example

The following example sets an initial value of 1.42 and tolerance set 2 for column 25 in the matrix.

```
double InitialValue;
int vartype, tolset;

InitialValue = 1.42;
tolset = 2;
vartype = 1<<1 | 1<<2;
```

```
XSLPchgvar(prob, 25, NULL, NULL, NULL, NULL,  
           NULL, &InitialValue, NULL, &tolset,  
           NULL, NULL, &vartype);
```

Note that bits 1 and 2 of `vartype` are set, indicating that the variable has a delta vector and an initial value. For columns already defined as SLP variables, use `XSLPgetvar` to obtain the current value of `vartype` because other bits may already have been set by the system.

Further information

If any of the arguments is `NULL` then the corresponding information for the variable will be left unaltered. If the information is new (i.e. the column was not previously defined as an SLP variable) then the default values will be used.

Changing `history` or `converged` is only effective during SLP iterations, changing `value` only during iterations and before calling functions that specifically use this like `XSLPcascade`.

Changing `initial` and `initstepbound` is only effective before `XSLPconstruct`.

If a value of `XPRS_PLUSINFINITY` is used in the value for `stepbound` or `initstepbound`, the delta will never have step bounds applied, and will almost always be regarded as converged.

Related topics

`XSLPaddvars`, `XSLPdelvars`, `XSLPgetvar`, `XSLPloadvars`

XSLPconstruct, XPRSslpconstruct

Purpose

Create the full augmented SLP matrix and data structures, ready for optimization

Topic areas

SLP, Solution Process

Synopsis

```
int XPRS_CC XSLPconstruct(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

The following example constructs the augmented matrix and then outputs the result in MPS format to a file called augment.mat

```
/* creation and/or loading of data */
/* precedes this segment of code */
...
XSLPconstruct(prob);
XSLPwriteprob(prob, "augment", "1");
```

The "1" flag causes output of the current linear problem (which is now the augmented structure and the current linearization) rather than the original nonlinear problem.

Further information

XSLPconstruct adds new rows and columns to the SLP matrix and calculates initial values for the non-linear coefficients. Which rows and columns are added will depend on the setting of XSLP_AUGMENTATION. Names for the new rows and columns are generated automatically, based on the existing names and the string control variables XSLP_XXXFORMAT.

Once XSLPconstruct has been called, no new rows, columns or non-linear coefficients can be added to the problem. Any rows or columns which will be required must be added first. Non-linear coefficients must not be changed; constant matrix elements can generally be changed after XSLPconstruct, but not after XSLPpresolve if used.

XSLPconstruct is called automatically by the SLP optimization procedure, and so only needs to be called explicitly if changes need to be made between the augmentation and the optimization.

Related topics

XSLPpresolve

XSLPcopycallbacks

Purpose

Copy the user-defined callbacks from one SLP problem to another

Topic area

Callback

Synopsis

```
int XPRS_CC XSLPcopycallbacks(XSLPprob dest, XSLPprob src);
```

Arguments

dest	The SLP problem to receive the callbacks.
src	The SLP problem from which the callbacks are to be copied.

Example

The following example creates a new problem and copies only the Xpress NonLinear callbacks from the existing problem (not the Optimizer library ones).

```
XSLPprob nProb;
XPRSprob xProb;
int Control;

XSLPcreateprob(&nProb, &xProb);

Control = 1<<2;
XSLPsetintcontrol(Prob, XSLP_CONTROL, Control);
XSLPcopycallbacks(nProb, Prob);
```

Note that XSLP_CONTROL is set in the *old* problem, not the new one.

Further information

Normally XSLPcopycallbacks copies both the Xpress NonLinear callbacks and the Optimizer Library callbacks for the underlying problem. If only the Xpress NonLinear callbacks are required, set the integer control variable XSLP_CONTROL appropriately.

Related topics

XSLP_CONTROL

XSLPcopycontrols

Purpose

Copy the values of the control variables from one SLP problem to another

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPcopycontrols(XSLPprob dest, XSLPprob src);
```

Arguments

dest	The SLP problem to receive the controls.
src	The SLP problem from which the controls are to be copied.

Example

The following example creates a new problem and copies only the Xpress NonLinear controls from the existing problem (not the Optimizer library ones).

```
XSLPprob nProb;
XPRSprob xProb;
int Control;

XSLPcreateprob(&nProb, &xProb);

Control = 1<<1;
XSLPsetintcontrol(Prob, XSLP_CONTROL, Control);
XSLPcopycontrols(nProb, Prob);
```

Note that XSLP_CONTROL is set in the *old* problem, not the new one.

Further information

Normally XSLPcopycontrols copies both the Xpress NonLinear controls and the Optimizer Library controls for the underlying problem. If only the Xpress NonLinear controls are required, set the integer control variable XSLP_CONTROL appropriately.

Since the objective sense is a control in Xpress nonlinear, XSLPcopycontrols will also copy the objective sense from the source to the destination.

Related topics

XSLP_CONTROL

XSLPcopyprob

Purpose

Copy an existing SLP problem to another

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XSLPcopyprob(XSLPprob dest, XSLPprob src, char *probname);
```

Arguments

dest	The SLP problem to receive the copy.
src	The SLP problem from which to copy.
probname	The name to be given to the problem.

Example

The following example creates a new Xpress NonLinear problem and then copies an existing problem to it. The new problem is named "ANewProblem".

```
XSLPprob nProb;
XPRSprob xProb;

XSLPcreateprob(&nProb, &xProb);
XSLPcopyprob(nProb, Prob, "ANewProblem");
```

Further information

Normally XSLPcopyprob copies both the Xpress NonLinear problem and the underlying Optimizer Library problem. If only the Xpress NonLinear problem is required, set the integer control variable XSLP_CONTROL appropriately.

This function does not copy callbacks. These must be copied separately using XSLPcopycallbacks if required.

Related topics

XSLP_CONTROL

XSLPcreateprob

Purpose

Create a new SLP problem

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XSLPcreateprob(XSLPprob *p_prob, XPRSprob *p_xprob);
```

Arguments

p_prob	The address of the SLP problem variable.
p_xprob	The address of the underlying Optimizer Library problem variable.

Example

The following example creates an optimizer problem, and then a new Xpress NonLinear problem.

```
XSLPprob nProb;  
XPRSprob xprob;  
  
XPRScreateprob(&xprob);  
XSLPcreateprob(&nProb, &xprob);
```

Further information

An Xpress NonLinear problem includes an underlying optimizer problem which is used to solve the successive linear approximations. The user is responsible for creating and destroying the underlying linear problem, and can also access it using the normal optimizer library functions. When an SLP problem is to be created, the underlying problem is created first, and the SLP problem is then created, knowing the address of the underlying problem.

Related topics

XSLPdestroyprob

XSLPdelcoefs, XPRSslpdelcoefs

Purpose

Delete coefficients from the current problem. For a simpler version of this function see XSLPdelformulas.

Topic areas

SLP, Problem Modification

Synopsis

```
int XPRS_CC XSLPdelcoefs(XSLPprob prob, int ncoefs, const int[] rowind,
    const int[] colind);
```

Arguments

prob	The current SLP problem.
ncoefs	Number of SLP coefficients to delete.
rowind	Row indices of the SLP coefficients to delete.
colind	Column indices of the SLP coefficients to delete.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPdelformulas, XPRSnlpdelformulas

Purpose

Delete nonlinear formulas from the current problem

Topic area

Problem Modification

Synopsis

```
int XPRS_CC XSLPdelformulas(XSLPprob prob, int nformulas, const int[]  
    rowind);
```

Arguments

prob	The current SLP problem.
nformulas	Number of SLP nonlinear formulas to delete.
rowind	Row indices of the SLP nonlinear formulas to delete.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula,
XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPdeltolsets

Purpose

This subroutine is deprecated and will be removed in a future release.

Delete tolerance sets from the current problem

Topic areas

SLP, SLP-convergence

Synopsis

```
int XPRS_CC XSLPdeltolsets(XSLPprob prob, int ntolsets, int *tolind);
```

Arguments

prob	The current SLP problem.
ntolsets	Number of tolerance sets to delete.
tolind	Indices of tolerance sets to delete.

Related topics

XSLPaddtolsets, XSLPchgtolset, XSLPgettolset, XSLPloadtolsets

XSLPdeluserfunction, XPRSnlpdeluserfunction

Purpose

Delete a user function from the current problem

Topic area

User Functions

Synopsis

```
int XPRS_CC XSLPdeluserfunction(XSLPprob prob, int type);
```

Arguments

prob	The current SLP problem.
type	The identifier of the user function as returned by XSLPadduserfunction.

Related topics

XSLPadduserfunction, XSLPimportlibfunc

XSLPdelvars

Purpose

This subroutine is deprecated and will be removed in a future release.

Convert SLP variables to normal columns. Variables must not appear in SLP structures

Topic area

Data Input

Synopsis

```
int XPRS_CC XSLPdelvars(XSLPprob prob, int nvars, int *colind);
```

Arguments

prob	The current SLP problem.
nvars	Number SLP variables to be converted to linear columns.
colind	Column indices of the SLP vars to be converted to linear ones.

Further information

The SLP variables to be converted to linear, non SLP columns must not be in use by any other SLP structure (coefficients, initial value formulae, delayed columns). Use the appropriate deletion or change functions to remove them first.

Related topics

XSLPaddvars, XSLPchgvar, XSLPgetvar, XSLPloadvars

XSLPdestroyprob

Purpose

Delete an SLP problem and release all the associated memory

Topic area

Problem Creation

Synopsis

```
int XPRS_CC XSLPdestroyprob(XSLPprob prob);
```

Argument

prob The SLP problem.

Example

The following example creates an SLP problem and then destroys it together with the underlying optimizer problem.

```
XSLPprob nProb;  
XPRSprob xProb;  
  
XPRScreateprob(&xProb);  
XSLPcreateprob(&nProb, &xProb);  
...  
XSLPdestroyprob(nProb);  
XPRSdestroyprob(xProb);
```

Further information

When you have finished with the SLP problem, it should be "destroyed" so that the memory used by the problem can be released. Note that this does not destroy the underlying optimizer problem, so a call to `XPRSdestroyprob` should follow `XSLPdestroyprob` as and when you have finished with the underlying optimizer problem.

Related topics

`XSLPcreateprob`

XSLPevaluatecoef, XPRSslpevaluatecoef

Purpose

Evaluate a coefficient using the current values of the variables

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPevaluatecoef(XSLPprob prob, int row, int col, double
    *p_value);
```

Arguments

prob	The current SLP problem.
row	Integer index of the row.
col	Integer index of the column.
p_value	Address of a double precision value to receive the result of the calculation.

Example

The following example sets the value of column 5 to 1.42 and then calculates the coefficient in row 2, column 3. If the coefficient depends on column 5, then a value of 1.42 will be used in the calculation.

```
double value, dValue;

value = 1.42;
XSLPchgvar(prob, 5, NULL, NULL, NULL, NULL,
    NULL, NULL, &value, NULL, NULL, NULL,
    NULL);
XSLPevaluatecoef(prob, 2, 3, &dValue);
```

Further information

The values of the variables are obtained from the solution, or from the p_value setting of an SLP variable (see XSLPchgvar and XSLPgetvar).

Related topics

XSLPchgvar, XSLPevaluateformula XSLPgetvar

XSLPevaluateformula, XPRSnlpevaluateformula

Purpose

Evaluate a formula using the current values of the variables

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPevaluateformula(XSLPprob prob, int parsed, const int[]
                               type, const double[] values, double *p_value);
```

Arguments

prob	The current SLP problem.
parsed	integer indicating whether the formula of the item is in internal unparsed format (parsed=0) or parsed (reverse Polish) format (parsed=1).
type	Integer array of token types for the formula.
values	Double array of values corresponding to type.
p_value	Address of a double precision value to receive the result of the calculation.

Example

The following example calculates the value of column 3 divided by column 6.

```
int n, type[10];
double value, values[10];

n = 0;
type[n] = XSLP_COL; values[n++] = 3;
type[n] = XSLP_COL; values[n++] = 6;
type[n] = XSLP_OP; values[n++] = XSLP_DIVIDE;
type[n++] = XSLP_EOF;

XSLPevaluateformula(prob, 1, type, values, &value);
```

Further information

The formula in type and values must be terminated by an XSLP_EOF token.

The formula cannot include "complicated" functions, such as user functions which return more than one value

Related topics

XSLPevaluatecoef

XSLPfixpenalties, XPRSslpfixpenalties

Purpose

Fixe the values of the error vectors

Topic area

Solution Process

Synopsis

```
int XPRS_CC XSLPfixpenalties(XSLPprob prob, int *p_status);
```

Arguments

prob	The current SLP problem.
p_status	Return status after fixing the penalty variables: 0 is successful, nonzero otherwise.

Further information

The function fixes the values of all error vectors on their current values. It also removes their objective cost contribution.

The function is intended to support post optimization analysis, by removing any possible direct effect of the error vectors from the dual and reduced cost values.

The XSLPfixpenalties will automatically reoptimize the linearization. However, as the XSLP convergence and infeasibility checks (regarding the original non-linear problem) will not be carried out, this function will not update the SLP solution itself. The updated values will be accessible using XPRSgetlpsol instead.

XSLPfree

Purpose

Free any memory allocated by Xpress NonLinear and close any open Xpress NonLinear files

Topic area

Licensing

Synopsis

```
int XPRS_CC XSLPfree(void);
```

Example

The following code frees the Xpress NonLinear memory and then frees the optimizer memory:

```
XSLPfree();  
XPRSfree();
```

Further information

A call to `XSLPfree` only frees the items specific to Xpress NonLinear. `XPRSfree` must be called after `XSLPfree` to free the optimizer structures.

Related topics

`XSLPinit`

XSLPgetbanner

Purpose

This subroutine is deprecated and will be removed in a future release.

This function is deprecated and may be removed in future releases. This function has the same effect as XPRSgetbanner

Topic area

Logging

Synopsis

```
int XPRS_CC XSLPgetbanner(char *banner);
```

Argument

banner	Character buffer to hold the banner. This will be at most 512 characters including the null terminator.
--------	---

XSLPgetccoef, XPRSslpgetccoef

Purpose

This subroutine is deprecated and will be removed in a future release. Use XPRSslpgetcoefstr instead.

Retrieve a single matrix coefficient as a formula in a character string. For a simpler version of this function see XPRSnlpgetformulastr.

Topic areas

SLP, Problem Information

Synopsis

```
int XPRS_CC XSLPgetccoef(XSLPprob prob, int row, int col, double *p_factor,
    char *formula, int maxbytes);
```

Arguments

prob	The current SLP problem.
row	Integer holding the row index for the coefficient.
col	Integer holding the column index for the coefficient.
p_factor	Address of a double precision variable to receive the value of the constant factor multiplying the formula in the coefficient.
formula	Character buffer in which the formula will be placed in the same format as used for input from a file. The formula will be null terminated.
maxbytes	Maximum length of returned formula.

Return value

0	Normal return.
32	formula is too long for the buffer and has been truncated.
other	Error.

Example

The following example displays the formula for the coefficient of column 3 in row 2:

```
char Buffer[60];
double factor;
int Code;

Code = XSLPgetccoef(prob, 2, 3, &factor, Buffer, 60);
switch (Code) {
case 0:  printf("\nFormula is %s", Buffer);
        printf("\nFactor = %lg", factor);
        break;
case 32: printf("\nFormula is too long for the buffer");
        break;
default: printf("\nError accessing coefficient");
        break;
}
```

Further information

If the requested coefficient is constant, then p_factor will be set to 1.0 and the value will be formatted in formula.

If the length of the formula would exceed maxbytes-1, the formula is truncated to the last token that will fit, and the (partial) formula is terminated with a null character.

Related topics

XPRSslpgetcoefstr, XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XPRSSlpgetcoefstr

Purpose

Retrieve a single matrix coefficient as a formula in a character string. For a simpler version of this function see `XPRSSnlpgetformulastr`.

Topic areas

SLP, Problem Information

Synopsis

```
int XPRS_CC XPRSSlpgetcoefstr(XPRSprob prob, int row, int col, double
    *p_factor, char *formula, int maxbytes, int *p_nbytes);
```

Arguments

<code>prob</code>	The current SLP problem.
<code>row</code>	Integer holding the row index for the coefficient.
<code>col</code>	Integer holding the column index for the coefficient.
<code>p_factor</code>	Address of a double precision variable to receive the value of the constant factor multiplying the formula in the coefficient.
<code>formula</code>	Character buffer in which the formula will be placed in the same format as used for input from a file. The formula will be null terminated. This argument may be <code>NULL</code> if <code>maxbytes</code> is zero.
<code>maxbytes</code>	Length of the formula buffer.
<code>p_nbytes</code>	Will be set to the length of the formula, not including the null terminator.

Example

The following example displays the formula for the coefficient of column 3 in row 2:

```
char *buffer;
double factor;
int len;

XPRSSlpgetcoefstr(prob, 2, 3, &factor, NULL, 0, &len);
buffer = malloc(len + 1);
XPRSSlpgetcoefstr(prob, 2, 3, &factor, buffer, len + 1, NULL);
printf("\nFormula is %s", buffer);
```

Further information

If the requested coefficient is constant, then `p_factor` will be set to 1.0 and the value will be formatted in `formula`.

If the length of the formula would exceed `maxbytes-1`, the formula is truncated to the last token that will fit, and the (partial) formula is terminated with a null character.

Related topics

`XPRSSnlpgetformulastr`, `XSLPaddformulas`, `XPRSSnlpchgformulastr`, `XSLPchgformula`, `XSLPloadformulas`, `XSLPgetformularows`, `XSLPgetformula`, `XSLPdelformulas`

XSLPgetcoeffformula, XPRSslpgetcoeffformula

Purpose

Retrieve a single matrix coefficient as a formula split into tokens. For a simpler version of this function see XSLPgetformula.

Topic areas

SLP, Problem Information

Synopsis

```
int XPRS_CC XSLPgetcoeffformula(XSLPprob prob, int row, int col, double
    *p_factor, int parsed, int maxtypes, int *p_ntypes, int[] type,
    double[] value);
```

Arguments

prob	The current SLP problem.
row	Integer holding the row index for the coefficient.
col	Integer holding the column index for the coefficient.
p_factor	Address of a double precision variable to receive the value of the constant factor multiplying the formula in the coefficient.
parsed	Integer indicating whether the formula of the item is to be returned in internal unparsed format (parsed=0) or parsed (reverse Polish) format (parsed=1).
maxtypes	Maximum number of tokens to return, i.e. length of the type and value arrays.
p_ntypes	Number of tokens returned in type and value.
type	Integer array to hold the token types for the formula.
value	Double array of values corresponding to type.

Example

The following example displays the formula for the coefficient of column 3 in row 2 in unparsed form:

```
int n, type[10];
double value[10];
double factor;
int ntypes;

XSLPgetcoeffformula(prob, 2, 3, &factor, 0, 10, &ntypes, type, value);

for (n=0; type[n] != XSLP_EOF; n++)
    printf("\nType=%-3d  value=%lg", type[n], value[n]);
```

Further information

1. The type and value arrays are terminated by an XSLP_EOF token.
2. If type and value are both null, the number of tokens available will be returned in p_ntypes.
3. If type or value are non-null and the number of tokens available exceeds maxtypes, an error code will be returned: this behaviour may change in a future release. To determine the number of available tokens, pass null for these arrays.
4. If the requested coefficient is constant, then p_factor will be set to 1.0 and the value will be returned with token type XSLP_CON.

Related topics

XPRSslpgetformulastr, XSLPaddformulas, XPRSslpchgformulastr, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPgetcoefs, XPRSslpgetcoefs

Purpose

Retrieve the list of positions of the nonlinear coefficients in the problem. For a simpler version of this function see XSLPgetformularows.

Topic areas

SLP, Problem Information

Synopsis

```
int XPRS_CC XSLPgetcoefs(XSLPprob prob, int *p_ncoefs, int[] rowind, int[]  
    colind);
```

Arguments

prob	The current SLP problem.
p_ncoefs	Integer used to return the total number of nonlinear coefficients in the problem.
rowind	Integer array used for returning the row positions of the coefficients. May be NULL if not required.
colind	Integer array used for returning the column positions of the coefficients. May be NULL if not required.

Related topics

XPRSslpgetcoefstr, XSLPgetcoeffformula

XSLPgetcolinfo, XPRSslpgetcolinfo

Purpose

Get current column information.

Topic areas

SLP, Solution

Synopsis

```
int XSLP_CC XSLPgetcolinfo(XSLPprob prob, int type, int col, XSLPalltype
    *p_info);
```

Arguments

prob	The current SLP problem
type	Type of information (see below)
col	Index of the column whose information is to be handled
p_info	Address of information to be set or retrieved

Example

The following example prints the current value of variable 3:

```
XSLPalltype alltype;
XSLPgetcolinfo(prob, XSLP_COLINFO_VALUE, 3, &alltype);
printf("value[3]=%f\n", alltype.value.real);
```

Further information

1. If the data is not available, the type of the returned p_info is set to XSLPtype_undefined. Please refer to the header file `xprs.h` for the definition of XPRSalltype which is identical to XSLPalltype. The following constants are provided for column information handling:

XSLP_COLINFO_VALUE	Get the current value of the column
XSLP_COLINFO_RDJ	Get the current reduced cost of the column
XSLP_COLINFO_DELTAINDEX	Get the delta variable index associated to the column
XSLP_COLINFO_DELTA	Get the delta value (change since previous value) of the column
XSLP_COLINFO_DELTADJ	Get the delta variables reduced cost
XSLP_COLINFO_UPDATEROW	Get the index of the update (or step bound) row associated to the column
XSLP_COLINFO_SB	Get the step bound on the variable
XSLP_COLINFO_SBDUAL	Get the dual multiplier of the step bound row for the variable
2. This function can be used to get the current SLP iteration solution in the case where this is not available using XSLPgetslpsol due to solution filtering via XSLP_FILTER.

XSLPgetdblattrib

Purpose

Retrieve the value of a double precision problem attribute

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPgetdblattrib(XSLPprob prob, int attrib, double *p_value);
```

Arguments

prob	The current SLP problem.
attrib	attribute (SLP or optimizer) whose value is to be returned.
p_value	Address of a double precision variable to receive the value.

Example

The following example retrieves the value of the Xpress NonLinear attribute XSLP_CURRENTDELTACOST and of the optimizer attribute XPRS_LPOBJVAL:

```
double DeltaCost, ObjVal;
XSLPgetdblattrib(prob, XSLP_CURRENTDELTACOST, &DeltaCost);
XSLPgetdblattrib(prob, XPRS_LPOBJVAL, &ObjVal);
```

Further information

Both SLP and optimizer attributes can be retrieved using this function. If an optimizer attribute is requested, the return value will be the same as that from [XPRSgetdblattrib](#), which can similarly be used to obtain both Optimizer and SLP attributes.

Related topics

XSLPgetintattrib, XSLPgetstrattrib

XSLPgetdblcontrol

Purpose

Retrieve the value of a double precision problem control

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPgetdblcontrol(XSLPprob prob, int control, double *p_value);
```

Arguments

prob	The current SLP problem.
control	control (SLP or optimizer) whose value is to be returned.
p_value	Address of a double precision variable to receive the value.

Example

The following example retrieves the value of the Xpress NonLinear control XSLP_CTOL and of the optimizer control XPRS_FEASTOL:

```
double CTol, FeasTol;
XSLPgetdblcontrol(prob, XSLP_CTOL, &CTol);
XSLPgetdblcontrol(prob, XPRS_FEASTOL, &FeasTol);
```

Further information

Both SLP and optimizer controls can be retrieved using this function. If an optimizer control is requested, the return value will be the same as that from [XPRSgetdblcontrol](#), which can similarly be used to obtain both Optimizer and SLP controls.

Related topics

XSLPgetintcontrol, XSLPgetstrcontrol, XSLPsetdblcontrol

XSLPgetdf

Purpose

This subroutine is deprecated and will be removed in a future release.
Get a distribution factor.

Topic areas

SLP, Problem Information

Synopsis

```
int XSLP_CC XSLPgetdf(XSLPprob prob, int col, int row, double *p_value)
```

Arguments

prob	The current SLP problem.
col	The index of the column whose distribution factor is to be retrieved.
row	The index of the row from which the distribution factor is to be taken.
p_value	Address of a double precision variable to receive the value of the distribution factor. May be NULL if not required.

Example

The following example retrieves the value of the distribution factor for column 282 in row 134 and changes it to be twice as large.

```
double value;
XSLPgetdf(prob, 282, 134, &value);
value = value * 2;
XSLPchgdf(prob, 282, 134, &value);
```

Further information

The *distribution factor* of a column in a row is the matrix coefficient of the corresponding delta vector in the row. Distribution factors are used in conventional recursion models, and are essentially normalized first-order derivatives. Xpress-SLP can accept distribution factors instead of initial values, provided that the values of the variables involved can all be calculated after optimization using determining rows, or by a callback.

Related topics

XSLPadddfs, XSLPchgdf, XSLPloddfs

XSLPgetformula, XPRSnlpgetformula

Purpose

Retrieve a single matrix formula as a formula split into tokens.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XSLPgetformula(XSLPprob prob, int row, int parsed, int
    maxtypes, int *p_ntypes, int[] type, double[] value);
```

Arguments

prob	The current SLP problem.
row	Integer holding the row index for the formula.
parsed	Integer indicating whether the formula of the row is to be returned in internal unparsed format (parsed=0) or parsed (reverse Polish) format (parsed=1).
maxtypes	Maximum number of tokens to return, i.e., the length of the type and value arrays.
p_ntypes	Will be set to the length of the formula, including the XSLP_EOF token.
type	Integer array to hold the token types for the formula. May be NULL if maxtypes is zero.
value	Double array of values corresponding to type. May be NULL if maxtypes is zero.

Example

The following example displays the nonlinear formula in row 2, column 3 in unparsed form:

```
int ntypes, n;
int *type;
double *value;

XSLPgetformula(prob, 2, 0, 0, &ntypes, NULL, NULL);
type = malloc(ntypes * sizeof(int));
value = malloc(ntypes * sizeof(double));
XSLPgetformula(prob, 2, 0, 10, &ntypes, type, value);

for (n = 0; type[n] != XSLP_EOF; n++)
    printf("\nType=%-3d  value=%lg", type[n], value[n]);
```

Further information

1. The type and value arrays are terminated by an XSLP_EOF token.
2. If type and value are both null, the number of tokens available will be returned in p_ntypes.
3. If type or value are non-null and the number of tokens available exceeds maxtypes, an error code will be returned: this behaviour may change in a future release. To determine the number of available tokens, pass null for these arrays.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula,
XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XPRSnlpgetformulastr

Purpose

Retrieve a single matrix formula in a character string.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XPRSnlpgetformulastr(XPRSprob prob, int row, char *formula, int
    maxbytes, int *p_nbytes);
```

Arguments

prob	The current SLP problem.
row	Integer holding the row index for the formula.
formula	Character buffer in which the formula will be placed in the same format as used for input from a file. The formula will be null terminated. This argument may be NULL if maxbytes is zero.
maxbytes	Length of the formula buffer.
p_nbytes	Will be set to the length of the formula, not including the null terminator.

Example

The following retrieves a formula in text form:

```
char *buffer;
int len;

XPRSnlpgetformulastr(prob, 2, NULL, 0, &len);
buffer = malloc(len + 1);
XPRSnlpgetformulastr(prob, 2, buffer, len + 1, NULL);
```

Further information

If the length of the formula would exceed `maxbytes-1`, the formula is truncated to the last token that will fit, and the (partial) formula is terminated with a null character.

Related topics

XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPgetformulastring, XPRSnlpgetformulastring

Purpose

This subroutine is deprecated and will be removed in a future release. Use XPRSnlpgetformulastring instead.

Retrieve a single matrix formula in a character string.

Topic area

Problem Information

Synopsis

```
int XPRS_CC XSLPgetformulastring(XSLPprob prob, int row, char *formula, int
    maxbytes);
```

Arguments

prob	The current SLP problem.
row	Integer holding the row index for the formula.
formula	Character buffer in which the formula will be placed in the same format as used for input from a file. The formula will be null terminated.
maxbytes	Maximum length of returned formula.

Return value

0	Normal return.
32	Formula is too long for the buffer and has been truncated.
other	Error.

Example

The following retrieves a formula in text form:

```
char Buffer[60];
int Code;

Code = XSLPgetformulastring(prob, 2, Buffer, 60);
switch (Code) {
case 0:  printf("\nFormula is %s", Buffer);
        break;
case 32: printf("\nFormula is too long for the buffer");
        break;
default: printf("\nError accessing formula");
        break;
}
```

Further information

If the length of the formula would exceed `maxbytes-1`, the formula is truncated to the last token that will fit, and the (partial) formula is terminated with a null character.

Related topics

XPRSnlpgetformulastring, XPRSnlpchgformulastring, XSLPchgformula, XSLPaddformulas, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPgetformularows, XPRSnlpgetformularows

Purpose

Retrieve the list of positions of the nonlinear formulas in the problem

Topic area

Problem Information

Synopsis

```
int XPRS_CC XSLPgetformularows(XSLPprob prob, int *p_nformulas, int[]
    rowind);
```

Arguments

prob	The current SLP problem.
p_nformulas	Integer used to return the total number of nonlinear formulas in the problem.
rowind	Integer array used for returning the row positions of the nonlinear formulas. May be NULL if not required.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula, XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPgetindex

Purpose

Retrieve the index of an Xpress NonLinear entity with a given name

Topic areas

Problem Information, User Functions

Synopsis

```
int XPRS_CC XSLPgetindex(XSLPprob prob, int type, char *name, int
                        *p_index);
```

Arguments

prob	The current SLP problem.
type	type of entity. The following are defined: XSLP_USERFUNCNames (=6) User functions; XSLP_INTERNALFUNCNames (=7) Internal functions; XSLP_USERFUNCNamesNOCASE (=8) User functions, case insensitive; XSLP_INTERNALFUNCNamesNOCASE (=9) Internal functions, case insensitive; The constants 1 (for row names) and 2 (for column names) may also be used.
name	Character string containing the name, terminated by a null character.
p_index	Integer to receive the index of the item.

Example

The following example retrieves the index of the internal SIN function using both an upper-case and a lower case version of the name.

```
int UpperIndex, LowerIndex;
XSLPgetindex(prob, XSLP_INTERNALFUNCNamesNOCASE,
             "SIN", &UpperIndex);
XSLPgetindex(prob, XSLP_INTERNALFUNCNamesNOCASE,
             "sin", &LowerIndex);
```

UpperIndex and LowerIndex will contain the same value because the search was made using case-insensitive matching.

Further information

All entities count from 1. This includes the use of 1 or 2 (row or column) for type. A value of zero returned in p_index means there is no matching item. The case-insensitive types will find the first match regardless of the case of name or of the defined function.

XSLPgetintattrib

Purpose

Retrieve the value of an integer problem attribute

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPgetintattrib(XSLPprob prob, int attrib, int *p_value);
```

Arguments

prob	The current SLP problem.
attrib	attribute (SLP or optimizer) whose value is to be returned.
p_value	Address of an integer variable to receive the value.

Further information

Both SLP and optimizer attributes can be retrieved using this function. If an optimizer attribute is requested, the return value will be the same as that from [XPRSgetintattrib](#), which can similarly be used to obtain both Optimizer and SLP attributes.

Related topics

XSLPgetdblattrib, XSLPgetstrattrib

XSLPgetintcontrol

Purpose

Retrieve the value of an integer problem control

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPgetintcontrol(XSLPprob prob, int control, int *p_value);
```

Arguments

prob	The current SLP problem.
control	control (SLP or optimizer) whose value is to be returned.
p_value	Address of an integer variable to receive the value.

Example

The following example retrieves the value of the Xpress NonLinear control XSLP_ALGORITHM and of the optimizer control XPRS_DEFAULTALG:

```
int Algorithm, DefaultAlg;
XSLPgetintcontrol(prob, XSLP_ALGORITHM, &Algorithm);
XSLPgetintcontrol(prob, XPRS_DEFAULTALG, &DefaultAlg);
```

Further information

Both SLP and optimizer controls can be retrieved using this function. If an optimizer control is requested, the return value will be the same as that from [XPRSgetintcontrol](#), which can similarly be used to obtain both Optimizer and SLP controls.

Related topics

XSLPgetdblcontrol, XSLPgetstrcontrol, XSLPsetintcontrol

XSLPgetlasterror

Purpose

Retrieve the error message corresponding to the last Xpress NonLinear error during an SLP run

Topic area

Misc

Synopsis

```
int XPRS_CC XSLPgetlasterror(XSLPprob prob, int *p_code, char *msg);
```

Arguments

prob	The current SLP problem.
p_code	Address of an integer to receive the message number of the last error. May be NULL if not required.
msg	Character buffer to receive the error message. The error message will never be longer than 256 characters. May be NULL if not required.

Example

The following example checks the return code from reading a matrix. If the code is nonzero then an error has occurred, and the error number is retrieved for further processing.

```
int Error, code;
if (Error=XSLPreadprob(prob, "Matrix", "")) {
    XSLPgetlasterror(prob, &code, NULL);
    MyErrorHandler(code);
}
```

Further information

In general, Xpress NonLinear functions return a value of 32 to indicate a non-recoverable error. XSLPgetlasterror can retrieve the actual error number and message.

In case no SLP error code was returned, the function will check the underlying XPRS library for any errors reported.

XSLPgetptrattrib

Purpose

Retrieve the value of a problem pointer attribute

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPgetptrattrib(XSLPprob prob, int attrib, void **p_value);
```

Arguments

prob	The current SLP problem.
attrib	attribute whose value is to be returned.
p_value	Address of a pointer to receive the value.

Example

The following example retrieves the value of the Xpress NonLinear pointer attribute XSLP_XPRSPROBLEM which is the underlying optimizer problem pointer:

```
XPRSprob xprob;  
XSLPgetptrattrib(prob, XSLP_XPRSPROBLEM, &xprob);
```

Further information

This function is normally used to retrieve the underlying optimizer problem pointer, as shown in the example.

Related topics

XSLPgetdblattrib, XSLPgetintattrib, XSLPgetstrattrib

XSLPgetrowinfo, XPRSslpgetrowinfo

Purpose

Get current row information.

Topic areas

SLP, Solution

Synopsis

```
int XSLP_CC XSLPgetrowinfo(XSLPprob prob, int type, int row, XSLPalltype
    *p_info);
```

Arguments

prob	The current SLP problem
type	Type of information (see below)
row	Index of the row whose information is to be handled
p_info	Address of information to be set or retrieved

Example

The following example prints the current penalty factor of row 1:

```
XSLPalltype alltype;
XSLPgetrowinfo(prob, XSLP_ROWINFO_CURRENTPENALTYFACTOR, 1, &alltype);
printf("penaltyfactor[1]=%f\n", alltype.value.real);
```

Further information

1. If the data is not available, the type of the returned p_info is set to XSLPtype_undefined.
Please refer to the header file `xprs.h` for the definition of XPRSalltype which is identical to XSLPalltype.
The following constants are provided for row information handling:

XSLP_ROWINFO_SLACK	Get the current slack value of the row
XSLP_ROWINFO_DUAL	Get the current dual multiplier of the row
XSLP_ROWINFO_NUMPENALTYERRORS	Get the number of times the penalty error vector has been active for the row
XSLP_ROWINFO_MAXPENALTYERROR	Get the maximum size of the penalty error vector activity for the row
XSLP_ROWINFO_TOTALPENALTYERROR	Get the total size of the penalty error vector activity for the row
XSLP_ROWINFO_CURRENTPENALTYERROR	Get the size of the penalty error vector activity in the current iteration for the row
XSLP_ROWINFO_CURRENTPENALTYFACTOR	Set the size of the penalty error factor for the current iteration for the row
XSLP_ROWINFO_PENALTYCOLUMNPLUS	Get the index of the positive penalty column for the row (+)
XSLP_ROWINFO_PENALTYCOLUMNPLUSVALUE	Get the value of the positive penalty column for the row (+)
XSLP_ROWINFO_PENALTYCOLUMNPLUSDJ	Get the reduced cost of the positive penalty column for the row (+)
XSLP_ROWINFO_PENALTYCOLUMNMINUS	Get the index of the negative penalty column for the row (-)
XSLP_ROWINFO_PENALTYCOLUMNMINUSVALUE	Get the value of the negative penalty column for the row (-)
XSLP_ROWINFO_PENALTYCOLUMNMINUSDJ	Get the reduced cost of the negative penalty column for the row (-)
2. This function can be used to get the current SLP iteration solution in the case where this is not available using XSLPgetslpsol due to solution filtering via XSLP_FILTER.

XSLPgetrowstatus, XPRSslpgetrowstatus

Purpose

Retrieve the status setting of a constraint

Topic areas

SLP, Solution, Bit-vector

Synopsis

```
int XPRS_CC XSLPgetrowstatus(XSLPprob prob, int row, int *p_status);
```

Arguments

prob	The current SLP problem.
row	The index of the matrix row whose data is to be obtained.
p_status	Address of an integer holding a bitmap to receive the status settings.

Example

This recovers the status of the rows of the matrix of the current problem and reports those which are flagged as enforced constraints.

```
int iRow, nRow, status;
XSLPgetintattrib(prob, XPRS_ROWS, &nRow);
for (iRow=0; iRow<nRow; iRow++) {
    XSLPgetrowstatus(prob, iRow, &status);
    if (status & 0x800) printf("\nRow %d is enforced");
}
```

Further information

See the section on bitmap settings for details on the possible information in p_status.

Related topics

XSLPchgrowstatus

XSLPgetrowwt, XPRSslpgetrowwt

Purpose

Get the initial penalty error weight for a row

Topic areas

SLP, Data Information

Synopsis

```
int XSLP_CC XSLPgetrowwt(XSLPprob prob, int row, double *p_weight)
```

Arguments

prob	The current SLP problem.
row	The index of the row whose weight is to be retrieved.
p_weight	Address of a double precision variable to receive the value of the weight.

Example

The following example gets the initial weight of row number 2.

```
double weight;  
XSLPgetrowwt(prob, 2, &weight)
```

Further information

The initial row weight is used only when the augmented structure is created. After that, the current weighting can be accessed using `XSLPgetrowinfo`.

Related topics

`XSLPchgrowwt`, `XSLPgetrowinfo`

XSLPgetslpsol, XPRSgetnlpsol

Purpose

This subroutine is deprecated and will be removed in a future release. Use XPRSgetsolution and related functions instead.

Obtain the current SLP solution values

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPgetslpsol(XSLPprob prob, double *x, double *slack, double
    *duals, double *djs);
```

Arguments

prob	The current SLP problem.
x	Double array of length XSLP_ORIGINALCOLS to hold the values of the primal variables. May be NULL if not required.
slack	Double array of length XSLP_ORIGINALROWS to hold the values of the slack variables. May be NULL if not required.
duals	Double array of length XSLP_ORIGINALROWS to hold the values of the dual variables. May be NULL if not required.
djs	Double array of length XSLP_ORIGINALCOLS to hold the reduced costs of the primal variables. May be NULL if not required.

Example

The following code fragment recovers the values and reduced costs of the primal variables for the current SLP solution:

```
XSLPprob prob;
int nCol;
double *val, *djs;
XSLPgetintattrib(prob, XSLP_ORIGINALCOLS, &nCol);
val = malloc(nCol*sizeof(double));
djs = malloc(nCol*sizeof(double));
XSLPgetslpsol(prob, val, NULL, NULL, djs);
```

Further information

The behavior of this function depends on the XSLP_FILTER control. If XSLP_FILTER_KEEPEST is set (which is the default), then XSLPgetslpsol will return the best solution seen so far, otherwise the solution of the most recent SLP iteration will be returned. The current iteration solution (in the presolved space) can be obtained from XSLPgetcolinfo.

XSLPgetslpsol can be called at any time after an SLP iteration has completed, and will return the same values even if the problem is subsequently changed. XSLPgetslpsol returns the solution in the original space. To access the values of any augmentation columns or rows, use XPRSgetslpsol; accessing the augmented solution is only recommended if XSLP_PRESOLVELEVEL indicates that the problem dimensions should not be changed in presolve.

XSLPgetstrattrib

Purpose

Retrieve the value of a string problem attribute

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPgetstrattrib(XSLPprob prob, int attrib, char *value);
```

Arguments

prob	The current SLP problem.
attrib	attribute (SLP or optimizer) whose value is to be returned.
value	Character buffer to receive the value.

Example

The following example retrieves the value of the Xpress NonLinear attribute XSLP_VERSIONDATE and of the optimizer attribute XPRS_MATRIXNAME:

```
char VersionDate[200], MatrixName[200];  
XSLPgetstrattrib(prob, XSLP_VERSIONDATE, VersionDate);  
XSLPgetstrattrib(prob, XPRS_MATRIXNAME, MatrixName);
```

Further information

Both SLP and optimizer attributes can be retrieved using this function. If an optimizer attribute is requested, the return value will be the same as that from [XPRSgetstrattrib](#), which can similarly be used to obtain both Optimizer and SLP attributes.

Related topics

[XSLPgetdblattrib](#), [XSLPgetintattrib](#)

XSLPgetstrcontrol

Purpose

Retrieve the value of a string problem control

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPgetstrcontrol(XSLPprob prob, int control, char *value);
```

Arguments

prob	The current SLP problem.
control	control (SLP or optimizer) whose value is to be returned.
value	Character buffer to receive the value.

Further information

Both SLP and optimizer controls can be retrieved using this function. If an optimizer control is requested, the return value will be the same as that from [XPRSgetstrcontrol](#), which can similarly be used to obtain both Optimizer and SLP controls.

Related topics

[XSLPgetdblcontrol](#), [XSLPgetintcontrol](#), [XSLPsetstrcontrol](#)

XSLPgettolset

Purpose

This subroutine is deprecated and will be removed in a future release.

Retrieve the values of a set of convergence tolerances for an SLP problem.

Topic areas

SLP, SLP-convergence

Synopsis

```
int XPRS_CC XSLPgettolset(XSLPprob prob, int tolset, int *p_status, double
    *tols);
```

Arguments

prob	The current SLP problem.
tolset	The index of the tolerance set.
p_status	Address of integer to receive the bit-map of status settings. May be NULL if not required.
tols	Array of 9 double-precision values to hold the tolerances. May be NULL if not required.

Example

The following example retrieves the values for tolerance set 3 and prints those which are set:

```
double tols[9];
int i, status;
XSLPgettolset(prob, 3, &status, tols);
for (i=0; i<9; i++)
    if (status & (1<<i))
        printf("\nTolerance %d = %lg", i, tols[i]);
```

Further information

If p_status or tols is NULL, then the corresponding information will not be returned.

If tols is not NULL, then a set of 9 values will always be returned. p_status indicates which of these values is active as follows. Bit n of p_status is set if tols[n] is active, where n is:

Entry / Bit	Tolerance	XSLP constant	XSLP bit constant
0	Closure tolerance (TC)	XSLP_TOLSET_TC	XSLP_TOLSETBIT_TC
1	Absolute delta tolerance (TA)	XSLP_TOLSET_TA	XSLP_TOLSETBIT_TA
2	Relative delta tolerance (RA)	XSLP_TOLSET_RA	XSLP_TOLSETBIT_RA
3	Absolute coefficient tolerance (TM)	XSLP_TOLSET_TM	XSLP_TOLSETBIT_TM
4	Relative coefficient tolerance (RM)	XSLP_TOLSET_RM	XSLP_TOLSETBIT_RM
5	Absolute impact tolerance (TI)	XSLP_TOLSET_TI	XSLP_TOLSETBIT_TI
6	Relative impact tolerance (RI)	XSLP_TOLSET_RI	XSLP_TOLSETBIT_RI
7	Absolute slack tolerance (TS)	XSLP_TOLSET_TS	XSLP_TOLSETBIT_TS
8	Relative slack tolerance (RS)	XSLP_TOLSET_RS	XSLP_TOLSETBIT_RS

The XSLP_TOLSET constants can be used to access the corresponding entry in the value arrays, while the XSLP_TOLSETBIT constants are used to set or retrieve which tolerance values are used for a given SLP variable.

Related topics

Related topics

XSLPaddtolsets, XSLPchgtolset, XSLPdeltolsets, XSLPloadtolsets

XSLPgetvar

Purpose

This subroutine is deprecated and will be removed in a future release.

Retrieve information about an SLP variable.

Topic area

Data Information

Synopsis

```
int XPRS_CC XSLPgetvar(XSLPprob prob, int col, int *p_detrow, double
    *p_initstepbound, double *p_stepbound, double *p_penalty, double
    *p_damp, double *p_initial, double *p_value, int *p_tolset, int
    *p_history, int *p_converged, int *p_vartype, int *p_delta, int
    *p_penaltydelta, int *p_updaterow, double *p_old);
```

Arguments

prob	The current SLP problem.
col	The index of the column.
p_detrow	Address of an integer to receive the index of the determining row. May be NULL if not required.
p_initstepbound	Address of a double precision variable to receive the value of the initial step bound of the variable. May be NULL if not required.
p_stepbound	Address of a double precision variable to receive the value of the current step bound of the variable. May be NULL if not required.
p_penalty	Address of a double precision variable to receive the value of the penalty delta weighting of the variable. May be NULL if not required.
p_damp	Address of a double precision variable to receive the value of the current damping factor of the variable. May be NULL if not required.
p_initial	Address of a double precision variable to receive the value of the initial value of the variable. May be NULL if not required.
p_value	Address of a double precision variable to receive the current activity of the variable. May be NULL if not required.
p_tolset	Address of an integer to receive the index of the tolerance set of the variable. May be NULL if not required.
p_history	Address of an integer to receive the SLP history of the variable. May be NULL if not required.
p_converged	Address of an integer to receive the convergence status of the variable as defined in the "Convergence Criteria" section (The returned value will match the numbering of the tolerances). May be NULL if not required.
p_vartype	Address of an integer to receive the status settings (a bitmap defining the existence of certain properties for this variable). The following bits are defined: Bit 1: Variable has a delta vector Bit 2: Variable has an initial value Bit 14: Variable is the reserved "=" column Other bits are reserved for internal use. May be NULL if not required.
p_delta	Address of an integer to receive the index of the delta vector for the variable. May be NULL if not required.
p_penaltydelta	Address of an integer to receive the index of the first penalty delta vector for the variable. The second penalty delta immediately follows the first. May be NULL if not required.
p_updaterow	Address of an integer to receive the index of the update row for the variable. May be

	NULL if not required.
p_old	Address of a double precision variable to receive the value of the variable at the previous SLP iteration. May be NULL if not required.

Example

The following example retrieves the current value, convergence history and status for column 3.

```
int converged, history;
double value;

XSLPgetvar(prob, 3, NULL, NULL, NULL,
           NULL, NULL, NULL, &value,
           NULL, &history, &converged,
           NULL, NULL, NULL, NULL, NULL);
```

Further information

If col refers to a column which is not an SLP variable, then all the return values will indicate that there is no corresponding data.

p_detrow will be set to -1 if there is no determining row.

p_delta, p_penaltydelta and p_updaterow will be set to -1 if there is no corresponding item.

Current values, deltas and step bounds can also be obtained individually through XSLPgetcolinfo.

Related topics

XSLPaddvars, XSLPgetcolinfo, XSLPchgvar, XSLPdelvars, XSLPloadvars

XSLPimportlibfunc, XPRSnlpimportlibfunc

Purpose

Imports a function from a library file to be called as a user function

Topic area

User Functions

Synopsis

```
int XPRS_CC XSLPimportlibfunc(XSLPprob prob, const char * libname, const
    char * funcname, XPRSfunctionptraddr p_function, int * p_status );
```

Arguments

prob	The current SLP problem.
libname	Filename of the library.
funcname	Fucntion name inside the library.
p_function	Function pointer to return the loaded function.
p_status	Outcome of the load operation
	0 success.
	1 library file not found.
	2 library function in library file not found.

Further information

On systems where necessary, Xpress will hold the handle of the library opened and free up when the problem object `prob` is destroyed. The type `XPRSfunctionptraddr` is a pointer to a generic function pointer.

Related topics

`XSLPadduserfunction`, `XSLPdeluserfunction`

XSLPinit

Purpose

Initializes the Xpress NonLinear system

Topic area

Licensing

Synopsis

```
int XPRS_CC XSLPinit();
```

Argument

none

Example

The following example initiates the Xpress NonLinear system and prints the banner.

```
char Buffer[256];
XPRSinit();
XSLPinit();
XSLPgetbanner(Buffer);
```

XPRSinit initializes the Xpress optimizer; XSLPinit then initializes the SLP module, so that the banner contains information from both systems.

Further information

XSLPinit must be the first call to the Xpress NonLinear system except for XSLPgetbanner and XSLPgetversion. It initializes any global parts of the system if required. The call to XSLPinit must be preceded by a call to XPRSinit to initialize the Optimizer Library part of the system first.

Related topics

XSLPfree

XSLPInterrupt

Purpose

Interrupts the current SLP optimization

Topic area

Solution Process

Synopsis

```
int XPRS_CC XSLPInterrupt(int reason);
```

Arguments

prob	The current SLP problem.
reason	Interrupt code to be propagated.

Further information

Provides functionality to stop the SLP optimization process from inside a callback. The following constants are provided for the parameter value:

Value 1	XSLP_STOP_TIMELIMIT
Value 2	XSLP_STOP_CTRLC
Value 3	XSLP_STOP_NODELIMIT
Value 4	XSLP_STOP_ITERLIMIT
Value 5	XSLP_STOP_MIPGAP
Value 6	XSLP_STOP_SOLLIMIT
Value 9	XSLP_STOP_USER

XSLPitemname

Purpose

Retrieves the name of an Xpress NonLinear entity or the value of a function token as a character string.

Topic area

Names Manager

Synopsis

```
int XPRS_CC XSLPitemname(XSLPprob prob, int type, double value, char
    *buffer);
```

Arguments

prob	The current SLP problem.
type	Integer holding the type of Xpress NonLinear entity. This can be any one of the token types described in the section on 13.
value	Double precision value holding the index or value of the token. The use and meaning of the value is as described in the section on 13.
buffer	Character buffer to hold the result, which will be terminated with a null character.

Example

The following example displays the formula for the coefficient in row 2, column 3 in unparsed form:

```
int n, type[10];
double value[10];
char buffer[60];
int TokenCount;

XSLPgetcoeffformula(prob, 2, 3, &Factor, 0, 10, &TokenCount, type, value);

printf("\n");
for (n=0; type[n] != XSLP_EOF; n++) {
    XSLPitemname(prob, type[n], value[n], buffer);
    printf(" %s", buffer);
}
```

Further information

If a name has not been provided for an Xpress NonLinear entity, then an internally-generated name will be used.

Numerical values will be formatted as fixed-point or floating-point depending on their size.

XSLPloadcoefs, XPRSslploadcoefs

Purpose

Load non-linear coefficients into the SLP problem. For a simpler version of this function see XSLPloadformulas.

Topic areas

SLP, Problem Information

Synopsis

```
int XPRS_CC XSLPloadcoefs(XSLPprob prob, int ncoefs, const int[] rowind,
    const int[] colind, const double[] factor, const int[] formulastart,
    int parsed, const int[] type, const double[] coef);
```

Arguments

prob	The current SLP problem.
ncoefs	Number of non-linear coefficients to be loaded.
rowind	Integer array holding index of row for the coefficient.
colind	Integer array holding index of column for the coefficient.
factor	Double array holding factor by which formula is scaled. If this is NULL, then a value of 1.0 will be used.
formulastart	Integer array of length ncoefs+1 holding the start position in the arrays type and coef of the formula for the coefficients. formulastart[ncoefs] should be set to the next position after the end of the last formula.
parsed	Integer indicating whether the token arrays are formatted as internal unparsed (parsed=0) or internal parsed reverse Polish (parsed=1).
type	Array of token types providing the formula for each coefficient.
coef	Array of values corresponding to the types in type.

Example

Assume that the rows and columns of prob are named Row1, Row2 ..., Col1, Col2 ... The following example loads coefficients representing:

Col2 * Col3 + Col6 * Col2^2 into Row1 and

Col2 ^ 2 into Row3.

```
int rowind[3], colind[3], formulastart[4], type[8];
int n, ncoefs;
double coef[8];

rowind[0] = 1; colind[0] = 2;
rowind[1] = 1; colind[1] = 6;
rowind[2] = 3; colind[2] = 2;

n = ncoefs = 0;
formulastart[ncoefs++] = n;
type[n] = XSLP_COL; coef[n++] = 3;
type[n++] = XSLP_EOF;

formulastart[ncoefs++] = n;
type[n] = XSLP_COL; coef[n++] = 2;
type[n] = XSLP_COL; coef[n++] = 2;
type[n] = XSLP_OP; coef[n++] = XSLP_MULTIPLY;
type[n++] = XSLP_EOF;

formulastart[ncoefs++] = n;
type[n] = XSLP_COL; coef[n++] = 2;
type[n++] = XSLP_EOF;
```

```

formulastart[ncoefs] = n;

XSLPloadcoefs(prob, ncoefs, rowind, colind,
              NULL, formulastart, 1, type, coef);

```

The first coefficient in Row1 is in Col2 and has the formula Col3, so it represents $\text{Col2} * \text{Col3}$.

The second coefficient in Row1 is in Col6 and has the formula $\text{Col2} * \text{Col2}$ so it represents $\text{Col6} * \text{Col2}^2$.

The formulae are described as *parsed* (parsed=1), so the formula is written as

Col2 Col2 *

rather than the unparsed form

Col2 * Col2

The last coefficient, in Row3, is in Col2 and has the formula Col2, so it represents $\text{Col2} * \text{Col2}$.

Further information

1. The j^{th} coefficient is made up of two parts: factor and Formula. factor is a constant multiplier, which can be provided in the factor array. If Xpress NonLinear can identify a constant factor in Formula, then it will use that as well, to minimize the size of the formula which has to be calculated. Formula is made up of a list of tokens in type and coef starting at formulastart[j]. The tokens follow the rules for parsed or unparsed formulae as indicated by the setting of parsed. The formula must be terminated with an XSLP_EOF token. If several coefficients share the same formula, they can have the same value in formulastart. For possible token types and values see the chapter on 13.
2. The load functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding add functions add or replace items leaving other items of the same type unchanged.

Related topics

XPRSnlpgetformulastr, XSLPaddformulas, XPRSnlpchgformulastr, XSLPchgformula,
XSLPloadformulas, XSLPgetformularows, XSLPgetformula, XSLPdelformulas

XSLPloaddfs

Purpose

This subroutine is deprecated and will be removed in a future release.

Load a set of distribution factors.

Topic areas

SLP, Data Input

Synopsis

```
int XSLP_CC XSLPloaddfs(XSLPprob prob, int ndfs, const int *colind, const
    int *rowind, const double *value)
```

Arguments

prob	The current SLP problem.
ndfs	The number of distribution factors.
colind	Array of indices of columns whose distribution factor is to be changed.
rowind	Array of indices of the rows where each distribution factor applies.
value	Array of double precision variables holding the new values of the distribution factors.

Example

The following example loads distribution factors as follows:

column 282 in row 134 = 0.1

column 282 in row 136 = 0.15

column 285 in row 133 = 1.0.

Any other first-order derivative placeholders are set to XSLP_DELTA_Z.

```
int colind[3], rowind[3];
double value[3];
colind[0] = 282; rowind[0] = 134; value[0] = 0.1;
colind[1] = 282; rowind[1] = 136; value[1] = 0.15;
colind[2] = 285; rowind[2] = 133; value[2] = 1.0;
XSLPloaddfs(prob, 3, colind, rowind, value);
```

Further information

The *distribution factor* of a column in a row is the matrix coefficient of the corresponding delta vector in the row. Distribution factors are used in conventional recursion models, and are essentially normalized first-order derivatives. Xpress-SLP can accept distribution factors instead of initial values, provided that the values of the variables involved can all be calculated after optimization using determining rows, or by a callback.

The add functions load additional items into the SLP problem. The corresponding load functions delete any existing items first.

Related topics

XSLPadddfs, XSLPchgdf, XSLPgetdf

XSLPloadformulas, XPRSnlploadformulas

Purpose

Load non-linear formulas into the SLP problem

Topic area

Problem Information

Synopsis

```
int XPRS_CC XSLPloadformulas(XSLPprob prob, int nnlpcoefs, const int[]
    rowind, const int[] formulastart, int parsed, const int[] type, const
    double[] value);
```

Arguments

prob	The current SLP problem.
nnlpcoefs	Number of non-linear coefficients to be loaded.
rowind	Integer array holding index of row for the coefficient.
formulastart	Integer array of length nnlpcoefs+1 holding the start position in the arrays type and value of the formula for the coefficients. formulastart[nnlpcoefs] should be set to the next position after the end of the last formula.
parsed	Integer indicating whether the token arrays are formatted as internal unparsed (parsed=0) or internal parsed reverse Polish (parsed=1).
type	Array of token types providing the formula for each coefficient.
value	Array of values corresponding to the types in type.

Example

Assume that the rows and columns of prob are named Row0, Row1 ..., Col0, Col1 ... The following example adds coefficients representing:

Col2 * Col3 * Col6^2 into Row1 and

Col2 * Col3 ^ 2 into Row3.

```
int rowind[3], formulastart[4], type[8];
int n, nnlpcoefs;
double value[8];

rowind[0] = 1;
rowind[1] = 1;
rowind[2] = 3;

n = nnlpcoefs = 0;
formulastart[nnlpcoefs++] = n;
type[n] = XSLP_COL; value[n++] = 3;
type[n++] = XSLP_EOF;

formulastart[nnlpcoefs++] = n;
type[n] = XSLP_COL; value[n++] = 2;
type[n] = XSLP_COL; value[n++] = 3;
type[n] = XSLP_OP; value[n++] = XSLP_MULTIPLY;
type[n] = XSLP_COL; value[n++] = 6;
type[n] = XSLP_CON; value[n++] = 2;
type[n] = XSLP_OP; value[n++] = XSLP_EXPONENT;
type[n] = XSLP_OP; value[n++] = XSLP_MULTIPLY;
type[n++] = XSLP_EOF;

formulastart[nnlpcoefs++] = n;
type[n] = XSLP_COL; value[n++] = 2;
type[n] = XSLP_COL; value[n++] = 3;
type[n] = XSLP_CON; value[n++] = 2;
```

```

type[n] = XSLP_OP;  value[n++] = XSLP_EXPONENT;
type[n] = XSLP_OP;  value[n++] = XSLP_MULTIPLY;
type[n++] = XSLP_EOF;

formulastart[nnlpcoefs] = n;

XSLPloadformulas(prob, nnlpcoefs, rowind, formulastart, 1, type, value);

```

Further information

Formula *j* is made up of a list of tokens in *type* and *value* starting at *formulastart[j]*. The tokens follow the rules for parsed or unparsed formulae as indicated by the setting of *parsed*. The formula must be terminated with an *XSLP_EOF* token. If several formulas share the same nonlinear expressions, they can have the same value in *formulastart*. For possible token types and values see the chapter on 13.

The *load* functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding *add* functions add or replace items leaving other items of the same type unchanged.

Related topics

XPRSnlpgetformulastr, *XSLPaddformulas*, *XPRSnlpchgformulastr*, *XSLPchgformula*,
XSLPloadformulas, *XSLPgetformularows*, *XSLPgetformula*, *XSLPdelformulas*

XSLPloadtolsets

Purpose

This subroutine is deprecated and will be removed in a future release.
Load sets of standard tolerance values into an SLP problem.

Topic areas

SLP, SLP-convergence

Synopsis

```
int XPRS_CC XSLPloadtolsets(XSLPprob prob, int ntolsets, double *tols);
```

Arguments

prob	The current SLP problem.
ntolsets	The number of tolerance sets to be loaded.
tols	Double array of (ntolsets * 9) items containing the 9 tolerance values for each set in order.

Example

The following example creates two tolerance sets: the first has values of 0.005 for all tolerances; the second has values of 0.001 for relative tolerances (numbers 2,4,6,8), values of 0.01 for absolute tolerances (numbers 1,3,5,7) and zero for the closure tolerance (number 0).

```
double tols[18];
for (i=0;i<9;i++) tols[i] = 0.005;
tol[9] = 0;
for (i=10;i<18;i=i+2) tols[i] = 0.01;
for (i=11;i<18;i=i+2) tols[i] = 0.001;
XSLPloadtolsets(prob, 2, tols);
```

Further information

A tolerance set is an array of 9 values containing the following tolerances:

Entry / Bit	Tolerance	XSLP constant	XSLP bit constant
0	Closure tolerance (TC)	XSLP_TOLSET_TC	XSLP_TOLSETBIT_TC
1	Absolute delta tolerance (TA)	XSLP_TOLSET_TA	XSLP_TOLSETBIT_TA
2	Relative delta tolerance (RA)	XSLP_TOLSET_RA	XSLP_TOLSETBIT_RA
3	Absolute coefficient tolerance (TM)	XSLP_TOLSET_TM	XSLP_TOLSETBIT_TM
4	Relative coefficient tolerance (RM)	XSLP_TOLSET_RM	XSLP_TOLSETBIT_RM
5	Absolute impact tolerance (TI)	XSLP_TOLSET_TI	XSLP_TOLSETBIT_TI
6	Relative impact tolerance (RI)	XSLP_TOLSET_RI	XSLP_TOLSETBIT_RI
7	Absolute slack tolerance (TS)	XSLP_TOLSET_TS	XSLP_TOLSETBIT_TS
8	Relative slack tolerance (RS)	XSLP_TOLSET_RS	XSLP_TOLSETBIT_RS

The XSLP_TOLSET constants can be used to access the corresponding entry in the value arrays, while the XSLP_TOLSETBIT constants are used to set or retrieve which tolerance values are used for a given SLP variable.

Once created, a tolerance set can be used to set the tolerances for any SLP variable.

If a tolerance value is zero, then the default tolerance will be used instead. To force the use of a tolerance, use the XSLPchgtolset function and set the Status variable appropriately.

See the section "Convergence Criteria" for a fuller description of tolerances and their uses.

The load functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding add functions add or replace items leaving other items of the same type unchanged.

Related topics

XSLPaddtolsets, XSLPdeltolsets, XSLPchgtolset, XSLPgettolset

XSLPloadvars

Purpose

This subroutine is deprecated and will be removed in a future release.

Load SLP variables defined as matrix columns into an SLP problem.

Topic area

Data Input

Synopsis

```
int XPRS_CC XSLPloadvars(XSLPprob prob, int nvars, int *colind, int
    *vartype, int *detrow, int *seqnum, int *tolind, double *initial,
    double *stepbound);
```

Arguments

prob	The current SLP problem.
nvars	The number of SLP variables to be loaded.
colind	Integer array holding the index of the matrix column corresponding to each SLP variable.
vartype	Bitmap giving information about the SLP variable as follows: Bit 1 Variable has a delta vector; Bit 2 Variable has an initial value; Bit 14 Variable is the reserved "=" column; May be NULL if not required.
detrow	Integer array holding the index of the determining row for each SLP variable (a negative value means there is no determining row) May be NULL if not required.
seqnum	Integer array holding the index sequence number for cascading for each SLP variable (a zero value means there is no pre-defined order for this variable) May be NULL if not required.
tolind	Integer array holding the index of the tolerance set for each SLP variable (a zero value means the default tolerances are used) May be NULL if not required.
initial	Double array holding the initial value for each SLP variable (use the vartype bit map to indicate if a value is being provided) May be NULL if not required.
stepbound	Double array holding the initial step bound size for each SLP variable (a zero value means that no initial step bound size has been specified). If a value of XPRS_PLUSINFINITY is used for a value in stepbound, the delta will never have step bounds applied, and will almost always be regarded as converged. May be NULL if not required.

Example

The following example loads two SLP variables into the problem. They correspond to columns 23 and 25 of the underlying LP problem. Column 25 has an initial value of 1.42; column 23 has no specific initial value

```
int colind[2], vartype[2];
double initial[2];

colind[0] = 23; vartype[0] = 0;
colind[1] = 25; vartype[1] = 4; initial[1] = 1.42;

XSLPloadvars(prob, 2, colind, vartype, NULL, NULL,
    NULL, initial, NULL);
```

`initial` is not set for the first variable, because it is not used (`vartype = 0`). Bit 1 of `vartype` is set for the second variable to indicate that the initial value has been set.

The arrays for determining rows, sequence numbers, tolerance sets and step bounds are not used at all, and so have been passed to the function as `NULL`.

Further information

The `load` functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding `add` functions add or replace items leaving other items of the same type unchanged.

Related topics

`XSLPaddvars`, `XSLPchgvar`, `XSLPdelvars`, `XSLPgetvar`

XSLPmaxim

Purpose

This subroutine is deprecated and will be removed in a future release. Please call `XPRSchgobjsense` followed by `XSLPnlpoptimize` or `XPRSoptimize` (noting that the latter will by default solve the MINLP if any integers are present unless providing the `-l` flag) instead.

Maximize an SLP problem

Topic area

Solution Process

Synopsis

```
int XPRS_CC XSLPmaxim(XSLPprob prob, char *flags);
```

Arguments

`prob` The current SLP problem.

`flags` These have the same meaning as for `XPRSmaxim` with the exception of the extra 'c' flag which makes `XSLPmaxim` continue a solve.

Example

The following example reads an SLP problem from file and then maximizes it using the primal simplex optimizer.

```
XSLPreadprob("Matrix", "");
XSLPmaxim(prob, "p");
```

Related controls

Integer

`XSLP_ALGORITHM` Bit map determining the SLP algorithm(s) used in the optimization.

`XSLP_AUGMENTATION` Bit map determining the type of augmentation used to create the linearization.

`XSLP_CASCADE` Bit map determining the type of cascading (recalculation of SLP variable values) used during the SLP optimization.

`XSLP_LOG` Determines the amount of iteration logging information produced.

`XSLP_PRESOLVE` Bit map determining the type of nonlinear presolve used before the SLP optimization starts.

Further information

If `XSLPconstruct` has not already been called, it will be called first, using the augmentation defined by the control variable `XSLP_AUGMENTATION`. If determining rows are provided, then cascading will be invoked in accordance with the setting of the control variable `XSLP_CASCADE`.

Related topics

`XSLPconstruct`, `XSLPminim`, `XSLPnlpoptimize`, `XSLPpresolve`

XSLPminim

Purpose

This subroutine is deprecated and will be removed in a future release. Please call `XPRSchgobjsense` followed by `XSLPnloptimize` or `XPRSoptimize` (noting that the latter will by default solve the MINLP if any integers are present unless providing the `-l` flag) instead.

Minimize an SLP problem

Topic area

Solution Process

Synopsis

```
int XPRS_CC XSLPminim(XSLPprob prob, char *flags);
```

Arguments

`prob` The current SLP problem.

`flags` These have the same meaning as for `XPRSminim` with the exception of the extra 'c' flag which makes `XSLPminim` continue a solve.

Example

The following example reads an SLP problem from file and then minimizes it using the Newton barrier optimizer.

```
XSLPreadprob("Matrix", "");
XSLPminim(prob, "b");
```

Related controls

Integer

`XSLP_ALGORITHM` Bit map determining the SLP algorithm(s) used in the optimization.

`XSLP_AUGMENTATION` Bit map determining the type of augmentation used to create the linearization.

`XSLP_CASCADE` Bit map determining the type of cascading (recalculation of SLP variable values) used during the SLP optimization.

`XSLP_LOG` Determines the amount of iteration logging information produced.

`XSLP_PRESOLVE` Bit map determining the type of nonlinear presolve used before the SLP optimization starts.

Further information

If `XSLPconstruct` has not already been called, it will be called first, using the augmentation defined by the control variable `XSLP_AUGMENTATION`. If determining rows are provided, then cascading will be invoked in accordance with the setting of the control variable `XSLP_CASCADE`.

Related topics

`XSLPconstruct`, `XSLPmaxim`, `XSLPnloptimize`, `XSLPpresolve`

XSLPmsaddcustompreset, XPRSmsaddcustompreset

Purpose

A combined version of XSLPmsaddjob and XSLPmsaddpreset. The preset described is loaded, topped up with the specific settings supplied

Topic areas

Multistart, Data Input

Synopsis

```
int XSLP_CC XSLPmsaddcustompreset( XSLPprob prob, const char *description,
    const int preset, const int maxjobs, const int ninitial, const int
    *colind, const double *initial, const int nintcontrols, const int
    *intcontrolid, const int *intcontrolval, const int ndblcontrols,
    const int *dblcontrolid, const double *dblcontrolval, void *data);
```

Arguments

prob	The current SLP problem.
description	Text description of the job. Used for messaging, may be NULL if not required.
preset	Which preset to load.
maxjobs	Maximum number of jobs to be added to the multistart pool.
ninitial	Number of initial values to set.
colind	Indices of the variables for which to set an initial value. May be NULL if ninitial is zero.
initial	Initial values for the variables for which to set an initial value. May be NULL if ninitial is zero.
nintcontrols	Number of integer controls to set.
intcontrolid	The indices of the integer controls to be set. May be NULL if nintcontrols is zero.
intcontrolval	The values of the integer controls to be set. May be NULL if nintcontrols is zero.
ndblcontrols	Number of double controls to set.
dblcontrolid	The indices of the double controls to be set. May be NULL if ndblcontrols is zero.
dblcontrolval	The values of the double controls to be set. May be NULL if ndblcontrols is zero.
data	Job specific user context pointer to be passed to the multistart callbacks.

Further information

This function allows for repeatedly calling the same multistart preset (e.g. initial values) using different basic controls.

Related topics

XSLPmsaddpreset, XSLPmsaddjob, XSLPmsclear

XSLPmsaddjob, XPRSmsaddjob

Purpose

Adds a multistart job to the multistart pool

Topic areas

Multistart, Data Input

Synopsis

```
int XSLP_CC XSLPmsaddjob( XSLPprob prob, const char *description, const int
    ninitial, const int *colind, const double *initial, const int
    nintcontrols, const int *intcontrolid, const int *intcontrolval,
    const int ndblcontrols, const int *dblcontrolid, const double
    *dblcontrolval, void *data);
```

Arguments

<code>prob</code>	The current SLP problem.
<code>description</code>	Text description of the job. Used for messaging, may be NULL if not required.
<code>ninitial</code>	Number of initial values to set.
<code>colind</code>	Indices of the variables for which to set an initial value. May be NULL if <code>ninitial</code> is zero.
<code>initial</code>	Initial values for the variables for which to set an initial value. May be NULL if <code>ninitial</code> is zero.
<code>nintcontrols</code>	Number of integer controls to set.
<code>intcontrolid</code>	The indices of the integer controls to be set. May be NULL if <code>nintcontrols</code> is zero.
<code>intcontrolval</code>	The values of the integer controls to be set. May be NULL if <code>nintcontrols</code> is zero.
<code>ndblcontrols</code>	Number of double controls to set.
<code>dblcontrolid</code>	The indices of the double controls to be set. May be NULL if <code>ndblcontrols</code> is zero.
<code>dblcontrolval</code>	The values of the double controls to be set. May be NULL if <code>ndblcontrols</code> is zero.
<code>data</code>	Job specific user context pointer to be passed to the multistart callbacks.

Further information

Adds a multistart job, applying the specified initial point and option combinations on top of the base problem, i.e. the options and initial values specified to the function is applied on top of the existing settings.

This function allows for loading empty template jobs, that can then be identified using the data variable.

Related topics

XSLPmsaddpreset, XSLPmsaddcustompreset, XSLPmsclear

XSLPmsaddpreset, XPRSmsaddpreset

Purpose

Loads a preset of jobs into the multistart job pool.

Topic areas

Multistart, Data Input

Synopsis

```
int XSLP_CC XSLPmsaddpreset( XSLPprob prob, const char *description, const
    int preset, const int maxjobs, void *data);
```

Arguments

prob	The current SLP problem.
description	Text description of the preset. Used for messaging, may be NULL if not required.
preset	Which preset to load.
maxjobs	Maximum number of jobs to be added to the multistart pool.
data	Job specific user context pointer to be passed to the multistart callbacks.

Further information

The following presets are defined:

XSLP_MSSET_INITIALVALUES: generate maxjobs number of random base points.

XSLP_MSSET_SOLVERS: load all solvers.

XSLP_MSSET_SLP_BASIC: load the most typical SLP tuning settings. A maximum of maxjobs jobs are loaded.

XSLP_MSSET_SLP_EXTENDED: load a comprehensive set of SLP tuning settings. A maximum of maxjobs jobs are loaded.

XSLP_MSSET_KNITRO_BASIC: load the most typical Knitro tuning settings. A maximum of maxjobs jobs are loaded.

XSLP_MSSET_KNITRO_EXTENDED: load a comprehensive set of Knitro tuning settings. A maximum of maxjobs jobs are loaded.

XSLP_MSSET_INITIALFILTERED: generate maxjobs number of random base points, filtered by a merit function centred on initial feasibility.

See XSLP_MS_MAXBOUNDRANGE for controlling the range in which initial values are generated.

Related topics

XSLPmsaddjob, XSLPmsaddcustompreset, XSLPmsclear

XSLPmsclear, XPRSmsclear

Purpose

Removes all scheduled jobs from the multistart job pool

Topic area

Multistart

Synopsis

```
int XSLP_CC XSLPmsclear( XSLPprob prob);
```

Argument

prob	The current SLP problem.
------	--------------------------

Related topics

XSLPmsaddjob, XSLPmsaddpreset, XSLPmsaddcustompreset

XSLPnlpoptimize, XPRSnlpoptimize

Purpose

Maximize or minimize an SLP problem

Topic area

Solution Process

Synopsis

```
int XPRS_CC XSLPnlpoptimize(XSLPprob prob, const char *flags);
```

Arguments

prob The current SLP problem.
flags These have the same meaning as for XPRSmxim and XPRSminim.

Related controls

Double

XSLP_OBJSENSE Determines the direction of optimization: +1 is for minimization, -1 is for maximization.

Integer

XSLP_ALGORITHM Bit map determining the SLP algorithm(s) used in the optimization.

XSLP_AUGMENTATION Bit map determining the type of augmentation used to create the linearization.

XSLP_CASCADE Bit map determining the type of cascading (recalculation of SLP variable values) used during the SLP optimization.

XSLP_LOG Determines the amount of iteration logging information produced.

XSLP_PRESOLVE Bit map determining the type of nonlinear presolve used before the SLP optimization starts.

Further information

XSLPnlpoptimize is equivalent to XSLPmxim (if XSLP_OBJSENSE = -1) or XSLPminim (if XSLP_OBJSENSE = +1).

If XSLPconstruct has not already been called, it will be called first, using the augmentation defined by the control variable XSLP_AUGMENTATION. If determining rows are provided, then cascading will be invoked in accordance with the setting of the control variable XSLP_CASCADE.

Related topics

XSLPconstruct, XSLPmxim, XSLPminim, XSLPpresolve

XSLPpostsolve, XPRSnlppostsolve

Purpose

Restores the problem to its pre-solve state

Topic area

Presolve

Synopsis

```
int XPRS_CC XSLPpostsolve(XSLPprob prob);
```

Argument

prob The current SLP problem.

Related controls

Integer

XSLP_POSTSOLVE Determines if postsolve is applied automatically.

Further information

If Xpress-SLP was used to solve the problem, postsolve will unconstruct the problem before postsolving (including any reformulation that might have been applied).

Related topics

XSLP_POSTSOLVE

XSLPpresolve

Purpose

Perform a nonlinear presolve on the problem

Topic area

Presolve

Synopsis

```
int XPRS_CC XSLPpresolve(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

The following example reads a problem from file, sets the presolve control, presolves the problem and then solves it (using the "s" flag to enforce a local solve).

```
XSLPreadprob(prob, "Matrix", "");
XSLPsetintcontrol(prob, XSLP_PRESOLVE, 1);
XSLPpresolve(prob);
XSLPnloptimize(prob, "s")
```

Related controls

Integer

XSLP_PRESOLVE Bitmap containing nonlinear presolve options.

Further information

If bit 1 of XSLP_PRESOLVE is not set, no nonlinear presolve will be performed. Otherwise, the presolve will be performed in accordance with the bit settings.. XSLPpresolve is called automatically by XSLPconstruct, so there is no need to call it explicitly unless there is a requirement to interrupt the process between presolve and optimization. XSLPpresolve must be called before XSLPconstruct or any of the SLP optimization procedures..

Related topics

XSLP_PRESOLVE

XSLPprintmemory

Purpose

Print the dimensions and memory allocations for a problem

Topic area

Logging

Synopsis

```
int XPRS_CC XSLPprintmemory(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

The following example loads a problem from file and then prints the dimensions of the arrays.

```
XSLPreadprob(prob, "Matrix1", "");
XSLPprintmemory(prob);
```

The output is similar to the following:

```
Arrays and dimensions:
Array      Item  Used  Max  Allocated  Memory
          Size Items Items      Memory  Control
MemList    28   103   129      4K
String      1  8779 13107     13K  XSLP_MEM_STRING
Xv          16     2   1000     16K  XSLP_MEM_XV
Xvitem     48    11   1000     47K  XSLP_MEM_XVITEM
....
```

Further information

XSLPprintmemory lists the current sizes and amounts used of the variable arrays in the current problem. For each array, the size of each item, the number used and the number allocated are shown, together with the size of memory allocated and, where appropriate, the name of the memory control variable to set the array size. Loading and execution of some problems can be speeded up by setting the memory controls immediately after the problem is created. If an array has to be moved to re-allocate it with a larger size, there may be insufficient memory to hold both the old and new versions; pre-setting the memory controls reduces the number of such re-allocations which take place and may allow larger problems to be solved.

XSLPpintevalinfo, XPRSnlpintevalinfo

Purpose

Print a summary of any evaluation errors that may have occurred during solving a problem

Topic area

Logging

Synopsis

```
int XPRS_CC XSLPpintevalinfo(XSLPprob prob);
```

Argument

prob	The current SLP problem.
------	--------------------------

Related topics

XSLPsetcbcofevalerror

XSLPreadprob

Purpose

Read an Xpress NonLinear extended MPS format matrix from a file into an SLP problem

Topic areas

File IO, Problem Creation

Synopsis

```
int XPRS_CC XSLPreadprob(XSLPprob prob, char *filename, char *flags);
```

Arguments

<code>prob</code>	The current SLP problem.
<code>filename</code>	Character string containing the name of the file from which the matrix is to be read.
<code>flags</code>	Character string containing any flags needed for the input routine:
l	only <code>filename.lp</code> is searched for;
v	use the provided filename verbatim, without appending the <code>.mps</code> , <code>.mat</code> or <code>.lp</code> extension;
z	read a compressed input file.

Example

The following example reads the problem from file "Matrix.mat".

```
XSLPreadprob(prob, "Matrix", "");
```

Further information

XSLPreadprob tries to open the file with an extension of "mat" or, failing that, an extension of "mps". If both fail, the file name will be tried with no extension.

XSLPreadprob is capable to read most Ampl .nl files. To specify that a .nl file is to be read, provide the full filename including the .nl extension.

For details of the format of the file, see the section on Extended MPS format.

Related topics

Extended MPS format, XSLPwriteprob

XSLPmaxim

Purpose

This subroutine is deprecated and will be removed in a future release. Please use XSLPmaxim with the 'c' flag instead.

Continue the maximization of an SLP problem.

Topic area

Solution Process

Synopsis

```
int XPRS_CC XSLPmaxim(XSLPprob prob, char *flags);
```

Arguments

prob	The current SLP problem.
flags	These have the same meaning as for XSLPmaxim.

Example

The following example optimizes the SLP problem for up to 10 SLP iterations. If it has not converged, it saves the file and continues for another 10.

```
int Status;

XSLPsetintcontrol(prob, XSLP_ITERLIMIT, 10);
XSLPmaxim(prob, "");
XSLPgetintattrib(prob, XSLP_STATUS, &Status);
if (Status & XSLP_MAXSLPITERATIONS) {
    XSLPsave(prob);
    XSLPsetintcontrol(prob, XSLP_ITERLIMIT, 20);
    XSLPmaxim(prob, "");
}
```

Further information

This allows Xpress NonLinear to continue the maximization of a problem after it has been terminated, without re-initializing any of the parameters. In particular, the iteration count will resume at the point where it previously stopped, and not at 1.

Related topics

XSLPmaxim, XSLPminim

XSLPpreminim

Purpose

This subroutine is deprecated and will be removed in a future release. Please use XSLPminim with the 'c' flag instead.

Continue the minimization of an SLP problem.

Topic area

Solution Process

Synopsis

```
int XPRS_CC XSLPpreminim(XSLPprob prob, char *flags);
```

Arguments

prob	The current SLP problem.
flags	These have the same meaning as for XSLPminim.

Example

The following example optimizes the SLP problem for up to 10 SLP iterations (using the "s" flag to enforce a local solve). If it has not converged, it saves the file and continues for another 10.

```
int Status;

XSLPsetintcontrol(prob, XSLP_ITERLIMIT, 10);
XSLPminim(prob, "s");
XSLPgetintattrib(prob, XSLP_STATUS, &Status);
if (Status & XSLP_MAXSLPITERATIONS) {
    XSLPsave(prob);
    XSLPsetintcontrol(prob, XSLP_ITERLIMIT, 20);
    XSLPpreminim(prob, "s");
}
```

Further information

This allows Xpress NonLinear to continue the minimization of a problem after it has been terminated, without re-initializing any of the parameters. In particular, the iteration count will resume at the point where it previously stopped, and not at 1.

Related topics

XSLPminim, XSLPpremaxim

XSLPrestore

Purpose

Restore the Xpress NonLinear problem from a file created by XSLPsave

Topic areas

File IO, Save Restore

Synopsis

```
int XPRS_CC XSLPrestore(XSLPprob prob, char *filename);
```

Arguments

prob	The current SLP problem.
filename	Character string containing the name of the problem which is to be restored.

Example

The following example restores a problem originally saved on file "MySave"

```
XSLPrestore(prob, "MySave");
```

Further information

Normally XSLPrestore restores both the Xpress NonLinear problem and the underlying optimizer problem. If only the Xpress NonLinear problem is required, set the integer control variable XSLP_CONTROL appropriately.

Related topics

XSLP_CONTROL, XSLPsave

XSLPreinitialize, XPRSslpreinitialize

Purpose

Reset the SLP problem to match a just augmented system

Topic areas

SLP, Solution Process

Synopsis

```
int XPRS_CC XSLPreinitialize(XSLPprob prob);
```

Argument

prob The current SLP problem.

Further information

Can be used to rerun the SLP optimization process with updated parameters, penalties or initial values, but unchanged augmentation.

Related topics

XSLPcreateprob, XSLPdestroyprob, XSLPunconstruct, XSLPsetcurrentiv,

XSLPsave

Purpose

Save the Xpress NonLinear problem to file

Topic areas

File IO, Save Restore

Synopsis

```
int XPRS_CC XSLPsave(XSLPprob prob);
```

Argument

prob The current SLP problem.

Further information

Normally XSLPsave saves both the Xpress NonLinear problem and the underlying optimizer problem. If only the Xpress NonLinear problem is required, set the integer control variable XSLP_CONTROL appropriately.

Related topics

XSLP_CONTROL, XSLPrestore XSLPsaveas

XSLPsaveas

Purpose

Save the Xpress NonLinear problem to a named file

Topic areas

File IO, Save Restore

Synopsis

```
int XPRS_CC XSLPsaveas(XSLPprob prob, const char *filename);
```

Arguments

prob	The current SLP problem.
filename	The name of the file (without extension) in which the problem is to be saved.

Further information

Normally XSLPsaveas saves both the Xpress NonLinear problem and the underlying optimizer problem. If only the Xpress NonLinear problem is required, set the integer control variable XSLP_CONTROL appropriately.

Related topics

XSLP_CONTROL, XSLPrestore XSLPsave

XSLPscaling

Purpose

Analyze the current matrix for largest/smallest coefficients and ratios

Topic area

Numerics

Synopsis

```
int XPRS_CC XSLPscaling(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

The following example analyzes the matrix

```
XSLPscaling(prob);
```

Further information

The current matrix (including augmentation if it has been carried out) is scanned for the absolute and relative sizes of elements. The following information is reported:

- Largest and smallest elements in the matrix;
- Counts of the ranges of row ratios in powers of 10 (e.g. number of rows with ratio between 1.0E+01 and 1.0E+02);
- List of the rows (with largest and smallest elements) which appear in the highest range;
- Counts of the ranges of column ratios in powers of 10 (e.g. number of columns with ratio between 1.0E+01 and 1.0E+02);
- List of the columns (with largest and smallest elements) which appear in the highest range;
- Element ranges in powers of 10 (e.g. number of elements between 1.0E+01 and 1.0E+02).

Where any of the reported items (largest or smallest element in the matrix or any reported row or column element) is in a penalty error vector, the results are repeated, excluding all penalty error vectors.

XSLPsetcbcascadeend, XPRSsetcbslpcascadeend

Purpose

Set a user callback to be called at the end of the cascading process, after the last variable has been cascaded

Topic areas

Callback, SLP, Cascading

Synopsis

```
int XPRS_CC XSLPsetcbcascadeend(XSLPprob prob, int (XPRS_CC *cascadeend)
    (XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
cascadeend	The function to be called at the end of the cascading process. <code>cascadeend</code> returns an integer value. The return value is noted by Xpress-SLP but it has no effect on the optimization.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to <code>XSLPsetcbcascadeend</code> .
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to <code>cascadeend</code> as <code>cbdata</code> .

Example

The following example sets up a callback to be executed at the end of the cascading process which checks if any of the values have been changed significantly:

```
double *cSol;
XSLPsetcbcascadeend(prob, CBCascEnd, &cSol);
```

A suitable callback function might resemble this:

```
int XPRS_CC CBCascEnd(XSLPprob MyProb, void *Obj) {
    int iCol, nCol;
    double *cSol, Value;
    cSol = * (double **) Obj;
    XSLPgetintcontrol(MyProb, XPRS_COLS, &nCol);
    for (iCol=0; iCol<nCol; iCol++) {
        XSLPgetvar(MyProb, iCol, NULL, NULL, NULL,
            NULL, NULL, NULL, &Value,
            NULL, NULL, NULL, NULL,
            NULL, NULL, NULL, NULL);
        if (fabs(Value-cSol[iCol]) > .01)
            printf("\nCol %d changed from %lg to %lg",
                iCol, cSol[iCol], Value);
    }
    return 0;
}
```

The data argument is used here to hold the address of the array `cSol` which we assume has been populated with the original solution values.

Further information

This callback can be used at the end of the cascading, when all the solution values have been recalculated.

Related topics

XSLPcascade, XSLPsetcbcascadestart, XSLPsetcbcascadevar, XSLPsetcbcascadevarfail

XSLPsetcbcascadestart, XPRSsetcbslpcascadestart

Purpose

Set a user callback to be called at the start of the cascading process, before any variables have been cascaded

Topic areas

Callback, SLP, Cascading

Synopsis

```
int XPRS_CC XSLPsetcbcascadestart(XSLPprob prob, int (XPRS_CC
    *cascadestart) (XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

<code>prob</code>	The current SLP problem.
<code>cascadestart</code>	The function to be called at the start of the cascading process. <code>cascadestart</code> returns an integer value. If the return value is nonzero, the cascading process will be omitted for the current SLP iteration, but the optimization will continue.
<code>cbprob</code>	The problem passed to the callback function.
<code>cbdata</code>	The user-defined object passed as data to <code>XSLPsetcbcascadestart</code> .
<code>data</code>	Address of a user-defined object, which can be used for any purpose by the function. data is passed to <code>cascadestart</code> as <code>cbdata</code> .

Example

The following example sets up a callback to be executed at the start of the cascading process to save the current values of the variables:

```
double *cSol;
XSLPsetcbcascadestart(prob, CBCascStart, &cSol);
```

A suitable callback function might resemble this:

```
int XPRS_CC CBCascStart(XSLPprob MyProb, void *Obj) {
    int iCol, nCol;
    double *cSol;
    cSol = * (double **) Obj;
    XSLPgetintcontrol(MyProb, XPRS_COLS, &nCol);
    for (iCol=0; iCol<nCol; iCol++) {
        XSLPgetvar(MyProb, iCol, NULL, NULL, NULL,
            NULL, NULL, NULL, &cSol[iCol],
            NULL, NULL, NULL, NULL,
            NULL, NULL, NULL, NULL);
    }
    return 0;
}
```

The data argument is used here to hold the address of the array `cSol` which we populate with the solution values.

Further information

This callback can be used at the start of the cascading, before any of the solution values have been recalculated.

Related topics

`XSLPcascade`, `XSLPsetcbcascadeend`, `XSLPsetcbcascadevar`, `XSLPsetcbcascadevarfail`

XSLPsetcbcascadevar, XPRSsetcbslpcascadevar

Purpose

Set a user callback to be called after each column has been cascaded

Topic areas

Callback, SLP, Cascading

Synopsis

```
int XPRS_CC XSLPsetcbcascadevar(XSLPprob prob, int (XPRS_CC *cascadevar)
    (XSLPprob cbprob, void *cbdata, int col), void *data);
```

Arguments

prob	The current SLP problem.
cascadevar	The function to be called after each column has been cascaded. <code>cascadevar</code> returns an integer value. If the return value is nonzero, the cascading process will be omitted for the remaining variables during the current SLP iteration, but the optimization will continue.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to <code>XSLPsetcbcascadevar</code> .
col	The number of the column which has been cascaded.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to <code>cascadevar</code> as <code>cbdata</code> .

Example

The following example sets up a callback to be executed after each variable has been cascaded:

```
double *cSol;
XSLPsetcbcascadevar(prob, CBCascVar, &cSol);
```

The following sample callback function resets the value of the variable if the cascaded value is of the opposite sign to the original value:

```
int XPRS_CC CBCascVar(XSLPprob MyProb, void *Obj, int iCol) {
    double *cSol, Value;
    cSol = * (double **) Obj;
    XSLPgetvar(MyProb, iCol, NULL, NULL, NULL,
        NULL, NULL, NULL, &Value,
        NULL, NULL, NULL, NULL,
        NULL, NULL, NULL, NULL);
    if (Value * cSol[iCol] < 0) {
        Value = cSol[iCol];
        XSLPchgvar(MyProb, ColNum, NULL, NULL, NULL, NULL,
            NULL, NULL, &Value, NULL, NULL, NULL,
            NULL);
    }
    return 0;
}
```

The data argument is used here to hold the address of the array `cSol` which we assume has been populated with the original solution values.

Further information

This callback can be used after each variable has been cascaded and its new value has been calculated.

Related topics

XSLPcascade, XSLPsetcbcascadeend, XSLPsetcbcascadestart, XSLPsetcbcascadevarfail

XSLPsetcbcascadevarfail, XPRSsetcbslpcascadevarfail

Purpose

Set a user callback to be called after cascading a column was not successful

Topic areas

Callback, SLP, Cascading

Synopsis

```
int XPRS_CC XSLPsetcbcascadevarfail(XSLPprob prob, int (XPRS_CC
    *cascadevarfail) (XSLPprob cbprob, void *cbdata, int col), void
    *data);
```

Arguments

prob	The current SLP problem.
cascadevarfail	The function to be called after cascading a column was not successful. cascadevarfail returns an integer value. If the return value is nonzero, the cascading process will be omitted for the remaining variables during the current SLP iteration, but the optimization will continue.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbcascadevarfail.
col	The number of the column which has been cascaded.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to cascadevarfail as cbdata.

Further information

This callback can be used to provide user defined updates for SLP variables having a determining row that were not successfully cascaded due to the determining row being close to singular around the current values. This callback will always be called in place of the cascadevar callback in such cases, and in no situation will both the cascadevar and the cascadevarfail callback be called in the same iteration for the same variable.

Related topics

XSLPcascade, XSLPsetcbcascadeend, XSLPsetcbcascadestart, XSLPsetcbcascadevar

XSLPsetcbcofevalerror, XPRSsetcbnlpcfevalerror

Purpose

Set a user callback to be called when an evaluation of a coefficient fails during the solve

Topic areas

Callback, Numerics

Synopsis

```
int XPRS_CC XSLPsetcbcofevalerror(XSLPprob prob, int (XPRS_CC
    *cofevalerror) (XSLPprob cbprob, void *cbdata, int row, int col),
    void *data);
```

Arguments

prob	The current SLP problem.
cofevalerror	The function to be called when an evaluation fails.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbcofevalerror.
row	The row position of the coefficient.
col	The column position of the coefficient.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to cofevalerror as cbdata.

Further information

This callback can be used to capture when an evaluation of a coefficient fails. The callback is called only once for each coefficient.

Related topics

XSLPprintevalinfo

XSLPsetcbconstruct, XPRSsetcbslpconstruct

Purpose

Set a user callback to be called during the Xpress-SLP augmentation process

Topic areas

Callback, SLP

Synopsis

```
int XPRS_CC XSLPsetcbconstruct(XSLPprob prob, int (XPRS_CC *construct)
    (XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
construct	The function to be called during problem augmentation. <code>construct</code> returns an integer value. See below for an explanation of the values.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to <code>XSLPsetcbconstruct</code> .
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to <code>construct</code> as <code>cbdata</code> .

Example

The following example sets up a callback to be executed during the Xpress-SLP problem augmentation:

```
double *cValue;
cValue = NULL;
XSLPsetcbconstruct(prob, CBConstruct, &cValue);
```

The following sample callback function sets values for the variables the first time the function is called and returns to `XSLPconstruct` to recalculate the initial matrix. The second time it is called it frees the allocated memory and returns to `XSLPconstruct` to proceed with the rest of the augmentation.

```
int XPRS_CC CBConstruct(XSLPprob MyProb, void *Obj) {
    double *cValue;
    int i, n;
    /* if data is NULL, this is first-time entry */
    if (*(void**)Obj == NULL) {
        XSLPgetintattrib(MyProb, XPRS_COLS, &n);
        cValue = malloc(n*sizeof(double));
        /* ... initialize with values (not shown here) and then ... */
        for (i=0;i<n;i++)
            /* store into SLP structures */
            XSLPchgvar(MyProb, i, NULL, NULL, NULL, NULL,
                NULL, NULL, &cValue[i], NULL, NULL, NULL,
                NULL);
        /* set data non-null to indicate we have processed data */
        *(void**)Obj = cValue;
        return -1;
    }
    else {
        /* free memory, clear marker and continue */
        free(*(void**)Obj);
        *(void**)Obj = NULL;
    }
    return 0;
}
```

Further information

This callback can be used during the problem augmentation, generally (although not exclusively) to change the initial values for the variables.

The following return codes are accepted:

0	Normal return: augmentation continues
-1	Return to recalculate matrix values
-2	Return to recalculate row weights and matrix entries
other	Error return: augmentation terminates, <code>XSLPconstruct</code> terminates with a nonzero error code.

The return values -1 and -2 will cause the callback to be called a second time after the matrix has been recalculated. It is the responsibility of the callback to ensure that it does ultimately exit with a return value of zero.

Related topics

`XSLPconstruct`

XSLPsetcbdestroy

Purpose

Set a user callback to be called when an SLP problem is about to be destroyed

Topic area

Callback

Synopsis

```
int XPRS_CC XSLPsetcbdestroy(XSLPprob prob, int (XPRS_CC *destroy)
                             (XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
destroy	The function to be called when the SLP problem is about to be destroyed. destroy returns an integer value. At present the return value is ignored.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbdestroy.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to destroy as cbdata.

Example

The following example sets up a callback to be executed before the SLP problem is destroyed:

```
double *cSol;
XSLPsetcbdestroy(prob, CBDestroy, &cSol);
```

The following sample callback function frees the memory associated with the user-defined object:

```
int XPRS_CC CBDestroy(XSLPprob MyProb, void *Obj) {
    if (*(void**)Obj) free(*(void**)Obj);
    return 0;
}
```

The data argument is used here to hold the address of the array `cSol` which we assume was assigned using one of the `malloc` functions.

Further information

This callback can be used when the problem is about to be destroyed to free any user-defined resources which were allocated during the life of the problem.

Related topics

XSLPdestroyprob

XSLPsetcbdrcol, XPRSsetcbslpdrcol

Purpose

Set a user callback used to override the update of variables with small determining column

Topic areas

Callback, SLP, Cascading

Synopsis

```
int XPRS_CC XSLPsetcbdrcol(XSLPprob prob, int (XPRS_CC *drcol) (XSLPprob
    prob, void *data, int col, int detcol, double detval, double *
    p_value, double lb, double ub), void *data);
```

Arguments

prob	The current SLP problem.
drcol	The function to be called after each column has been cascaded. drcol returns an integer value. If the return value is positive, it will indicate that the value has been fixed, and cascading should be omitted for the variable. A negative value indicates that a previously fixed value has been relaxed. If no action is taken, a 0 return value should be used.
prob	The problem passed to the callback function.
data	The user-defined object passed as data to XSLPsetcbcascadevar.
col	The index of the column for which the determining columns is checked.
detcol	The index of the determining column for the column that is being updated.
detval	The value of the determining column in the current SLP iteration.
p_value	Used to return the new value for column col, should it need to be updated, in which case the callback must return a positive value to indicate that this value should be used.
lb	The original lower bound of column col. The callback provides this value as a reference, should the bound be updated or changed during the solution process.
ub	The original upper bound of column col. The callback provides this value as a reference, should the bound be updated or changed during the solution process.
data	Address of a user-defined object, which can be used for any purpose. by the function. data is passed to drcol as data.

Further information

If set, this callback is called as part of the cascading procedure. Please see Chapter *Cascading* for more information.

Related topics

XSLP_DRCOLTOL, XSLPcascade, XSLPsetcbcascadeend, XSLPsetcbcascadestart

XSLPsetcbintsol, XPRSsetcbintsol

Purpose

Set a user callback to be called during MISLP when an integer solution is obtained

Topic areas

Callback, MISLP

Synopsis

```
int XPRS_CC XSLPsetcbintsol(XSLPprob prob, int (XPRS_CC *intsol) (XSLPprob
    cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
intsol	The function to be called when an integer solution is obtained. <code>intsol</code> returns an integer value. At present, the return value is ignored.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to <code>XSLPsetcbintsol</code> .
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to <code>intsol</code> as <code>cbdata</code> .

Example

The following example sets up a callback to be executed whenever an integer solution is found during MISLP:

```
double *cSol;
int nInputCol;
XPRSgetintattrib(xprob, XPRS_INPUTCOLS, &nInputCol);
cSol = (double*)malloc(nInputCol * sizeof(double));
XSLPsetcbintsol(prob, CBIntSol, &cSol);
```

The following sample callback function saves the solution values for the integer solution just found:

```
int XPRS_CC CBIntSol(XSLPprob MyProb, void *Obj) {
    XPRSprob xprob;
    int nInputCol;
    int solStatus;
    double *cSol;

    cSol = * (double **) Obj;
    XSLPgetptrattrib(MyProb, XSLP_XPRSPROBLEM, &xprob);
    XPRSgetintattrib(xprob, XPRS_INPUTCOLS, &nInputCol);
    XPRSgetsolution(xprob, &solStatus, cSol, 0, nInputCol-1);
    return 0;
}
```

The data argument is used here to hold the address of the array `cSol`.

Further information

This callback must be used during MISLP instead of the `XPRSsetcbintsol` callback which is used for MIP problems.

Further information

Note that for SLP-MIP-SLP (see `XSLP_MIPALGORITHM`), the `intsol` callback will only be fired once after the second SLP solve terminated, in which case the solution can also only be queried via `XPRSgetsolution` or `XSLPgetslpsol` but not via `XPRSgetslpsol` (unless `postsolve` is disabled).

Related topics

`XSLPsetcboptnode`, `XSLPsetcbprenode`

XSLPsetcbiterend, XPRSsetcbslpiterend

Purpose

Set a user callback to be called at the end of each SLP iteration

Topic areas

Callback, SLP

Synopsis

```
int XPRS_CC XSLPsetcbiterend(XSLPprob prob, int (XPRS_CC *iterend)
    (XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
iterend	The function to be called at the end of each SLP iteration. iterend returns an integer value. If the return value is nonzero, the SLP iterations will stop.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbiterend.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to iterend as cbdata.

Example

The following example sets up a callback to be executed at the end of each SLP iteration. It records the number of LP iterations in the latest optimization and stops if there were fewer than 10:

```
XSLPsetcbiterend(prob, CBIterEnd, NULL);
```

A suitable callback function might resemble this:

```
int XPRS_CC CBIterEnd(XSLPprob MyProb, void *Obj) {
    int nIter;
    XPRSprob xprob;
    XSLPgetptrattrib(MyProb, XSLP_XPRSPROBLEM, &xprob);
    XSLPgetintattrib(xprob, XPRS_SIMPLEXITER, &nIter);
    if (nIter < 10) return 1;
    return 0;
}
```

The data argument is not used here, and so is passed as NULL.

Further information

This callback can be used at the end of each SLP iteration to carry out any further processing and/or stop any further SLP iterations.

Related topics

XSLPsetcbiterstart, XSLPsetcbitervar

XSLPsetcbiterstart, XPRSsetcbslpiterstart

Purpose

Set a user callback to be called at the start of each SLP iteration

Topic areas

Callback, SLP

Synopsis

```
int XPRS_CC XSLPsetcbiterstart(XSLPprob prob, int (XPRS_CC *iterstart)
                               (XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
iterstart	The function to be called at the start of each SLP iteration. iterstart returns an integer value. If the return value is nonzero, the SLP iterations will stop.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbiterstart.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to iterstart as cbdata.

Example

The following example sets up a callback to be executed at the start of the optimization to save the values of the variables from the previous iteration:

```
double *cSol;
XSLPsetcbiterstart(prob, CBIterStart, &cSol);
```

A suitable callback function might resemble this:

```
int XPRS_CC CBIterStart(XSLPprob MyProb, void *Obj) {
    XPRSprob xprob;
    double *cSol;
    int nIter;
    cSol = * (double **) Obj;
    XSLPgetintattrib(MyProb, XSLP_ITER, &nIter);
    if (nIter == 0) return 0; /* no previous solution */
    XSLPgetptrattrib(MyProb, XSLP_XPRSPROBLEM, &xprob);
    XPRSgetlp_sol(xprob, cSol, NULL, NULL, NULL);
    return 0;
}
```

The data argument is used here to hold the address of the array cSol which we populate with the solution values.

Further information

This callback can be used at the start of each SLP iteration before the optimization begins.

Related topics

XSLPsetcbiterend, XSLPsetcbitervar

XSLPsetcbitervar, XPRSsetcbslpitervar

Purpose

Set a user callback to be called after each column has been tested for convergence

Topic areas

Callback, SLP, SLP-convergence

Synopsis

```
int XPRS_CC XSLPsetcbitervar(XSLPprob prob, int (XPRS_CC *itervar)
                             (XSLPprob cbprob, void *cbdata, int col), void *data);
```

Arguments

prob	The current SLP problem.
itervar	The function to be called after each column has been tested for convergence. <code>itervar</code> returns an integer value. The return value is interpreted as a convergence status. The possible values are: <ul style="list-style-type: none"> < 0 The variable has not converged; 0 Keep the internal convergence status; 1 to 10 The column has converged on a system-defined convergence criterion (these values should not normally be returned); > 10 The variable has converged on user criteria.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to <code>XSLPsetcbitervar</code> .
col	The number of the column which has been tested for convergence.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to <code>itervar</code> as <code>cbdata</code> .

Example

The following example sets up a callback to be executed after each variable has been tested for convergence. The user object `Important` is an integer array which has already been set up and holds a flag for each variable indicating whether it is important that it converges.

```
int *Important;
XSLPsetcbitervar(prob, CBIterVar, &Important);
```

The following sample callback function tests if the variable is already converged. If not, then it checks if the variable is important. If it is not important, the function returns a convergence status of 99.

```
int XPRS_CC CBIterVar(XSLPprob MyProb, void *Obj, int iCol) {
    int *Important, Converged;
    Important = *(int **) Obj;
    XSLPgetvar(MyProb, iCol, NULL, NULL, NULL,
               NULL, NULL, NULL, NULL,
               NULL, NULL, &Converged, NULL,
               NULL, NULL, NULL, NULL);
    if (Converged) return 0;
    if (!Important[iCol]) return 99;
    return -1;
}
```

The data argument is used here to hold the address of the array `Important`.

Further information

This callback can be used after each variable has been checked for convergence, and allows the convergence status to be reset if required.

Related topics

`XSLPsetcbiterend`, `XSLPsetcbiterstart`

XSLPsetcbmessage

Purpose

Set a user callback to be called whenever Xpress NonLinear outputs a line of text according to XSLP_ECHOXPRSMESSAGES.

Topic areas

Callback, Logging

Synopsis

```
int XPRS_CC XSLPsetcbmessage(XSLPprob prob, void (XPRS_CC *message)
                             (XSLPprob cbprob, void *cbdata, char *msg, int msglen, int msgtype),
                             void *data);
```

Arguments

prob	The current SLP problem.						
message	The function to be called whenever Xpress NonLinear outputs a line of text. message does not return a value.						
cbprob	The problem passed to the callback function.						
cbdata	The user-defined object passed as data to XSLPsetcbmessage.						
msg	Character buffer holding the string to be output.						
msglen	Length in characters of msg excluding the null terminator.						
msgtype	Type of message. The following are system-defined: <table data-bbox="451 919 844 1018"> <tr><td>1</td><td>Information message</td></tr> <tr><td>3</td><td>Warning message</td></tr> <tr><td>4</td><td>Error message</td></tr> </table> A negative value indicates that the Optimizer is about to finish and any buffers should be flushed at this time.	1	Information message	3	Warning message	4	Error message
1	Information message						
3	Warning message						
4	Error message						
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to message as cbdata.						

Example

The following example creates a log file into which all messages are placed. System messages are also printed on standard output:

```
FILE *logfile;
logfile = fopen("myLog", "w");
XSLPsetcbmessage(prob, CBMessage, logfile);
```

A suitable callback function could resemble the following:

```
void XPRS_CC CBMessage(XSLPprob prob, void *Obj,
                      char *msg, int msglen, int msgtype) {

    FILE *logfile;
    logfile = (FILE *) Obj;
    if (msgtype < 0) {
        fflush(stdout);
        if (logfile) fflush(logfile);
        return;
    }
    switch (msgtype) {
        case 1: /* information */
        case 3: /* warning */
        case 4: /* error */
            printf("%s\n", msg);
        default: /* user */
            if (logfile)
                fprintf(logfile, "%s\n", msg);
    }
```

```
        break;  
    }  
    return;  
}
```

Further information

If a user message callback is defined then screen output is automatically disabled.

Output can be directed into a log file by using `XSLPsetlogfile`.

Related topics

`XSLPsetlogfile`

XSLPsetcbmsjobend, XPRSsetcbmsjobend

Purpose

Set a user callback to be called every time a new multistart job finishes. Can be used to overwrite the default solution ranking function

Topic areas

Callback, Multistart

Synopsis

```
int XSLP_CC XSLPsetcbmsjobend(XSLPprob prob, int (XSLP_CC
    *msjobend)(XSLPprob cbprob, void *cbdata,void *jobdata,const char
    *jobdesc,int *p_status), void *data);
```

Arguments

prob	The current SLP problem.
msjobend	The function to be called when a new multistart job is created
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbmsjobend.
jobdata	Job specific user-defined object, as specified by the multistart job creating API functions.
jobdesc	The description of the problem as specified by the multistart job creating API functions.
p_status	User return status variable: 0 - use the default evaluation of the finished job 1 - disregard the result and continue 2 - stop the multistart search
data	User data passed to the callback function.

Further information

The multistart pool is dynamic, and this callback can be used to load new multistart jobs using the normal API functions.

Related topics

XSLPsetcbmsjobstart, XSLPsetcbmswinner

XSLPsetcbmsjobstart, XPRSsetcbmsjobstart

Purpose

Set a user callback to be called every time a new multistart job is created, and the pre-loaded settings are applied

Topic areas

Callback, Multistart

Synopsis

```
int XSLP_CC XSLPsetcbmsjobstart(XSLPprob prob, int (XSLP_CC
    *msjobstart)(XSLPprob cbprob, void *cbdata,void *jobdata,const char
    *jobdesc,int *p_status), void *data);
```

Arguments

prob	The current SLP problem.
msjobstart	The function to be called when a new multistart job is created
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbmsjobstart.
jobdata	Job specific user-defined object, as specified by the multistart job creating API functions.
jobdesc	The description of the problem as specified by the multistart job creating API functions.
p_status	User return status variable: 0 - normal return, solve the job, 1 - disregard this job and continue, 2 - Stop multistart.
data	User data passed to the callback function.

Further information

All multistart jobs operation on an independent copy of the original problem, and any modification to the problem is allowed, including structural changes. Please note however, that any modification will be carried over to the base problem, should a modified problem be declared the winner prob.

Related topics

XSLPsetcbmsjobend, XSLPsetcbmswinner

XSLPsetcbmswinner, XPRSsetcbmswinner

Purpose

Set a user callback to be called every time a multistart winner has been declared

Topic areas

Callback, Multistart

Synopsis

```
int XSLP_CC XSLPsetcbmswinner(XSLPprob prob, int (XSLP_CC
    *mswinner)(XSLPprob cbprob, void *cbdata, void *jobdata, const char
    *jobdesc), void *data);
```

Arguments

prob	The current SLP problem.
mswinner	The function to be called when a multistart winner is declared. The return code is used to indicate whether callbacks should be copied over from the winner ($\neq 0$) or not (0).
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbmswinner.
jobdata	Job specific user-defined object, as specified by the multistart job creating API functions.
jobdesc	The description of the problem as specified by the multistart job creating API functions.
data	User data passed to the callback function.

Related topics

XSLPsetcbmsjobstart, XSLPsetcbmsjobend

XSLPsetcboptnode

Purpose

Set a user callback to be called during MISLP when an optimal SLP solution is obtained at a node

Topic areas

Callback, MISLP

Synopsis

```
int XPRS_CC XSLPsetcboptnode(XSLPprob prob, int (XPRS_CC *optnode)
                             (XSLPprob cbprob, void *cbdata, int *p_infeasible), void *data);
```

Arguments

prob	The current SLP problem.
optnode	The function to be called when an optimal SLP solution is obtained at a node. optnode returns an integer value. If the return value is nonzero, or if the feasibility flag is set nonzero, then further processing of the node will be terminated (it is declared infeasible).
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcboptnode.
p_infeasible	Address of an integer containing the feasibility flag. If optnode sets the flag nonzero, the node is declared infeasible.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to optnode as cbdata.

Example

The following example defines a callback function to be executed at each node when an SLP optimal solution is found. If there are significant penalty errors in the solution, the node is declared infeasible.

```
XSLPsetcboptnode(prob, CBOptNode, NULL);
```

A suitable callback function might resemble the following:

```
int XPRS_CC CBOptNode(XSLPprob cbprob, void *Obj, int *p_infeasible) {
    double Total, ObjVal;
    XSLPgetdblattrib(cbprob, XSLP_ERRORCOSTS, &Total);
    XSLPgetdblattrib(cbprob, XSLP_OBJVAL, &ObjVal);
    if (fabs(Total) > fabs(ObjVal) * 0.001 &&
        fabs(Total) > 1) *p_infeasible = 1;
    return 0;
}
```

Further information

If a node is declared infeasible from the callback function, the cost of exploring the node further will be avoided.

This callback must be used in place of XPRSsetcboptnode when optimizing with MISLP.

Related topics

XSLPsetcbprenode, XSLPsetcbslpnode

XSLPsetcbprenode

Purpose

Set a user callback to be called during MISLP after the set-up of the SLP problem to be solved at a node, but before SLP optimization

Topic areas

Callback, MISLP

Synopsis

```
int XPRS_CC XSLPsetcbprenode(XSLPprob prob, int (XPRS_CC *prenode)
                             (XSLPprob cbprob, void *cbdata, int *p_infeasible), void *data);
```

Arguments

prob	The current SLP problem.
prenode	The function to be called after the set-up of the SLP problem to be solved at a node. prenode returns an integer value. If the return value is nonzero, or if the feasibility flag is set nonzero, then further processing of the node will be terminated (it is declared infeasible).
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbprenode.
p_infeasible	Address of an integer containing the feasibility flag. If prenode sets the flag nonzero, the node is declared infeasible.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to prenode as cbdata.

Example

The following example sets up a callback function to be executed at each node before the SLP optimization starts. The array IntList contains a list of integer variables, and the function prints the bounds on these variables.

```
int *IntList;
XSLPsetcbprenode(prob, CBPreNode, IntList);
```

A suitable callback function might resemble the following:

```
int XPRS_CC CBPreNode(XSLPprob cbprob, void *Obj, int *p_infeasible) {
    XPRSprob xprob;
    int i, *IntList;
    double LO, UP;
    IntList = (int *) Obj;
    XSLPgetptrattrib(cbprob, XSLP_XPRSPROBLEM, &xprob);
    for (i=0; IntList[i]>=0; i++) {
        XPRSgetlb(xprob,&LO,IntList[i],IntList[i]);
        XPRSgetub(xprob,&UP,IntList[i],IntList[i]);
        if (LO > 0 || UP < XPRS_PLUSINFINITY)
            printf("\nCol %d: %lg <= %lg",LO,UP);
    }
    return 0;
}
```

Further information

If a node can be identified as infeasible by the callback function, then the initial optimization at the current node is avoided, as well as further exploration of the node.

This callback must be used in place of XPRSsetcbprenode when optimizing with MISLP.

Related topics

XSLPsetcboptnode, XSLPsetcbslpnode

XSLPsetcbpreupdatelinearization, XPRSsetcbslppreupdatelinearization

Purpose

Set a user callback to be called before the linearization is updated

Topic areas

Callback, SLP

Synopsis

```
int XPRS_CC XSLPsetcbpreupdatelinearization(XSLPprob prob, int (XPRS_CC
    *preupdatelinearization) (XSLPprob cbprob, void *cbdata, int *when),
    void *data);
```

Arguments

prob	The current SLP problem.
preupdatelinearization	The function to be called before the linearization is updated. preupdatelinearization returns an integer value. If the return value is nonzero, the optimization will return an error code and the "User Return Code" error will be set.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbpreupdatelinearization.
when	Indicates the call number, starting at 1. If returned nonzero, another call to the callback will be scheduled. If returned zero, a final call with *when == -1 will be made.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to preupdatelinearization as cbdata.

Further information

When the linearization is updated, all user functions are evaluated and their derivatives calculated at the current base point. In some models, it is cheaper to compute the derivatives for all user functions at the same time, thereby avoiding repeated calculations for each function. This callback is intended to be used in such cases.

During each SLP iteration, the callback is invoked repeatedly, with *when indicating the current call number (starting from 1), until the callback indicates that no further calls are needed, by setting *when = 0. Between each callback invocation, the solver evaluates the user functions without requesting derivatives. After the callback has set *when = 0, the user functions are evaluated one more time, this time requesting derivatives, and then finally the callback is called with *when == -1, marking the end of the linearization update. The only time derivatives will be requested outside of this sequence is during KKT validation. This can be disabled during the solve by clearing the XSLP_CONVERGEBIT_VALIDATION_K bit in XSLP_CONVERGENCEOPS, ensuring that derivatives can always be precomputed.

One way that this callback can be used to precompute derivatives for user functions is as follows:

1. On each SLP iteration, the callback is first called with *when = 1. This is a signal that derivatives will be needed soon. The callback sets a flag to indicate that user functions should capture their input values.
2. When the callback returns, the user functions are evaluated without requesting derivatives. Each user function captures its input values somewhere, and returns the correct function value.
3. The callback is called again, with *when == 2. The callback now computes derivatives for all user functions using the captured input values. The callback clears the flag so that user functions no longer capture their input values, and sets *when = 0 to indicate that no further calls are needed.
4. When the callback returns, the user functions are evaluated again. Derivatives are requested, and the user functions return the precomputed derivative values.
5. The callback is invoked one more time for this iteration with *when == -1, marking the end of the linearization update. User functions should behave normally from this point.

Related topics

User functions, `XSLPdeluserfunction`, `XSLPimportlibfunc`

XSLPsetcbpresolved

Purpose

Set a user callback to be called after the nonlinear presolver has been applied.

Topic areas

Callback, SLP

Synopsis

```
int XSLP_CC XSLPsetcbpresolved(XSLPprob prob, int (XSLP_CC  
    *presolved)(XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
presolved	The function to be invoked after the nlp presolver is completed.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object received by the callback.
data	The user-defined object passed as cbdata to presolved.

XSLPsetcbslpend

Purpose

Set a user callback to be called at the end of the SLP optimization

Topic areas

Callback, SLP

Synopsis

```
int XPRS_CC XSLPsetcbslpend(XSLPprob prob, int (XPRS_CC *slpend) (XSLPprob
    cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
slpend	The function to be called at the end of the SLP optimization. slpend returns an integer value. If the return value is nonzero, the optimization will return an error code and the "User Return Code" error will be set.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbslpend.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to slpend as cbdata.

Example

The following example sets up a callback to be executed at the end of the SLP optimization. It frees the memory allocated to the object created when the optimization began:

```
void *ObjData;
ObjData = NULL;
XSLPsetcbslpend(prob, CBSlpEnd, &ObjData);
```

A suitable callback function might resemble this:

```
int XPRS_CC CBSlpEnd(XSLPprob MyProb, void *Obj) {
    void *ObjData;
    ObjData = * (void **) Obj;
    if (ObjData) free(ObjData);
    * (void **) Obj = NULL;
    return 0;
}
```

Further information

This callback can be used at the end of the SLP optimization to carry out any further processing or housekeeping before the optimization function returns.

Related topics

XSLPsetcbslpstart

XSLPsetcbslpnode

Purpose

Set a user callback to be called during MISLP after the SLP optimization at each node.

Topic areas

Callback, MISLP

Synopsis

```
int XPRS_CC XSLPsetcbslpnode(XSLPprob prob, int (XPRS_CC *slpnode)
                             (XSLPprob cbprob, void *cbdata, int *p_infeasible), void *data);
```

Arguments

prob	The current SLP problem.
slpnode	The function to be called after the SLP optimization at a node. slpnode returns an integer value. If the return value is nonzero, or if the feasibility flag is set nonzero, then further processing of the node will be terminated (it is declared infeasible).
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to XSLPsetcbslpnode.
p_infeasible	Address of an integer containing the feasibility flag. If slpnode sets the flag nonzero, the node is declared infeasible.
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to slpnode as cbdata.

Example

The following example sets up a callback function to be executed at each node after the SLP optimization finishes. If the solution value is worse than a target value (referenced through the user object), the node is cut off (it is declared infeasible).

```
double OBJtarget;
XSLPsetcbslpnode(prob, CBSLPNode, &OBJtarget);
```

A suitable callback function might resemble the following:

```
int XPRS_CC CBSLPNode(XSLPprob cbprob, void *Obj, int *p_infeasible) {
    double TargetValue, LPValue;
    XSLPgetdblattrib(prob, XPRS_LPOBJVAL, &LPValue);
    TargetValue = * (double *) Obj;
    if (LPValue < TargetValue) *p_infeasible = 1;
    return 0;
}
```

Further information

If a node can be cut off by the callback function, then further exploration of the node is avoided.

Related topics

XSLPsetcboptnode, XSLPsetcbprenode

XSLPsetcbslpstart

Purpose

Set a user callback to be called at the start of the SLP optimization

Topic areas

Callback, SLP

Synopsis

```
int XPRS_CC XSLPsetcbslpstart(XSLPprob prob, int (XPRS_CC *slpstart)
                             (XSLPprob cbprob, void *cbdata), void *data);
```

Arguments

prob	The current SLP problem.
slpstart	The function to be called at the start of the SLP optimization. <code>slpstart</code> returns an integer value. If the return value is nonzero, the optimization will not be carried out.
cbprob	The problem passed to the callback function.
cbdata	The user-defined object passed as data to <code>XSLPsetcbslpstart</code> .
data	Address of a user-defined object, which can be used for any purpose by the function. data is passed to <code>slpstart</code> as <code>cbdata</code> .

Example

The following example sets up a callback to be executed at the start of the SLP optimization. It allocates memory to a user-defined object to be used during the optimization:

```
void *ObjData;
ObjData = NULL;
XSLPsetcbslpstart(prob, CBSlpStart, &ObjData);
```

A suitable callback function might resemble this:

```
int XPRS_CC CBSlpStart(XSLPprob MyProb, void *Obj) {
    void *ObjData;
    ObjData = * (void **) Obj;
    if (ObjData) free(ObjData);
    * (void **) Obj = malloc(99*sizeof(double));
    return 0;
}
```

Further information

This callback can be used at the start of the SLP optimization to carry out any housekeeping before the optimization actually starts. Note that a nonzero return code from the callback will terminate the optimization immediately.

Related topics

XSLPsetcbslpend

XSLPsetcurrentiv, XPRSnlpsetcurrentiv

Purpose

Transfer the current solution to initial values

Topic area

Data Input

Synopsis

```
int XPRS_CC XSLPsetcurrentiv(XSLPprob prob);
```

Argument

prob The current SLP problem.

Further information

Provides a way to set the current iterates solution as initial values, make changes to parameters or to the underlying nonlinear problem and then rerun the SLP optimization process.

Related topics

XSLPreinitialize, XSLPunconstruct

XSLPsetdblcontrol

Purpose

Set the value of a double precision problem control

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPsetdblcontrol(XSLPprob prob, int control, double value);
```

Arguments

prob	The current SLP problem.
control	control (SLP or optimizer) whose value is to be returned.
value	Double precision value to be set.

Example

The following example sets the value of the Xpress NonLinear control XSLP_CTOL and of the optimizer control XPRS_FEASTOL:

```
XSLPsetdblcontrol(prob, XSLP_CTOL, 0.001);
XSLPgetdblcontrol(prob, XPRS_FEASTOL, 0.005);
```

Further information

Both SLP and optimizer controls can be set using this function. If an optimizer control is set, the return value will be the same as that from [XPRSsetdblcontrol](#), which can similarly be used to set both Optimizer and SLP controls.

Related topics

XSLPgetdblcontrol, XSLPsetintcontrol, XSLPsetstrcontrol

XSLPsetdefaultcontrol

Purpose

Set the values of one SLP control to its default value

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPsetdefaultcontrol(XSLPprob prob, int control);
```

Arguments

prob	The current SLP problem.
control	The number of the control to be reset to its default.

Example

The following example reads a problem from file, sets the XSLP_LOG control, optimizes the problem (using the "s" flag to enforce a local solve) and then reads and optimizes another problem using the default setting.

```
XSLPreadprob(prob, "Matrix1", "");
XSLPsetintcontrol(prob, XSLP_LOG, 4);
XSLPnlptimize(prob, "s");
XSLPsetdefaultcontrol(prob, XSLP_LOG);
XSLPreadprob(prob, "Matrix2", "");
XSLPnlptimize(prob, "s");
```

Further information

This also resets optimizer controls.

Related topics

XSLPsetdblcontrol, XSLPsetdefaults, XSLPsetintcontrol, XSLPsetstrcontrol, XSLP_CONTROL

XSLPsetdefaults

Purpose

Set the values of all SLP controls to their default values

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPsetdefaults(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

The following example reads a problem from file, sets some controls, optimizes the problem (using the "s" flag to enforce a local solve) and then reads and optimizes another problem using the default settings.

```
XSLPreadprob(prob, "Matrix1", "");
XSLPsetintcontrol(prob, XSLP_LOG, 4);
XSLPsetdblcontrol(prob, XSLP_CTOL, 0.001);
XSLPsetdblcontrol(prob, XSLP_ATOL_A, 0.005);
XSLPnlpoptimize(prob, "s");
XSLPsetdefaults(prob);
XSLPreadprob(prob, "Matrix2", "");
XSLPnlpoptimize(prob, "s");
```

Further information

1. This function also resets all Knitro parameters to their default values.
2. This also resets optimizer controls unless bit 4 of XSLP_CONTROL is set.

Related topics

XSLPsetdblcontrol, XSLPsetintcontrol, XSLPsetstrcontrol

XSLPsetdetrow, XPRSslpsetdetrow

Purpose

Set the determining row of a variable

Topic areas

SLP, Cascading, Data Input

Synopsis

```
int XPRS_CC XSLPsetdetrow(XSLPprob prob, int nvars, const int[] colind,  
    const int[] rowind);
```

Arguments

prob	The current SLP problem.
nvars	The number of variables for which determining rows are set.
colind	Array of length <code>nvars</code> with the index of the column for which the determining row is set.
rowind	Array of length <code>nvars</code> with the index of the determining row.

Related topics

XSLPsetinitval

XSLPsetfunctionerror, XPRSnlpsetfunctionerror

Purpose

Set the function error flag for the problem

Topic area

Misc

Synopsis

```
int XPRS_CC XSLPsetfunctionerror(XSLPprob prob);
```

Argument

prob	The current SLP problem.
------	--------------------------

Further information

Once the function error has been set, calculations generally stop and the routines will return to their caller with a nonzero return code.

XSLPsetinitval, XPRSnlpsetinitval

Purpose

Set the initial value of a variable

Topic area

Data Input

Synopsis

```
int XPRS_CC XSLPsetinitval(XSLPprob prob, int nvars, const int[] colind,  
                           const double[] initial);
```

Arguments

<code>prob</code>	The current SLP problem.
<code>nvars</code>	Number of variables for which the initial value is to be set.
<code>colind</code>	Array of length <code>nvars</code> with index of the column for which the initial value is provided.
<code>initial</code>	Array of length <code>nvars</code> with the initial value.

Related topics

XSLPsetdetrow

XSLPsetintcontrol

Purpose

Set the value of an integer problem control

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPsetintcontrol(XSLPprob prob, int control, int value);
```

Arguments

prob	The current SLP problem.
control	control (SLP or optimizer) whose value is to be returned.
value	The value to be set.

Example

The following example sets the value of the Xpress NonLinear control XSLP_ALGORITHM and of the optimizer control XPRS_DEFAULTALG:

```
XSLPsetintcontrol(prob, XSLP_ALGORITHM, 934);  
XSLPsetintcontrol(prob, XPRS_DEFAULTALG, 3);
```

Further information

Both SLP and optimizer controls can be set using this function. If an optimizer control is set, the return value will be the same as that from [XPRSsetintcontrol](#), which can similarly be used to set both Optimizer and SLP controls.

Related topics

XSLPgetintcontrol, XSLPsetdblcontrol, XSLPsetintcontrol, XSLPsetstrcontrol

XSLPsetlogfile

Purpose

Define an output file to be used to receive messages from Xpress NonLinear

Topic areas

Logging, File IO

Synopsis

```
int XPRS_CC XSLPsetlogfile(XSLPprob prob, char *Filename, int option);
```

Arguments

prob	The current SLP problem.
filename	Character string containing the name of the file to be used for output.
option	option to indicate whether the output is directed to the file only (option=0) or (in console mode) to the console as well (option=1).

Example

The following example defines a log file "MyLog1" and directs output to the file and to the console:

```
XSLPsetlogfile(prob, "MyLog1", 1);
```

Further information

If `Filename` is `NULL`, the current log file (if any) will be closed, and message handling will revert to the default mechanism.

Related topics

`XSLPsetcbmessage`

XSLPsetparam

Purpose

Set the value of a control parameter by name

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPsetparam(XSLPprob prob, const char *name, const char
                        *value);
```

Arguments

prob	The current SLP problem.
name	Name of the control or attribute whose value is to be returned.
value	Character buffer containing the value.

Example

The following example sets the value of XSLP_ALGORITHM:

```
XSLPprob prob;
int Algorithm;
char Buffer[32];
Algorithm = 934;
sprintf(Buffer, "%d", Algorithm);
XSLPsetparam(prob, "XSLP_ALGORITHM", Buffer);
```

Further information

This function can be used to set any Xpress NonLinear or Optimizer control. The value is always passed as a character string. It is the user's responsibility to create the character string in an appropriate format.

Related topics

XSLPsetdblcontrol, XSLPsetintcontrol, XSLPsetparam, XSLPsetstrcontrol

XSLPsetstrcontrol

Purpose

Set the value of a string problem control

Topic area

Controls and Attributes

Synopsis

```
int XPRS_CC XSLPsetstrcontrol(XSLPprob prob, int control, const char
                             *value);
```

Arguments

prob	The current SLP problem.
control	control (SLP or optimizer) whose value is to be returned.
value	Character buffer containing the value.

Further information

Both SLP and optimizer controls can be set using this function. If an optimizer control is set, the return value will be the same as that from [XPRSsetstrcontrol](#), which can similarly be used to set both Optimizer and SLP controls.

Related topics

[XSLPgetstrcontrol](#), [XSLPsetdblcontrol](#), [XSLPsetintcontrol](#), [XSLPsetstrcontrol](#)

XSLPunconstruct, XPRSslpunconstruct

Purpose

Removes the augmentation and returns the problem to its pre-linearization state

Topic areas

SLP, Solution Process

Synopsis

```
int XPRS_CC XSLPunconstruct(XSLPprob prob);
```

Argument

prob The current SLP problem.

Further information

Only limited changes are allowed to an augmented problem.

Related topics

XSLPconstruct

XSLPupdatelinearization, XPRSslpupdatelinearization

Purpose

Updates the current linearization

Topic areas

SLP, Solution Process

Synopsis

```
int XPRS_CC XSLPupdatelinearization(XSLPprob prob);
```

Argument

prob The current SLP problem.

Further information

Updates the augmented problem (the linearization) to match the current base point. The base point is the current SLP solution. The values of the SLP variables can be changed using `XSLPchgvar`.

The linearization must be present, and this function can only be called after the problem has been augmented by `XSLPconstruct`.

Related topics

`XSLPconstruct`

XSLPvalidate, XPRSnlpvalidate

Purpose

Validate the feasibility of constraints in a converged solution

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPvalidate(XSLPprob prob);
```

Argument

prob The current SLP problem.

Example

The following example sets the validation tolerance parameters, validates the converged solution and retrieves the validation indices.

```
double IndexA, IndexR;
XSLPsetdblcontrol(prob, XSLP_VALIDATIONTOL_A, 0.001);
XSLPsetdblcontrol(prob, XSLP_VALIDATIONTOL_R, 0.001);
XSLPvalidate(prob);
XSLPgetdblattrib(prob, XSLP_VALIDATIONINDEX_A, &IndexA);
XSLPgetdblattrib(prob, XSLP_VALIDATIONINDEX_R, &IndexR);
```

Further information

XSLPvalidate checks the feasibility of a converged solution against relative and absolute tolerances for each constraint. The left hand side and the right hand side of the constraint are calculated using the converged solution values. If the calculated values imply that the constraint is infeasible, then the difference (D) is tested against the absolute and relative validation tolerances.

If $D < \text{XSLP_VALIDATIONTOL_A}$

then the constraint is within the absolute validation tolerance. The total positive (TPos) and negative contributions (TNeg) to the left hand side are also calculated.

If $D < \text{MAX}(\text{ABS}(\text{TPos}), \text{ABS}(\text{TNeg})) * \text{XSLP_VALIDATIONTOL_R}$

then the constraint is within the relative validation tolerance. For each constraint which is outside both the absolute and relative validation tolerances, validation factors are calculated which are the factors by which the infeasibility exceeds the corresponding validation tolerance; the smallest factor is printed in the validation report.

The validation index XSLP_VALIDATIONINDEX_A is the largest absolute validation factor multiplied by the absolute validation tolerance; the validation index XSLP_VALIDATIONINDEX_R is the largest relative validation factor multiplied by the relative validation tolerance.

Related topics

XSLP_VALIDATIONINDEX_A, XSLP_VALIDATIONINDEX_R, XSLP_VALIDATIONTOL_A,
XSLP_VALIDATIONTOL_R

XSLPvalidatekkt, XPRSnlpvalidatekkt

Purpose

Validates the first order optimality conditions also known as the Karush-Kuhn-Tucker (KKT) conditions versus the current solution

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPvalidatekkt(XSLPprob prob, int mode, int respectbasis, int
    updatemult, double violtarget);
```

Arguments

prob	The current SLP problem.						
mode	The calculation mode can be: <table> <tr> <td>0</td><td>recalculate the reduced costs at the current solution using the current dual solution.</td></tr> <tr> <td>1</td><td>minimize the sum of KKT violations by adjusting the dual solution.</td></tr> <tr> <td>2</td><td>perform both.</td></tr> </table>	0	recalculate the reduced costs at the current solution using the current dual solution.	1	minimize the sum of KKT violations by adjusting the dual solution.	2	perform both.
0	recalculate the reduced costs at the current solution using the current dual solution.						
1	minimize the sum of KKT violations by adjusting the dual solution.						
2	perform both.						
respectbasis	The following ways are defined to assess if a constraint is active: <table> <tr> <td>0</td><td>evaluate the recalculated slack activity versus XSLP_ECFTOL_R.</td></tr> <tr> <td>1</td><td>use the basis status of the slack in the linearized problem if available.</td></tr> <tr> <td>2</td><td>use both.</td></tr> </table>	0	evaluate the recalculated slack activity versus XSLP_ECFTOL_R.	1	use the basis status of the slack in the linearized problem if available.	2	use both.
0	evaluate the recalculated slack activity versus XSLP_ECFTOL_R.						
1	use the basis status of the slack in the linearized problem if available.						
2	use both.						
updatemult	The calculated values can be: <table> <tr> <td>0</td><td>only used to calculate the XSLP_VALIDATIONINDEX_K measure.</td></tr> <tr> <td>1</td><td>used to update the current dual solution and reduced costs.</td></tr> </table>	0	only used to calculate the XSLP_VALIDATIONINDEX_K measure.	1	used to update the current dual solution and reduced costs.		
0	only used to calculate the XSLP_VALIDATIONINDEX_K measure.						
1	used to update the current dual solution and reduced costs.						
violtarget	When calculating the best KKT multipliers, it is possible to enforce an even distribution of reduced costs violations by enforcing a bound on them.						

Further information

The bounds enforced by violtarget are automatically relaxed if the desired accuracy cannot be achieved.

XSLPvalidateprob, XPRSnlpvalidateprob

Purpose

Validates the current problem formulation and statement

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPvalidateprob(XSLPprob prob, int *p_nerrors, int
                             *p_nwarnings);
```

Arguments

prob	The current SLP problem.
p_nerrors	Returns the number of errors found in the problem. Errors are expected to make the problem not solve.
p_nwarnings	Returns the number of potential issues found in the problem. The solver may be able to automatically recover during the solve.

Further information

This function is expected to be used in the development stage of a model.

XSLPvalidaterow, XPRSnlpvalidaterow

Purpose

Prints an extensive analysis on a given constraint of the SLP problem

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPvalidaterow(XSLPprob prob, int row);
```

Arguments

prob	The current SLP problem.
row	The index of the row to be analyzed

Further information

The analysis will include the readable format of the original constraint and the augmented constraint. For infeasible constraints, the absolute and relative infeasibility is calculated. Variables in the constraints are listed including their value in the solution of the last linearization, the internal value (e.g. cascaded), reduced cost, step bound and convergence status. Scaling analysis is also provided.

XSLPvalidatevector, XPRSnlpvalidatevector

Purpose

Validate the feasibility of constraints for a given solution

Topic area

Solution

Synopsis

```
int XPRS_CC XSLPvalidatevector(XSLPprob prob, const double[] solution,
    double *p_suminf, double *p_sumscaledinf, double *p_objval);
```

Arguments

<code>prob</code>	The current SLP problem.
<code>solution</code>	A vector of length <code>XPRS_COLS</code> containing the solution vector to be checked.
<code>p_suminf</code>	Pointer to double in which the sum of infeasibility will be returned. May be NULL if not required.
<code>p_sumscaledinf</code>	Pointer to double in which the sum of scaled (relative) infeasibility will be returned. May be NULL if not required.
<code>p_objval</code>	Pointer to double in which the net objective will be returned. May be NULL if not required.

Further information

`XSLPvalidatevector` works the same way as `XSLPvalidate`, and will update `XSLP_VALIDATIONINDEX_A` and `XSLP_VALIDATIONINDEX_R`.

Related topics

`XSLP_VALIDATIONINDEX_A`, `XSLP_VALIDATIONINDEX_R`, `XSLP_VALIDATIONTOL_A`, `XSLP_VALIDATIONTOL_R`

XSLPwriteprob

Purpose

Write the current problem to a file in extended MPS or text format

Topic areas

File IO, Problem Information

Synopsis

```
int XPRS_CC XSLPwriteprob(XSLPprob prob, char *filename, char *flags);
```

Arguments

prob	The current SLP problem.												
filename	Character string holding the name of the file to receive the output. The extension ".mat" will automatically be appended to the file name, except for LP format when ".lp" will be appended.												
flags	The following flags can be used: <table> <tr> <td>a</td><td>write the current approximation (linearized) matrix (the default is to write the non-linear matrix including formulae);</td></tr> <tr> <td>o</td><td>one coefficient per line (the default is up to two numbers or one formula per line);</td></tr> <tr> <td>l</td><td>write the matrix in LP format, similar to the LP format written by XPRSwriteprob, with more SLP specific information</td></tr> <tr> <td>s</td><td>"scrambled" names (the default is to use the names provided on input);</td></tr> <tr> <td>v</td><td>use the provided filename verbatim, without appending the .mps or .lp extension;</td></tr> <tr> <td>z</td><td>write a compressed file.</td></tr> </table>	a	write the current approximation (linearized) matrix (the default is to write the non-linear matrix including formulae);	o	one coefficient per line (the default is up to two numbers or one formula per line);	l	write the matrix in LP format, similar to the LP format written by XPRSwriteprob, with more SLP specific information	s	"scrambled" names (the default is to use the names provided on input);	v	use the provided filename verbatim, without appending the .mps or .lp extension;	z	write a compressed file.
a	write the current approximation (linearized) matrix (the default is to write the non-linear matrix including formulae);												
o	one coefficient per line (the default is up to two numbers or one formula per line);												
l	write the matrix in LP format, similar to the LP format written by XPRSwriteprob, with more SLP specific information												
s	"scrambled" names (the default is to use the names provided on input);												
v	use the provided filename verbatim, without appending the .mps or .lp extension;												
z	write a compressed file.												

Example

The following example reads a problem from file, augments it and writes the augmented (linearized) matrix in text form to file "output.lp":

```
XSLPreadprob(prob, "Matrix", "");
XSLPconstruct(prob);
XSLPwriteprob(prob, "output", "l");
```

Related topics

XSLPreadprob

XSLPwriteslxsol

Purpose

Write the current solution to an MPS like file format

Topic areas

File IO, Solution

Synopsis

```
int XPRS_CC XSLPwriteslxsol(XSLPprob prob, char *filename, char *flags);
```

Arguments

prob	The current SLP problem.
filename	Character string holding the name of the file to receive the output. The extension ".slx" will automatically be appended to the file name, unless an extension is already specified in the filename.
flags	The following flags can be used: <ul style="list-style-type: none">p use double precision numbers;v use the provided filename verbatim, without appending the .slx extension;z write a compressed file.

CHAPTER 22

Internal Functions

Xpress NonLinear provides a set of standard functions for use in formulae. Many are standard mathematical functions; there are a few which are intended for specialized applications.

The following is a list of all the Xpress NonLinear internal functions:

ABS	Absolute value	p. 434
ARCCOS	Arc cosine trigonometric function	p. 427
ARCSIN	Arc sine trigonometric function	p. 428
ARCTAN	Arc tangent trigonometric function	p. 429
COS	Cosine trigonometric function	p. 430
ERF	The error function	p. 435
ERFC	The complementary error function	p. 436
EXP	Exponential function (e raised to the power)	p. 437
LN	Natural logarithm	p. 438
LOG10	Logarithm to base 10	p. 439
MAX	Maximum value of two or more expressions	p. 440
MIN	Minimum value of two or more expressions	p. 441
PWL	The piecewise linear function	p. 442
SIGN	The sign function	p. 443
SIN	Sine trigonometric function	p. 431
SQRT	Square root	p. 444
TAN	Tangent trigonometric function	p. 432

22.1 Trigonometric functions

The trigonometric functions `SIN`, `COS` and `TAN` return the value corresponding to their argument in radians. `SIN` and `COS` are well-defined, continuous and differentiable for all values of their arguments; care must be exercised when using `TAN` because it is discontinuous.

The inverse trigonometric functions `ARCSIN` and `ARCCOS` are undefined for arguments outside the range -1 to +1 and special care is required to ensure that no attempt is made to evaluate them outside this range. Derivatives for the inverse trigonometric functions are always calculated numerically.

ARCCOS

Purpose

Arc cosine trigonometric function

Synopsis

ARCSIN (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

Return value

A value in the range 0 to $+\pi$.

Further information

value must be in the range -1 to +1. Values outside the range will return zero and produce an appropriate error message. If XSLP_STOPOUTOFRANGE is set then the function error flag will be set.

ARCSIN

Purpose

Arc sine trigonometric function

Synopsis

`ARCSIN (value)`

Argument

`value` One of the following: a constant; a variable; a formula evaluating to a single value

Return value

A value in the range $-\pi/2$ to $+\pi/2$.

Further information

`value` must be in the range -1 to +1. Values outside the range will return zero and produce an appropriate error message. If `XSLP_STOPOUTOFRANGE` is set then the function error flag will be set.

ARCTAN

Purpose

Arc tangent trigonometric function

Synopsis

`ARCTAN (value)`

Argument

`value` One of the following: a constant; a variable; a formula evaluating to a single value

Return value

A value in the range $-\pi/2$ to $+\pi/2$.

COS

Purpose

Cosine trigonometric function

Synopsis

COS (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

SIN

Purpose

Sine trigonometric function

Synopsis

`SIN (value)`

Argument

`value` One of the following: a constant; a variable; a formula evaluating to a single value

TAN

Purpose

Tangent trigonometric function

Synopsis

TAN (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

22.2 Other mathematical functions

Most of the mathematical functions are differentiable, although care should be taken in using analytic derivatives where the derivative is changing rapidly.

ABS

Purpose

Absolute value

Synopsis

ABS (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

Further information

ABS is not always differentiable and so alternative modeling approaches should be used where possible.

ERF

Purpose

The error function

Synopsis

ERF (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

ERFC

Purpose

The complementary error function

Synopsis

ERFC (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

EXP

Purpose

Exponential function (e raised to the power)

Synopsis

EXP (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

LN

Purpose

Natural logarithm

Synopsis

LN (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

Further information

value must be strictly positive (greater than 1.0E-300).

LOG10

Purpose

Logarithm to base 10

Synopsis

LOG10 (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

Further information

value must be strictly positive (greater than 1.0E-300).

MAX

Purpose

Maximum value of two or more expressions

Synopsis

`MAX(value1, value2)`

Argument

`value1, value2` Each argument is one of the following: a constant; a variable; a formula evaluating to a single value

Further information

MAX is not always differentiable and so alternative modeling approaches should be used where possible.

In `mmxnlp`, `fmax` is used to represent the max function.

MIN

Purpose

Minimum value of two or more expressions

Synopsis

`MIN(value1, value2)`

Argument

`value1, value2` Each argument is one of the following: a constant; a variable; a formula evaluating to a single value

Further information

MIN is not always differentiable and so alternative modeling approaches should be used where possible.

In `mmxnlp`, `fmin` is used to represent the min function.

PWL

Purpose

The piecewise linear function

Synopsis

```
PWL( variable, x1, y1, ..., xk, yk )
```

Arguments

`variable` is a single variable that describes where the pwl should be evaluated.

`x1, y2, ..., xk, yk` are the k breakpoints of the piecewise linear function. The pwl is extended to minus and plus infinity using the first and last 2 breakpoints respectively.

SIGN

Purpose

The sign function

Synopsis

SIGN (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

SQRT

Purpose

Square root

Synopsis

SQRT (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value

Further information

value must be non-negative.

CHAPTER 23

Error Messages

If the optimization procedure or some other library function encounters an error, then the procedure normally terminates with a nonzero return code and sets an error code. For most functions, the return code is 32 for an error; those functions which can return Optimizer return codes (such as the functions for accessing attributes and controls) will return the Optimizer code in such circumstances.

If an error message is produced, it will normally be output to the message handler; for console-based output, it will appear on the console. The error message and the error code can also be obtained using the function `XSLPgetlasterror`. This allows the user to retrieve the message number and/or the message text. The format is:

```
XSLPgetlasterror(Prob, &ErrorCode, &ErrorMessage);
```

The following is a list of the error codes and an explanation of the message. In the list, error numbers are prefixed by *E-* and warnings by *W-*. The printed messages are generally prefixed by *Xpress NonLinear error* and *Xpress NonLinear warning* respectively.

E-12001 *invalid parameter number num*

This message is produced by the functions which access SLP or Optimizer controls and attributes. The parameter numbers for SLP are given in the header file `xslp.h`. The parameter is of the wrong type for the function, or cannot be changed by the user.

E-12002 *internal hash error*

This is a non-recoverable program error. If this error is encountered, please contact your local Xpress support office.

E-12003 *XSLPprob problem pointer is NULL*

The problem pointer has not been initialized and contains a zero address. Initialize the problem using `XSLPcreateprob`.

E-12004 *XSLPprob is corrupted or is not a valid problem*

The problem pointer is not the address of a valid problem. The problem pointer has been corrupted, and no longer contains the correct address; or the problem has not been initialized correctly; or the problem has been corrupted in memory. Check that your program is using the correct pointer and is not overwriting part of the memory area.

E-12005 *memory manager error - allocation error*

This message normally means that the system has run out of memory when trying to allocate or reallocate arrays. Use `XSLPprintmemory` to obtain a list of the arrays and amounts of memory allocated by the system. Ensure that any memory allocated by user programs is freed at the appropriate time.

E-12006 *memory manager error - Array expansion size (num) ≤ 0*

This may be caused by incorrect setting of the `XSLP_EXTRA*` control parameters to negative numbers. Use `XSLPprintmemory` to obtain a list of the arrays and amounts of memory

allocated by the system for the specified array. If the problem persists, please contact your local Xpress support office.

E-12007 *memory manager error - object Obj size not defined*

This is a non-recoverable program error. If this error is encountered, please contact your local Xpress support office.

E-12008 *cannot open file name*

This message appears when Xpress NonLinear is required to open a file of any type and encounters an error while doing so. Check that the file name is spelt correctly (including the path, directory or folder) and that it is accessible (for example, not locked by another application).

E-12009 *cannot open problem file name*

This message is produced by XSLPreadprob if it cannot find name.mat, name.mps or name. Note that "lp" format files are not accepted for SLP input.

E-12010 *internal I/O error*

This error is produced by XSLPreadprob if it is unable to read or write intermediate files required for input.

E-12011 *XSLPreadprob unknown record type name*

This error is produced by XSLPreadprob if it encounters a record in the file which is not identifiable. It may be out of place (for example, a matrix entry in the *BOUNDS* section), or it may be a completely invalid record type.

E-12012 *XSLPreadprob invalid function argument type name*

This error is produced by XSLPreadprob if it encounters a user function definition with an argument type that is not one of NULL, DOUBLE, INTEGER, CHAR or VARIANT.

E-12013 *XSLPreadprob invalid function linkage type name*

This error is produced by XSLPreadprob if it encounters a user function with a linkage type that is not one of DLL, XLS, XLF, MOSEL or COM.

E-12014 *XSLPreadprob unrecognized function name*

This error is produced by XSLPreadprob if it encounters a function reference in a formula which is not a pre-defined internal function nor a defined user function. Check the formula and the function name, and define the function if required.

E-12015 *func: item num out of range*

This message is produced by the Xpress NonLinear function *func* which is referencing the SLP *item* (row, column variable, etc). The index provided is out of range (less than 1 unless zero is explicitly allowed, or greater than the current number of items of that type). Remember that most Xpress NonLinear items count from 1.

E-12016 *missing left bracket in formula*

This message is produced during parsing of formulae provided in character or unparsed internal format. A right bracket is not correctly paired with a corresponding left bracket. Check the formulae.

E-12017 *missing left operand in formula*

This message is produced during parsing of formulae provided in character or unparsed internal format. An operator which takes two operands is missing the left hand one (and so immediately follows another operator or a bracket). Check the formulae.

E-12018 *missing right operand in formula*

This message is produced during parsing of formulae provided in character or unparsed internal format. An operator is missing the right hand (following) operand (and so is immediately followed by another operator or a bracket). Check the formulae.

E-12019 *missing right bracket in formula*

This message is produced during parsing of formulae provided in character or unparsed internal format. A left bracket is not correctly paired with a corresponding right bracket. Check the formulae.

E-12020 *column #n is defined more than once as an SLP variable*

This message is produced by XSLPaddvars or XSLPloadvars if the same column appears more than once in the list, or has already been defined as an SLP variable. Although XSLPchgvar is less efficient, it can be used to set the properties of an SLP variable whether or not it has already been declared.

E-12022 *undefined tolerance type name*

This error is produced by XSLPreadprob if it encounters a tolerance which is not one of the 9 defined types (TC, TA, TM, TI, TS, RA, RM, RI, RS). Check the two-character code for the tolerance.

W-12023 *name has been given a tolerance but is not an SLP variable*

This error is produced by XSLPreadprob if it encounters a tolerance for a variable which is not an SLP variable (it is not in a coefficient, it does not have a non-constant coefficient and it has not been given an initial value). If the tolerance is required (that is, if the variable is to be monitored for convergence) then give it an initial value so that it becomes an SLP variable. Otherwise, the tolerance will be ignored.

W-12024 *name has been given SLP data of type ty but is not an SLP variable*

This error is produced by XSLPreadprob if it encounters *SLPDATA* for a variable which has not been defined as an SLP variable. Typically, this is because the variable would only appear in coefficients, and the relevant coefficients are missing. The data item will be ignored.

E-12025 *func has the same source and destination problems*

This message is produced by XSLPcopycallbacks, XSLPcopycontrols and XSLPcopyprob if the source and destination problems are the same. If they are the same, then there is no point in copying them.

E-12026 *invalid or corrupt SAVE file*

This message is produced by XSLPrestore if the SAVE file header is not valid, or if internal consistency checks fail. Check that the file exists and was created by XSLPsave.

E-12027 *SAVE file version is too old*

This message is produced by XSLPrestore if the SAVE file was produced by an earlier version of Xpress NonLinear. In general, it is not possible to restore a file except with the same version of the program as the one which SAVED it.

W-12028 *problem already has augmented SLP structure*

This message is produced by XSLPconstruct if it is called for a second time for the same problem. The problem can only be augmented once, which must be done after all the variables and coefficients have been loaded. XSLPconstruct is called automatically by XSLPmaxim and XSLPminim if it has not been called earlier.

E-12029 *zero divisor*

This message is produced by the formula evaluation routines if an attempt is made to divide by a value less than XSLP_ZERO. A value of +/-XSLP_INFINITY is returned as the result and the calculation continues.

E-12030 *negative number, fractional exponent - truncated to integer*

This message is produced by the formula evaluation routines if an attempt is made to raise a negative number to a non-integer exponent. The exponent is truncated to an integer value and the calculation continues.

- E-12031 *binary search failed***
This is a non-recoverable program error. If this error is encountered, please contact your local Xpress support office.
- E-12032 *wrong number (num) of arguments to function func***
This message is produced by the formula evaluation routines if a formula contains the wrong number of arguments for an internal function (for example, SIN(A, B)). Correct the formula.
- E-12033 *argument value out of range in function func***
This message is produced by the formula evaluation routines if an internal function is called with an argument outside the allowable range (for example, LOG of a negative number). The function will normally return zero as the result and, if XSLP_STOPOUTOFRANGE is set, will set the function error flag.
- W-12034 *terminated following user return code num***
This message is produced by XSLPmaxim and XSLPminim if a nonzero value is returned by the callback defined by XSLPsetcbiterend or XSLPsetcbslpend.
- E-12037 *failed to load library/file/program "name" containing function "func"***
This message is produced if a user function is defined to be in a file, but Xpress NonLinear cannot find the specified file. Check that the correct file name is specified (also check the search paths such as \$PATH and %path% if necessary).
This message may also be produced if the specified library exists but is dependent on another library which is missing.
- E-12038 *function "func" is not correctly defined or is not in the specified location***
This message is produced if a user function is defined to be in a file, but Xpress NonLinear cannot find it in the file. Check that the number and type of the arguments is correct, and that the (external) name of the user function matches the name by which it is known in the file.
- E-12084 *Xpress NonLinear has not been initialized***
An attempt has been made to use Xpress NonLinear functions without a previous call to XSLPinit. Only a very few functions can be called before initialization. Check the sequence of calls to ensure that XSLPinit is called first, and that it completed successfully. *This error message normally produces return code 279.*
- E-12085 *Xpress NonLinear has not been licensed for use here***
Either Xpress NonLinear is not licensed at all (although the Xpress Optimizer may be licensed), or the particular feature (such as MISLP) is not licensed. Check the license and contact the local Fair Isaac sales office if necessary. *This error message normally produces return code 352.*
- E-12105 *Xpress NonLinear error: I/O error on file***
The message is produced by XSLPsave or XSLPwriteprob if there is an I/O error when writing the output file (usually because there is insufficient space to write the file).
- E-12107 *Xpress NonLinear error: user function type name not supported on this platform***
This message is produced if a user function defined as being of type XLS, XLF or COM and is run on a non-Windows platform.
- E-12110 *Xpress NonLinear error: unidentified section in REVISE: name***
The file provided to XSLPrevis contains an unsupported MPS section.
- E-12111 *Xpress NonLinear error: unidentified row type in REVISE: name***
The file provided to XSLPrevis contains an unsupported row type.
- E-12112 *Xpress NonLinear error: unidentified row in REVISE: name***
The file provided to XSLPrevis contains a row name not found in the current problem.

- E-12113 *Xpress NonLinear error: unidentified bound in REVISE: name***
The file provided to XSLP`revise` contains an unsupported bound type.
- E-12114 *Xpress NonLinear error: unidentified column in REVISE: name***
The file provided to XSLP`revise` contains a column name not found in the current problem.
- E-12121 *Xpress NonLinear error: bad return code num from user function func***
This message is produced during evaluation of a complicated user function if it returns a value (-1) indicating that the system should estimate the result from a previous function call, but there has been no previous function call.
- E-12124 *Xpress NonLinear error: augmented problem not set up***
The message is produced by XSLP`validate` if an attempt is made to validate the problem without a preceding call to XSLP`construct`. In fact, unless a solution to the linearized problem is available, XSLP`validate` will not be able to give useful results.
- E-12125 *Xpress NonLinear error: user function func terminated with errors***
This message is produced during evaluation of a user function if it sets the function error flag (see XSLP`setfunctionerror`).
- W-12142 *Xpress NonLinear warning: invalid record: text***
This error is produced by XSLP`readprob` if it encounters a record in the file which is identifiable but invalid (for example, a *BOUNDS* record without a bound set name). The record is ignored.
- E-12147 *Xpress NonLinear error: incompatible arguments in user function func***
This message is produced if a user function is called without providing the arguments required by the function.
- E-12151 *Xpress NonLinear error: column is not defined for formula in row***
This error is produced when an mps file has an incomplete or incorrect record for a nonlinear formula.
- E-12156 *Xpress NonLinear error: invalid double value***
This error is produced when a string value that was expected to represent a double does not.
- E-12157 *Xpress NonLinear error: invalid integer value***
This error is produced when a string value that was expected to represent an integer does not.
- E-12150 *Xpress NonLinear error: problem contains undefined user functions***
This error is produced if an mps or lp file is read that contains user functions that are later not defined before optimization.
- E-12158 *Xpress NonLinear error: unknown parameter name name***
This message is produced if an attempt is made to set or retrieve a value for a control parameter or attribute given by name where the name is incorrect.
- E-12159 *Xpress NonLinear error: unknown parameter type name***
A parameter has an unexpected type and cannot be retrieved. This is an internal error, please contact FICO support.
- E-12159 *Xpress NonLinear error: parameter number is not writable***
This message is produced if an attempt is made to set a value for an attribute.
- E-12160 *Xpress NonLinear error: parameter num is not available***
This message is produced if an attempt is made to retrieve a value for a control or attribute which is not readable.
- E-12161 *Xpress NonLinear error: parameter num is not available***
The parameter corresponding to the provided ID is an internal, not readable parameter.

- E-12163 *Xpress NonLinear error: PWL functions can only take constant values following the determining variable: $PWL(x, base1, value1, \dots, basek, valuek)$***
The values following the variable in a piecewise linear function (PWL) describe the breakpoints of the PWL, and so can only be constants.
- E-12164 *Xpress NonLinear error: PWL functions must take an odd number of arguments: $PWL(x, base1, value1, \dots, basek, valuek)$***
The values following the variable in a piecewise linear function (PWL) describe the breakpoints of the PWL in pairs, making the total number of arguments odd.
- E-12165 *Xpress NonLinear error: PWL functions must take the following form: $PWL(x, base1, value1, \dots, basek, valuek)$***
The piecewise linear function (PWL) defined is likely missing any break points, or has unexpected tokens.
- E-12166 *Xpress NonLinear error: PWL functions's first argument must be a variable: $PWL(x, base1, value1, \dots, basek, valuek)$***
A piecewise linear function (PWL) argument list must start with the determining variable.
- E-12167 *Xpress NonLinear error: The base points of a PWL function must be monotonically increasing***
A piecewise linear function (PWL) must list its breakpoints first values in monotonically increasing order.
- E-12168 *Xpress NonLinear error: NLP with general constraints and piecewise linear not allowed with presolve off***
A model with general constraints and piecewise linear functions depends on presolve for reformulation.
- E-12169 *Xpress NonLinear error: error removing columns in unconstruct***
A call to unconstruct failed to remove columns added by construct. This is an unexpected error, please contact FICO support.
- E-12170 *Xpress NonLinear error: error removing rows in unconstruct***
A call to unconstruct failed to remove rows added by construct. This is an unexpected error, please contact FICO support.
- E-12171 *Xpress NonLinear error: compute server returned an error***
An error occurred while solving remotely using a compute server.
- E-12173 *Xpress NonLinear error: Multistart is not supported in compute mode***
Multistart is currently not supported in compute mode.
- E-12174 *Xpress NonLinear error: User functions with named arguments are no longer supported***
User functions with named input and output arguments are no longer supported in the C API.
- E-12192 *Xpress NonLinear error: no problem or solution read***
No problem or solution has been read. If a problem read fails, it is not valid to continue with any problem building or solving functions.
- E-12193 *Xpress NonLinear error: this version of SLP requires XPRS version num or newer***
Although not recommended, Xpress SLP can work with different xprs library versions. This error is issued when a tool old xprs library is found.
- E-12194 *Xpress NonLinear error: provided buffer is too short***
The provided buffer is too short. This error may occur if a formula is retrieved from Xpress, into a buffer that is not large enough.
- E-12195 *Xpress NonLinear error: type index value is invalid***
The index provided is not valid for the this type.

- E-12196 *Xpress NonLinear error: error in problem transformation***
An error occurred while the problem was attempted to be reformulated as part of the nonlinear presolver. Please contact FICO support.
- E-12197 *Xpress NonLinear error: request index invalid***
The requested information cannot be retrieved as it is not valid or not available.
- E-12198 *Xpress NonLinear error: error while cascading. Cannot evaluate coefficient at row rowname column.***
Evaluating an expression in cascading has returned an error. There is likely a user function in the expression returning an error.
- E-12199 *Xpress NonLinear error: nonlinear coefficient in neutral objective row 'rowname'. Please use an objective transfer row instead.***
A nonlinear objective function in SLP needs to be modelled using an objective transfer row.
- E-12200 *Xpress NonLinear error: problem is not augmented.***
The operation is only valid for augmented problems. Please call the construct method first, or solve using SLP.
- E-12201 *Xpress NonLinear error: an internal error has occurred.***
An internal error has occurred that is not expected to have been caused by incorrect input. Please contact FICO support.
- E-12202 *Xpress NonLinear error: attribute i cannot be changed.***
Attributes normally cannot be changed, as they are set up by the solver. There are a few exceptions to this rule, the requested attribute is not among the exceptions.
- E-12203 *Xpress NonLinear error: no problem or solution written.***
No problem or solution was written to disk due to an error processing the data.
- E-12204 *Xpress NonLinear error: Unexpected token name in formula.***
An unexpected token was encountered when parsing a string formula.
- E-12205 *Xpress NonLinear error: Duplicate user function names name are not allowed.***
A user function with this name was added twice. User function names have to be unique.
- E-12206 *Xpress NonLinear error: User function name name conflicts with an internal function name.***
The name of a user function conflicts with an internal function name (sin, cos, ...), it has to be different from all internal functions.
- E-12207 *Xpress NonLinear error: User function name name rejected: reason.***
A user function name was rejected due to invalid characters, for example whitespace, certain mathematical or control characters or due to starting with a digit.
- E-12208 *Xpress NonLinear error: cannot continue from current state. Consider disabling postsolve.***

The nonlinear solve could not be continued. A typical reason for this is that the problem was already postsolved.
- E-12209 *Xpress NonLinear error: bad flags flag.***
An invalid combination of flags was provided to a function.
- E-12210 *Xpress NonLinear error: feature is not supported in this release.***
This feature is not supported for global solves in the current release.

CHAPTER 24

Xpress Knitro Control Parameters

This chapter provides a full list of the controls accepted by Xpress for setting Knitro parameters. Knitro has a great number and variety of user option settings and although it tries to choose the best settings by default, often significant performance improvements can be realized by choosing some non-default option settings.

XKTR_PARAM_ALGORITHM, KNITRO_PARAM_ALGORITHM	Indicates which algorithm to use to solve the problem	
	[Knitro, Solution Process]	p. 461
XKTR_PARAM_BAR_DIRECTINTERVAL, KNITRO_PARAM_BAR_DIRECTINTERVAL	Controls the maximum number of consecutive conjugate gradient (CG) steps before Knitro will try to enforce that a step is taken using direct linear algebra.	
	[Knitro, Limits]	p. 461
XKTR_PARAM_BAR_FEASIBLE, KNITRO_PARAM_BAR_FEASIBLE	Specifies whether special emphasis is placed on getting and staying feasible in the interior-point algorithms.	
	[Knitro]	p. 461
XKTR_PARAM_BAR_FEASMODETOL, KNITRO_PARAM_BAR_FEASMODETOL	Specifies the tolerance in equation that determines whether Knitro will force subsequent iterates to remain feasible.	
	[Knitro, Tolerances]	p. 457
XKTR_PARAM_BAR_INITMU, KNITRO_PARAM_BAR_INITMU	Specifies the initial value for the barrier parameter : μ used with the barrier algorithms.	
	[Knitro]	p. 457
XKTR_PARAM_BAR_INITPT, KNITRO_PARAM_BAR_INITPT	Indicates whether an initial point strategy is used with barrier algorithms.	
	[Knitro]	p. 462
XKTR_PARAM_BAR_MAXBACKTRACK, KNITRO_PARAM_BAR_MAXBACKTRACK	Indicates the maximum allowable number of backtracks during the linesearch of the Interior/Direct algorithm before reverting to a CG step.	
	[Knitro, Limits]	p. 462
XKTR_PARAM_BAR_MAXCROSSIT, KNITRO_PARAM_BAR_MAXCROSSIT	Specifies the maximum number of crossover iterations before termination.	
	[Knitro, Limits]	p. 463
XKTR_PARAM_BAR_MAXREFACTOR, KNITRO_PARAM_BAR_MAXREFACTOR	Indicates the maximum number of refactorizations of the KKT system per iteration of the Interior/Direct algorithm before reverting to a CG step.	
	[Knitro, Limits]	p. 463

<code>XKTR_PARAM_BAR_MURULE</code> , <code>KNITRO_PARAM_BAR_MURULE</code>	Indicates which strategy to use for modifying the barrier parameter μ in the barrier algorithms.	
[Knitro]		p. 463
<code>XKTR_PARAM_BAR_PENCONS</code> , <code>KNITRO_PARAM_BAR_PENCONS</code>	Indicates whether a penalty approach is applied to the constraints.	
[Knitro]		p. 464
<code>XKTR_PARAM_BAR_PENRULE</code> , <code>KNITRO_PARAM_BAR_PENRULE</code>	Indicates which penalty parameter strategy to use for determining whether or not to accept a trial iterate.	
[Knitro]		p. 464
<code>XKTR_PARAM_BAR_SWITCHRULE</code> , <code>KNITRO_PARAM_BAR_SWITCHRULE</code>	Indicates whether or not the barrier algorithms will allow switching from an optimality phase to a pure feasibility phase.	
[Knitro]		p. 465
<code>XKTR_PARAM_DELTA</code> , <code>KNITRO_PARAM_DELTA</code>	Specifies the initial trust region radius scaling factor used to determine the initial trust region size.	
[Knitro]		p. 457
<code>XKTR_PARAM_FEASTOL</code> , <code>KNITRO_PARAM_FEASTOL</code>	Specifies the final relative stopping tolerance for the feasibility error.	
[Knitro, Tolerances]		p. 457
<code>XKTR_PARAM_FEASTOLABS</code> , <code>KNITRO_PARAM_FEASTOLABS</code>	Specifies the final absolute stopping tolerance for the feasibility error.	
[Knitro, Tolerances]		p. 458
<code>XKTR_PARAM_GRADOPT</code> , <code>KNITRO_PARAM_GRADOPT</code>	Specifies how to compute the gradients of the objective and constraint functions.	
[Derivatives, Knitro]		p. 465
<code>XKTR_PARAM_HESSOPT</code> , <code>KNITRO_PARAM_HESSOPT</code>	Specifies how to compute the (approximate) Hessian of the Lagrangian.	
[Derivatives, Knitro]		p. 465
<code>XKTR_PARAM_HONORBND</code> s, <code>KNITRO_PARAM_HONORBND</code> s	Indicates whether or not to enforce satisfaction of simple variable bounds throughout the optimization.	
[Knitro]		p. 466
<code>XKTR_PARAM_INFEASTOL</code> , <code>KNITRO_PARAM_INFEASTOL</code>	Specifies the (relative) tolerance used for declaring infeasibility of a model.	
[Knitro, Tolerances]		p. 458
<code>XKTR_PARAM_LMSIZE</code> , <code>KNITRO_PARAM_LMSIZE</code>	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option.	
[Knitro, Limits]		p. 466
<code>XKTR_PARAM_MAXCGIT</code> , <code>KNITRO_PARAM_MAXCGIT</code>	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option.	
[Knitro, Limits]		p. 467
<code>XKTR_PARAM_MAXIT</code> , <code>KNITRO_PARAM_MAXIT</code>	Specifies the maximum number of iterations before termination.	
[Knitro, Limits]		p. 467

<code>XKTR_PARAM_MIP_BRANCHRULE</code> , <code>KNITRO_PARAM_MIP_BRANCHRULE</code>	Specifies which branching rule to use for MIP branch and bound procedure. [Branching, Knitro-MINLP]	p. 467
<code>XKTR_PARAM_MIP_GUB_BRANCH</code> , <code>KNITRO_PARAM_MIP_GUB_BRANCH</code>	Specifies whether or not to branch on generalized upper bounds (GUBs). [Branching, Knitro-MINLP]	p. 468
<code>XKTR_PARAM_MIP_HEURISTIC</code> , <code>KNITRO_PARAM_MIP_HEURISTIC</code>	Specifies which MIP heuristic search approach to apply to try to find an initial integer feasible point. [Heuristics, Knitro-MINLP]	p. 468
<code>XKTR_PARAM_MIP_HEURISTIC_MAXIT</code> , <code>KNITRO_PARAM_MIP_HEURISTIC_MAXIT</code>	Specifies the maximum number of iterations to allow for MIP heuristic, if one is enabled. [Heuristics, Knitro-MINLP]	p. 468
<code>XKTR_PARAM_MIP_IMPLICATNS</code> , <code>KNITRO_PARAM_MIP_IMPLICATNS</code>	Specifies whether or not to add constraints to the MIP derived from logical implications. [Knitro-MINLP, Presolve]	p. 468
<code>XKTR_PARAM_MIP_INTEGERTOL</code> , <code>KNITRO_PARAM_INTEGERTOL</code>	This value specifies the threshold for deciding whether or not a variable is determined to be an integer. [Knitro, Knitro-MINLP, Tolerances]	p. 458
<code>XKTR_PARAM_MIP_INTGAPABS</code> , <code>KNITRO_PARAM_INTGAPABS</code>	The absolute integrality gap stop tolerance for MIP. [Knitro, Knitro-MINLP, Tolerances]	p. 458
<code>XKTR_PARAM_MIP_INTGAPREL</code> , <code>KNITRO_PARAM_INTGAPREL</code>	The relative integrality gap stop tolerance for MIP. [Knitro, Knitro-MINLP, Tolerances]	p. 459
<code>XKTR_PARAM_MIP_KNAPSACK</code> , <code>KNITRO_PARAM_MIP_KNAPSACK</code>	Specifies rules for adding MIP knapsack cuts. [Cuts, Knitro-MINLP]	p. 469
<code>XKTR_PARAM_MIP_LPALG</code> , <code>KNITRO_PARAM_MIP_LPALG</code>	Specifies which algorithm to use for any linear programming (LP) subproblem solves that may occur in the MIP branch and bound procedure. [Knitro-MINLP, Solution Process]	p. 469
<code>XKTR_PARAM_MIP_MAXNODES</code> , <code>KNITRO_PARAM_MIP_MAXNODES</code>	Specifies the maximum number of nodes explored. [Knitro-MINLP, Limits]	p. 469
<code>XKTR_PARAM_MIP_MAXSOLVES</code> , <code>KNITRO_PARAM_MIP_MAXSOLVES</code>	Specifies the maximum number of subproblem solves allowed (0 means no limit). [Knitro-MINLP, Limits]	p. 470
<code>XKTR_PARAM_MIP_METHOD</code> , <code>KNITRO_PARAM_MIP_METHOD</code>	Specifies which MIP method to use. [Knitro-MINLP, Solution Process]	p. 470
<code>XKTR_PARAM_MIP_OUTINTERVAL</code> , <code>KNITRO_PARAM_MIP_OUTINTERVAL</code>	Specifies node printing interval for <code>XKTR_PARAM_MIP_OUTLEVEL</code> when <code>XKTR_PARAM_MIP_OUTLEVEL > 0</code> . [Knitro-MINLP, Logging]	p. 470
<code>XKTR_PARAM_MIP_OUTLEVEL</code> , <code>KNITRO_PARAM_MIP_OUTLEVEL</code>	Specifies how much MIP information to print. [Knitro-MINLP, Logging]	p. 470

<code>XKTR_PARAM_MIP_PSEUDOINIT</code> , <code>KNITRO_PARAM_MIP_PSEUDOINIT</code>	Specifies the method used to initialize pseudo-costs corresponding to variables that have not yet been branched on in the MIP method.	
	[Branching, Knitro-MINLP]	p. 471
<code>XKTR_PARAM_MIP_ROOTALG</code> , <code>KNITRO_PARAM_MIP_ROOTALG</code>	Specifies which algorithm to use for the root node solve in MIP (same options as <code>XKTR_PARAM_ALGORITHM</code> user option).	
	[Knitro-MINLP, Solution Process]	p. 471
<code>XKTR_PARAM_MIP_ROUNDING</code> , <code>KNITRO_PARAM_MIP_ROUNDING</code>	Specifies the MIP rounding rule to apply.	
	[Knitro-MINLP]	p. 471
<code>XKTR_PARAM_MIP_SELECTRULE</code> , <code>KNITRO_PARAM_MIP_SELECTRULE</code>	Specifies the MIP select rule for choosing the next node in the branch and bound tree.	
	[Knitro-MINLP]	p. 471
<code>XKTR_PARAM_MIP_STRONG_CANDLIM</code> , <code>KNITRO_PARAM_MIP_STRONG_CANDLIM</code>	Specifies the maximum number of candidates to explore for MIP strong branching.	
	[Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_STRONG_LEVEL</code> , <code>KNITRO_PARAM_MIP_STRONG_LEVEL</code>	Specifies the maximum number of tree levels on which to perform MIP strong branching.	
	[Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_STRONG_MAXIT</code> , <code>KNITRO_PARAM_MIP_STRONG_MAXIT</code>	Specifies the maximum number of iterations to allow for MIP strong branching solves.	
	[Branching, Knitro-MINLP, Limits]	p. 472
<code>XKTR_PARAM_MIP_TERMINATE</code> , <code>KNITRO_PARAM_MIP_TERMINATE</code>	Specifies conditions for terminating the MIP algorithm.	
	[Knitro-MINLP]	p. 472
<code>XKTR_PARAM_OBJRANGE</code> , <code>KNITRO_PARAM_OBJRANGE</code>	Specifies the extreme limits of the objective function for purposes of determining unboundedness.	
	[Knitro, Limits]	p. 459
<code>XKTR_PARAM_OPTTOL</code> , <code>KNITRO_PARAM_OPTTOL</code>	Specifies the final relative stopping tolerance for the KKT (optimality) error.	
	[Knitro, Tolerances]	p. 459
<code>XKTR_PARAM_OPTTOLABS</code> , <code>KNITRO_PARAM_OPTTOLABS</code>	Specifies the final absolute stopping tolerance for the KKT (optimality) error.	
	[Knitro, Tolerances]	p. 459
<code>XKTR_PARAM_OUTLEV</code> , <code>KNITRO_PARAM_OUTLEV</code>	Controls the level of output produced by Knitro.	
	[Knitro, Logging]	p. 473
<code>XKTR_PARAM_PRESOLVE</code> , <code>KNITRO_PARAM_PRESOLVE</code>	Determine whether or not to use the Knitro presolver to try to simplify the model by removing variables or constraints.	
	[Knitro, Presolve]	p. 473
<code>XKTR_PARAM_PRESOLVE_TOL</code> , <code>KNITRO_PARAM_PRESOLVE_TOL</code>	Determines the tolerance used by the Knitro presolver to remove variables and constraints from the model.	
	[Knitro, Presolve, Tolerances]	p. 460
<code>XKTR_PARAM_SCALE</code> , <code>KNITRO_PARAM_SCALE</code>	Performs a scaling of the objective and constraint functions based on their values at the initial point.	
	[Knitro, Numerics]	p. 473

<code>XKTR_PARAM_SOC</code> , <code>KNITRO_PARAM_SOC</code>	Specifies whether or not to try second order corrections (SOC).	
	[Knitro]	p. 474
<code>XKTR_PARAM_XTOL</code> , <code>KNITRO_PARAM_XTOL</code>	The optimization process will terminate if the relative change in all components of the solution point estimate is less than <code>xtol</code> .	
	[Knitro, Tolerances]	p. 460

24.1 Double control parameters

These double control parameters can be set using `XSLPsetdblcontrol` using the Xpress NonLinear API, `XNLPsetsolverdoublecontrol` using the XNLP API and `setparam` in Mosel using module `mmxnlp`.

XKTR_PARAM_BAR_FEASMODETOL, KNITRO_PARAM_BAR_FEASMODETOL

Description	Specifies the tolerance in equation that determines whether Knitro will force subsequent iterates to remain feasible.
Type	Double
Category	Control
Topic areas	Knitro, Tolerances
Note	The tolerance applies to all inequality constraints in the problem. This option only has an effect if option <code>XKTR_PARAM_BAR_FEASIBLE = stay</code> or <code>XKTR_PARAM_BAR_FEASIBLE = get_stay</code> .
Default value	1.0e-4

XKTR_PARAM_BAR_INITMU, KNITRO_PARAM_BAR_INITMU

Description	Specifies the initial value for the barrier parameter : μ used with the barrier algorithms. This option has no effect on the Active Set algorithm.
Type	Double
Category	Control
Topic area	Knitro
Default value	1.0e-1

XKTR_PARAM_DELTA, KNITRO_PARAM_DELTA

Description	Specifies the initial trust region radius scaling factor used to determine the initial trust region size.
Type	Double
Category	Control
Topic area	Knitro
Default value	1.0e0

XKTR_PARAM_FEASTOL, KNITRO_PARAM_FEASTOL

Description	Specifies the final relative stopping tolerance for the feasibility error.
Type	Double
Category	Control
Topic areas	Knitro, Tolerances

Note Smaller values of feastol result in a higher degree of accuracy in the solution with respect to feasibility.

Default value 1.0e-6

XKTR_PARAM_FEASTOLABS, KNITRO_PARAM_FEASTOLABS

Description Specifies the final absolute stopping tolerance for the feasibility error.

Type Double

Category Control

Topic areas Knitro, Tolerances

Note Smaller values of feastol_abs result in a higher degree of accuracy in the solution with respect to feasibility.

Default value 0.0e0

XKTR_PARAM_INFEASTOL, KNITRO_PARAM_INFEASTOL

Description Specifies the (relative) tolerance used for declaring infeasibility of a model.

Type Double

Category Control

Topic areas Knitro, Tolerances

Note Smaller values of infeasol make it more difficult to satisfy the conditions Knitro uses for detecting infeasible models. If you believe Knitro incorrectly declares a model to be infeasible, then you should try a smaller value for infeasol.

Default value 1.0e-8

XKTR_PARAM_MIP_INTEGERTOL, KNITRO_PARAM_INTEGERTOL

Description This value specifies the threshold for deciding whether or not a variable is determined to be an integer.

Type Double

Category Control

Topic areas Knitro, Knitro-MINLP, Tolerances

Default value 1.0e-8

XKTR_PARAM_MIP_INTGAPABS, KNITRO_PARAM_INTGAPABS

Description The absolute integrality gap stop tolerance for MIP.

Type Double

Category Control

Topic areas Knitro, Knitro-MINLP, Tolerances

Default value 1.0e-6

XKTR_PARAM_MIP_INTGAPREL, KNITRO_PARAM_INTGAPREL

Description	The relative integrality gap stop tolerance for MIP.
Type	Double
Category	Control
Topic areas	Knitro, Knitro-MINLP, Tolerances
Default value	1.0e-6

XKTR_PARAM_OBJRANGE, KNITRO_PARAM_OBJRANGE

Description	Specifies the extreme limits of the objective function for purposes of determining unboundedness.
Type	Double
Category	Control
Topic areas	Knitro, Limits
Note	If the magnitude of the objective function becomes greater than objrange for a feasible iterate, then the problem is determined to be unbounded and Knitro proceeds no further.
Default value	1.0e20

XKTR_PARAM_OPTTOL, KNITRO_PARAM_OPTTOL

Description	Specifies the final relative stopping tolerance for the KKT (optimality) error.
Type	Double
Category	Control
Topic areas	Knitro, Tolerances
Note	Smaller values of opttol result in a higher degree of accuracy in the solution with respect to optimality.
Default value	1.0e-6

XKTR_PARAM_OPTTOLABS, KNITRO_PARAM_OPTTOLABS

Description	Specifies the final absolute stopping tolerance for the KKT (optimality) error.
Type	Double
Category	Control
Topic areas	Knitro, Tolerances
Note	Smaller values of opttol_abs result in a higher degree of accuracy in the solution with respect to optimality.
Default value	0.0e0

XKTR_PARAM_PRESOLVE_TOL, KNITRO_PARAM_PRESOLVE_TOL

Description	Determines the tolerance used by the Knitro presolver to remove variables and constraints from the model.
Type	Double
Category	Control
Topic areas	Knitro, Presolve, Tolerances
Note	If you believe the Knitro presolver is incorrectly modifying the model, use a smaller value for this tolerance (or turn the presolver off).
Default value	1.0e-6

XKTR_PARAM_XTOL, KNITRO_PARAM_XTOL

Description	The optimization process will terminate if the relative change in all components of the solution point estimate is less than xt看ol.
Type	Double
Category	Control
Topic areas	Knitro, Tolerances
Note	If using the Interior/Direct or Interior/CG algorithm and the barrier parameter is still large, Knitro will first try decreasing the barrier parameter before terminating.
Default value	1.0e-15

24.2 Integer control parameters

These integer control parameters can be set using XSLPsetintcontrol using the Xpress NonLinear API, XNLPsetsolverintcontrol using the XNLP API and setparam in Mosel using module mmxnlp.

XKTR_PARAM_ALGORITHM, KNITRO_PARAM_ALGORITHM

Description	Indicates which algorithm to use to solve the problem
Type	Integer
Category	Control
Topic areas	Knitro, Solution Process
Values	0 (auto) let Knitro automatically choose an algorithm, based on the problem characteristics. 1 (direct) use the Interior/Direct algorithm. 2 (cg) use the Interior/CG algorithm. 3 (active) use the Active Set algorithm. 4 (sqp) use the SQP algorithm. 5 (multi) run all algorithms, perhaps in parallel.
Default value	0

XKTR_PARAM_BAR_DIRECTINTERVAL, KNITRO_PARAM_BAR_DIRECTINTERVAL

Description	Controls the maximum number of consecutive conjugate gradient (CG) steps before Knitro will try to enforce that a step is taken using direct linear algebra.
Type	Integer
Category	Control
Topic areas	Knitro, Limits
Note	This option is only valid for the Interior/Direct algorithm and may be useful on problems where Knitro appears to be taking lots of conjugate gradient steps. Setting bar_directinterval to 0 will try to enforce that only direct steps are taken which may produce better results on some problems.
Default value	10

XKTR_PARAM_BAR_FEASIBLE, KNITRO_PARAM_BAR_FEASIBLE

Description	Specifies whether special emphasis is placed on getting and staying feasible in the interior-point algorithms.
Type	Integer
Category	Control
Topic area	Knitro

Values	0	(no) No special emphasis on feasibility.
	1	(stay) Iterates must satisfy inequality constraints once they become sufficiently feasible.
	2	(get) Special emphasis is placed on getting feasible before trying to optimize.
	3	(get_stay) Implement both options 1 and 2 above.
Note	This option can only be used with the Interior/Direct and Interior/CG algorithms. If <code>bar_feasible = stay</code> or <code>bar_feasible = get_stay</code> , this will activate the feasible version of Knitro. The feasible version of Knitro will force iterates to strictly satisfy inequalities, but does not require satisfaction of equality constraints at intermediate iterates. This option and the <code>honorbnds</code> option may be useful in applications where functions are undefined outside the region defined by inequalities. The initial point must satisfy inequalities to a sufficient degree; if not, Knitro may generate infeasible iterates and does not switch to the feasible version until a sufficiently feasible point is found. Sufficient satisfaction occurs at a point x if it is true for all inequalities that $cl + tol \leq c(x) \leq cu - tol$. The constant <code>tol</code> is determined by the option <code>bar_feasmodetol</code> . If <code>bar_feasible = get</code> or <code>bar_feasible = get_stay</code> , Knitro will place special emphasis on first trying to get feasible before trying to optimize.	
Default value	0	

XKTR_PARAM_BAR_INITPT, KNITRO_PARAM_BAR_INITPT

Description	Indicates whether an initial point strategy is used with barrier algorithms.	
Type	Integer	
Category	Control	
Topic area	Knitro	
Values	0	(auto) Let Knitro automatically choose the strategy.
	1	(yes) Shift the initial slacks and multipliers to improve barrier algorithm performance.
	2	(no) Do not alter the initial slacks and multipliers.
Note	This option has no effect on the Active Set algorithm.	
Default value	0	

XKTR_PARAM_BAR_MAXBACKTRACK, KNITRO_PARAM_BAR_MAXBACKTRACK

Description	Indicates the maximum allowable number of backtracks during the linesearch of the Interior/Direct algorithm before reverting to a CG step.	
Type	Integer	
Category	Control	
Topic areas	Knitro, Limits	
Note	Increasing this value will make the Interior/Direct algorithm less likely to take CG steps. If the Interior/Direct algorithm is taking a large number of CG steps (as indicated by a positive value for 'Gits' in the output), this may improve performance. This option has no effect on the Active Set algorithm.	
Default value	3	

XKTR_PARAM_BAR_MAXCROSSIT, KNITRO_PARAM_BAR_MAXCROSSIT

Description	Specifies the maximum number of crossover iterations before termination.
Type	Integer
Category	Control
Topic areas	Knitro, Limits
Note	If the value is positive and the algorithm in operation is Interior/Direct or Interior/CG, then Knitro will crossover to the Active Set algorithm near the solution. The Active Set algorithm will then perform at most bar_maxcrossit iterations to get a more exact solution. If the value is 0, no Active Set crossover occurs and the interior-point solution is the final result. If Active Set crossover is unable to improve the approximate interior-point solution, then Knitro will restore the interior-point solution. In some cases (especially on large-scale problems or difficult degenerate problems) the cost of the crossover procedure may be significant - for this reason, crossover is disabled by default. Enabling crossover generally provides a more accurate solution than Interior/Direct or Interior/CG.
Default value	0

XKTR_PARAM_BAR_MAXREFACTOR, KNITRO_PARAM_BAR_MAXREFACTOR

Description	Indicates the maximum number of refactorizations of the KKT system per iteration of the Interior/Direct algorithm before reverting to a CG step.
Type	Integer
Category	Control
Topic areas	Knitro, Limits
Note	These refactorizations are performed if negative curvature is detected in the model. Rather than reverting to a CG step, the Hessian matrix is modified in an attempt to make the subproblem convex and then the KKT system is refactorized. Increasing this value will make the Interior/Direct algorithm less likely to take CG steps. If the Interior/Direct algorithm is taking a large number of CG steps (as indicated by a positive value for "CGits" in the output), this may improve performance. This option has no effect on the Active Set algorithm.
Default value	-1

XKTR_PARAM_BAR_MURULE, KNITRO_PARAM_BAR_MURULE

Description	Indicates which strategy to use for modifying the barrier parameter mu in the barrier algorithms.
Type	Integer
Category	Control
Topic area	Knitro

Values	0	(auto) Let Knitro automatically choose the strategy.
	1	(monotone) Monotonically decrease the barrier parameter. Available for both barrier algorithms.
	2	(adaptive) Use an adaptive rule based on the complementarity gap to determine the value of the barrier parameter. Available for both barrier algorithms.
	3	(probing) Use a probing (affine-scaling) step to dynamically determine the barrier parameter. Available only for the Interior/Direct algorithm.
	4	(dampmpc) Use a Mehrotra predictor-corrector type rule to determine the barrier parameter, with safeguards on the corrector step. Available only for the Interior/Direct algorithm.
	5	(fullmpc) Use a Mehrotra predictor-corrector type rule to determine the barrier parameter, without safeguards on the corrector step. Available only for the Interior/Direct algorithm.
	6	(quality) Minimize a quality function at each iteration to determine the barrier parameter. Available only for the Interior/Direct algorithm.
Note	Not all strategies are available for both barrier algorithms. This option has no effect on the Active Set algorithm.	
Default value	0	

XKTR_PARAM_BAR_PENCONS, KNITRO_PARAM_BAR_PENCONS

Description	Indicates whether a penalty approach is applied to the constraints.	
Type	Integer	
Category	Control	
Topic area	Knitro	
Values	0	(auto) Let Knitro automatically choose the strategy.
	1	(none) No constraints are penalized.
	2	(all) A penalty approach is applied to all general constraints.
Note	Using a penalty approach may be helpful when the problem has degenerate or difficult constraints. It may also help to more quickly identify infeasible problems, or achieve feasibility in problems with difficult constraints. This option has no effect on the Active Set algorithm.	
Default value	0	

XKTR_PARAM_BAR_PENRULE, KNITRO_PARAM_BAR_PENRULE

Description	Indicates which penalty parameter strategy to use for determining whether or not to accept a trial iterate.	
Type	Integer	
Category	Control	
Topic area	Knitro	
Values	0	(auto) Let Knitro automatically choose the strategy.
	1	(single) Use a single penalty parameter in the merit function to weight feasibility versus optimality.
	2	(flex) Use a more tolerant and flexible step acceptance procedure based on a range of penalty parameter values.

Note This option has no effect on the Active Set algorithm.

Default value 0

XKTR_PARAM_BAR_SWITCHRULE, KNITRO_PARAM_BAR_SWITCHRULE

Description Indicates whether or not the barrier algorithms will allow switching from an optimality phase to a pure feasibility phase.

Type Integer

Category Control

Topic area Knitro

Values

0	(auto) Let Knitro determine the switching procedure.
1	(never) Never switch to feasibility phase.
2	(level1) Allow switches to feasibility phase.
3	(level2) Use a more aggressive switching rule.

Note This option has no effect on the Active Set algorithm.

Default value 0

XKTR_PARAM_GRADOPT, KNITRO_PARAM_GRADOPT

Description Specifies how to compute the gradients of the objective and constraint functions.

Type Integer

Category Control

Topic areas Knitro, Derivatives

Values

1	(exact) User provides a routine for computing the exact gradients.
2	(forward) Knitro computes gradients by forward finite-differences.
3	(central) Knitro computes gradients by central finite differences.

Note It is highly recommended to provide exact gradients if at all possible as this greatly impacts the performance of the code.

Default value 1

XKTR_PARAM_HESSOPT, KNITRO_PARAM_HESSOPT

Description Specifies how to compute the (approximate) Hessian of the Lagrangian.

Type Integer

Category Control

Topic areas Knitro, Derivatives

Values	1	(exact) User provides a routine for computing the exact Hessian.
	2	(bfgs) Knitro computes a (dense) quasi-Newton BFGS Hessian.
	3	(sr1) Knitro computes a (dense) quasi-Newton SR1 Hessian.
	4	(finite_diff) Knitro computes Hessian-vector products using finite-differences.
	5	(product) User provides a routine to compute the Hessian-vector products.
	6	(lbfgs) Knitro computes a limited-memory quasi-Newton BFGS Hessian (its size is determined by the option lmsize).

Note Options hessopt = 4 and hessopt = 5 are not available with the Interior/Direct algorithm. Knitro usually performs best when the user provides exact Hessians (hessopt = 1) or exact Hessian-vector products (hessopt = 5). If neither can be provided but exact gradients are available (i.e., gradopt = 1), then hessopt = 4 is recommended. This option is comparable in terms of robustness to the exact Hessian option and typically not much slower in terms of time, provided that gradient evaluations are not a dominant cost. If exact gradients cannot be provided, then one of the quasi-Newton options is preferred. Options hessopt = 2 and hessopt = 3 are only recommended for small problems ($n \leq 1000$) since they require working with a dense Hessian approximation. Option hessopt = 6 should be used for large problems.

Default value 1

XKTR_PARAM_HONORBND, KNITRO_PARAM_HONORBND

Description	Indicates whether or not to enforce satisfaction of simple variable bounds throughout the optimization.	
Type	Integer	
Category	Control	
Topic area	Knitro	
Values	0	(no) Knitro does not require that the bounds on the variables be satisfied at intermediate iterates.
	1	(always) Knitro enforces that the initial point and all subsequent solution estimates satisfy the bounds on the variables.
	2	(initpt) Knitro enforces that the initial point satisfies the bounds on the variables.
Note	This option and the bar_feasible option may be useful in applications where functions are undefined outside the region defined by inequalities.	
Default value	2	

XKTR_PARAM_LMSIZE, KNITRO_PARAM_LMSIZE

Description	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option.
Type	Integer
Category	Control
Topic areas	Knitro, Limits
Note	The value must be between 1 and 100 and is only used with <code>XKTR_PARAM_HESSOPT = 6</code> . Larger values may give a more accurate, but more expensive, Hessian approximation. Smaller values may give a less accurate, but faster, Hessian approximation. When using the limited memory BFGS approach it is recommended to experiment with different values of this parameter.

Default value 10

XKTR_PARAM_MAXCGIT, KNITRO_PARAM_MAXCGIT

Description	Specifies the number of limited memory pairs stored when approximating the Hessian using the limited-memory quasi-Newton BFGS option.	
Type	Integer	
Category	Control	
Topic areas	Knitro, Limits	
Values	0	Let Knitro automatically choose a value based on the problem size.
	n	At most n>0 CG iterations may be performed during one minor iteration of Knitro.
Default value	0	

XKTR_PARAM_MAXIT, KNITRO_PARAM_MAXIT

Description	Specifies the maximum number of iterations before termination.	
Type	Integer	
Category	Control	
Topic areas	Knitro, Limits	
Values	0	Let Knitro automatically choose a value based on the problem type. Currently Knitro sets this value to 10000 for LPs/NLPs and 3000 for MIP problems.
	n	At most n>0 iterations may be performed before terminating.
Default value	0	

XKTR_PARAM_MIP_BRANCHRULE, KNITRO_PARAM_MIP_BRANCHRULE

Description	Specifies which branching rule to use for MIP branch and bound procedure.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Branching	
Values	0	(auto) Let Knitro automatically choose the branching rule.
	1	(most_frac) Use most fractional (most infeasible) branching.
	2	(pseudocost) Use pseudo-cost branching.
	3	(strong) Use strong branching (see options XKTR_PARAM_MIP_STRONG_CANDLIM, XKTR_PARAM_MIP_STRONG_LEVEL, XKTR_PARAM_MIP_STRONG_MAXIT for further control of strong branching procedure).
Default value	0	

XKTR_PARAM_MIP_GUB_BRANCH, KNITRO_PARAM_MIP_GUB_BRANCH

Description	Specifies whether or not to branch on generalized upper bounds (GUBs).	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Branching	
Values	0	(no) Do not branch on GUBs.
	1	(yes) Allow branching on GUBs.
Default value	0	

XKTR_PARAM_MIP_HEURISTIC, KNITRO_PARAM_MIP_HEURISTIC

Description	Specifies which MIP heuristic search approach to apply to try to find an initial integer feasible point.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Heuristics	
Values	0	(auto) Let Knitro choose the heuristic to apply (if any).
	1	(none) No heuristic search applied.
	2	(feaspump) Apply feasibility pump heuristic.
	3	(mpec) Apply heuristic based on MPEC formulation.
Note	If a heuristic search procedure is enabled, it will run for at most mip_heuristic_maxit iterations, before starting the branch and bound procedure.	
Default value	0	

XKTR_PARAM_MIP_HEURISTIC_MAXIT, KNITRO_PARAM_MIP_HEURISTIC_MAXIT

Description	Specifies the maximum number of iterations to allow for MIP heuristic, if one is enabled.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Heuristics	
Default value	100	

XKTR_PARAM_MIP_IMPLICATNS, KNITRO_PARAM_MIP_IMPLICATNS

Description	Specifies whether or not to add constraints to the MIP derived from logical implications.	
Type	Integer	
Category	Control	

Topic areas	Knitro-MINLP, Presolve	
Values	0	(no) Do not add constraints from logical implications.
	1	(yes) Knitro adds constraints from logical implications.
Default value	1	

XKTR_PARAM_MIP_KNAPSACK, KNITRO_PARAM_MIP_KNAPSACK

Description	Specifies rules for adding MIP knapsack cuts.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Cuts	
Values	0	(none) Do not add knapsack cuts.
	1	(ineqs) Add cuts derived from inequalities only.
	2	(ineqs_eqs) Add cuts derived from both inequalities and equalities.
Default value	1	

XKTR_PARAM_MIP_LPALG, KNITRO_PARAM_MIP_LPALG

Description	Specifies which algorithm to use for any linear programming (LP) subproblem solves that may occur in the MIP branch and bound procedure.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Solution Process	
Values	0	(auto) Let Knitro automatically choose an algorithm, based on the problem characteristics.
	1	(direct) Use the Interior/Direct (barrier) algorithm.
	2	(cg) Use the Interior/CG (barrier) algorithm.
	3	(active) Use the Active Set (simplex) algorithm.
Note	LP subproblems may arise if the problem is a mixed integer linear program (MILP), or if using XKTR_PARAM_MIP_METHOD = HQG. (Nonlinear programming subproblems use the algorithm specified by the algorithm option.)	
Default value	0	

XKTR_PARAM_MIP_MAXNODES, KNITRO_PARAM_MIP_MAXNODES

Description	Specifies the maximum number of nodes explored.
Type	Integer
Category	Control
Topic areas	Knitro-MINLP, Limits
Note	Zero value means no limit.
Default value	100000

XKTR_PARAM_MIP_MAXSOLVES, KNITRO_PARAM_MIP_MAXSOLVES

Description	Specifies the maximum number of subproblem solves allowed (0 means no limit).
Type	Integer
Category	Control
Topic areas	Knitro-MINLP, Limits
Default value	200000

XKTR_PARAM_MIP_METHOD, KNITRO_PARAM_MIP_METHOD

Description	Specifies which MIP method to use.
Type	Integer
Category	Control
Topic areas	Knitro-MINLP, Solution Process
Values	0 (auto) Let Knitro automatically choose the method. 1 (BB) Use the standard branch and bound method. 2 (HQG) Use the hybrid Quesada-Grossman method (for convex, nonlinear problems only).
Default value	0

XKTR_PARAM_MIP_OUTINTERVAL, KNITRO_PARAM_MIP_OUTINTERVAL

Description	Specifies node printing interval for XKTR_PARAM_MIP_OUTLEVEL when XKTR_PARAM_MIP_OUTLEVEL > 0.
Type	Integer
Category	Control
Topic areas	Knitro-MINLP, Logging
Values	0 Print output every node. 2 Print output every 2nd node. N Print output every Nth node.
Default value	10

XKTR_PARAM_MIP_OUTLEVEL, KNITRO_PARAM_MIP_OUTLEVEL

Description	Specifies how much MIP information to print.
Type	Integer
Category	Control
Topic areas	Knitro-MINLP, Logging
Values	0 (none) Do not print any MIP node information. 1 (iters) Print one line of output for every node.
Default value	1

XKTR_PARAM_MIP_PSEUDOINIT, KNITRO_PARAM_MIP_PSEUDOINIT

Description	Specifies the method used to initialize pseudo-costs corresponding to variables that have not yet been branched on in the MIP method.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Branching	
Values	0	Let Knitro automatically choose the method.
	1	Initialize using the average value of computed pseudo-costs.
	2	Initialize using strong branching.
Default value	0	

XKTR_PARAM_MIP_ROOTALG, KNITRO_PARAM_MIP_ROOTALG

Description	Specifies which algorithm to use for the root node solve in MIP (same options as XKTR_PARAM_ALGORITHM user option).	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Solution Process	
Default value	0	

XKTR_PARAM_MIP_ROUNDING, KNITRO_PARAM_MIP_ROUNDING

Description	Specifies the MIP rounding rule to apply.	
Type	Integer	
Category	Control	
Topic area	Knitro-MINLP	
Values	0	(auto) Let Knitro choose the rounding rule.
	1	(none) Do not round if a node is infeasible.
	2	(heur_only) Round using a fast heuristic only.
	3	(nlp_sometimes) Round and solve a subproblem if likely to succeed.
	4	(nlp_always) Always round and solve a subproblem.
Default value	0	

XKTR_PARAM_MIP_SELECTRULE, KNITRO_PARAM_MIP_SELECTRULE

Description	Specifies the MIP select rule for choosing the next node in the branch and bound tree.	
Type	Integer	
Category	Control	

Topic area	Knitro-MINLP	
Values	0	(auto) Let Knitro choose the node selection rule.
	1	(depth_first) Search the tree using a depth first procedure.
	2	(best_bound) Select the node with the best relaxation bound.
	3	(combo_1) Use depth first unless pruned, then best bound.
Default value	0	

XKTR_PARAM_MIP_STRONG_CANDLIM, KNITRO_PARAM_MIP_STRONG_CANDLIM

Description	Specifies the maximum number of candidates to explore for MIP strong branching.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Branching, Limits	
Default value	10	

XKTR_PARAM_MIP_STRONG_LEVEL, KNITRO_PARAM_MIP_STRONG_LEVEL

Description	Specifies the maximum number of tree levels on which to perform MIP strong branching.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Branching, Limits	
Default value	10	

XKTR_PARAM_MIP_STRONG_MAXIT, KNITRO_PARAM_MIP_STRONG_MAXIT

Description	Specifies the maximum number of iterations to allow for MIP strong branching solves.	
Type	Integer	
Category	Control	
Topic areas	Knitro-MINLP, Branching, Limits	
Default value	1000	

XKTR_PARAM_MIP_TERMINATE, KNITRO_PARAM_MIP_TERMINATE

Description	Specifies conditions for terminating the MIP algorithm.	
Type	Integer	
Category	Control	
Topic area	Knitro-MINLP	
Values	0	(optimal) Terminate at optimum.
	1	(feasible) Terminate at first integer feasible point.
Default value	0	

XKTR_PARAM_OUTLEV, KNITRO_PARAM_OUTLEV

Description	Controls the level of output produced by Knitro.	
Type	Integer	
Category	Control	
Topic areas	Knitro, Logging	
Values	0	(none) Printing of all output is suppressed.
	1	(summary) Print only summary information.
	2	(iter_10) Print basic information every 10 iterations.
	3	(iter) Print basic information at each iteration.
	4	(iter_verbose) Print basic information and the function count at each iteration.
	5	(iter_x) Print all the above, and the values of the solution vector x.
	6	(all) Print all the above, and the values of the constraints c at x and the Lagrange multipliers lambda.
Default value	2	

XKTR_PARAM_PREOLVE, KNITRO_PARAM_PREOLVE

Description	Determine whether or not to use the Knitro presolver to try to simplify the model by removing variables or constraints. Specifies conditions for terminating the MIP algorithm.	
Type	Integer	
Category	Control	
Topic areas	Knitro, Presolve	
Values	0	(none) Do not use Knitro presolver.
	1	(basic) Use the Knitro basic presolver.
Default value	1	

XKTR_PARAM_SCALE, KNITRO_PARAM_SCALE

Description	Performs a scaling of the objective and constraint functions based on their values at the initial point.	
Type	Integer	
Category	Control	
Topic areas	Knitro, Numerics	
Values	0	(no) No scaling is performed.
	1	(yes) Knitro is allowed to scale the objective function and constraints.
Note	If scaling is performed, all internal computations, including the stopping tests, are based on the scaled values.	
Default value	1	

XKTR_PARAM_SOC, KNITRO_PARAM_SOC

Description	Specifies whether or not to try second order corrections (SOC).	
Type	Integer	
Category	Control	
Topic area	Knitro	
Values	0	(no) No second order correction steps are attempted.
	1	(maybe) Second order correction steps may be attempted on some iterations.
	2	(yes) Second order correction steps are always attempted if the original step is rejected and there are nonlinear constraints.
Note	A second order correction may be beneficial for problems with highly nonlinear constraints.	
Default value	1	

Appendix

APPENDIX A

The Xpress-SLP Log

The Xpress-SLP log consists of log lines of two different types: the output of the underlying XPRS optimizer, and the log of XSLP itself.

By default, messages produced by the nonlinear code are sent to the normal XPRS message callback as controlled by `XSLP_ECHOXPRSMESSAGES`. It may also be intercepted by a user function using the user output callback; see `XSLPsetcbmessage`. Users need to define a callback function and print messages to the screen themselves if they wish output to be displayed.

A.0.1 Logging controls

General SLP logging

<code>XPRS_OUTPUTLOG</code>	Logging level of the underlying XPRS problem
<code>XPRS_LPLOG</code>	Logging frequency for solving the linearization
<code>XPRS_MIPLOG</code>	Logging frequency for the MIP solver

Logging for the underlying XPRS problem

<code>XSLP_LOG</code>	Level of SLP logging (iteration, penalty, convergence)
<code>XSLP_SLPLOG</code>	Logging frequency for SLP iterations
<code>XSLP_MIPLOG</code>	MI-SLP specific logging

Special logging settings

<code>XPRS_DCLOG</code>	Logging of delayed constraint activation
<code>XSLP_ERRORTOL_P</code>	Absolute tolerance for printing error vectors

A.0.2 The structure of the log

The typical log with the default settings starts with statistics about the problem sizes. On the `Polygon1.mps` example, using the XSLP console program this looks like

```
[xpress mps] readprob Polygon1.mps
Reading Problem Polygon
Problem Statistics
    11 (    0 spare) rows
    10 (    4 spare) structural columns
     8 (    0 spare) non-zero elements
MIP Entity Statistics
    0 entities          0 sets          0 set members
    PV:      0    DC:      0    DR:      0    EC:      0
    IV:      0    RX:      0    TX:      0    SB:      0
    UF:      0    WT:      0    XV:      0    Total:      0
Xpress-SLP Statistics:
    7 coefficients
    9 SLP variables
```

The standard XPRS optimizer problem loading statistics is extended with a report about the special structures possibly present in the problem, including DC (delayed constraints), DR (determining rows), EC (enforced constraints), IV (initial values), RX/TX (relative and absolute tolerances), SB (initial step bounds), UF (user functions), WT (initial row weights), followed by a statistics about the number of SLP coefficients and variables.

```

-----
SLP iteration 1, 0s
Minimizing LP Polygon
Original problem has:
    20 rows          27 cols          68 elements
Presolved problem has:
    0 rows           0 cols           0 elements

    Its      Obj Value      S   Ninf  Nneg      Sum Inf  Time
    0      828864.7136      D     0    0      .000000    0
Uncrunching matrix
    0      828864.7136      D     0    0      .000000    0
Optimal solution found
8 unconverged values (at least 1 in active constraints)
Total feasibility error costs 829100.765742

Penalty Error Vectors - Penalties scaled by 200
Variable      Activity      Penalty
BE-V1V4      1381.836001      1.000000
BE-V2V4      1381.834610      1.000000
BE-V3V4      1381.833218      1.000000
Total:      4145.503829
Error Costs: 829100.765742  Penalty Delta Costs: 0.000000  Net Objective: -236.052107
-----

SLP iteration 2, 0s
Minimizing LP Polygon
Original problem has:
    20 rows          27 cols          73 elements
Presolved problem has:
    0 rows           0 cols           0 elements

    Its      Obj Value      S   Ninf  Nneg      Sum Inf  Time
    0     -3.13860E-05      D     0    0      .000000    0
Uncrunching matrix
    0     -3.13860E-05      D     0    0      .000000    0
Optimal solution found
4 unconverged values (at least 1 in active constraints)
-----

SLP iteration 3, 0s
Minimizing LP Polygon
Original problem has:
    20 rows          27 cols          72 elements
Presolved problem has:
    0 rows           0 cols           0 elements

    Its      Obj Value      S   Ninf  Nneg      Sum Inf  Time
    0     -1.56933E-05      D     0    0      .000000    0
Uncrunching matrix
    0     -1.56933E-05      D     0    0      .000000    0
Optimal solution found

```

The default solution log consists of the optimizer output of solving the linearizations, followed by statistics of the nonlinear infeasibilities, the penalty and the objective, and the convergence status.

```

-----
Iteration summary
Itr. LPS   NetObj      ErrorSum      ErrorCost      Unconv. Extended  Action
1   0   -236.052107  4145.503829  829100.7657      8           0

```

```

 2  O  -3.13860E-05      .000000      .000000      4      0
 3  O  -1.56932E-05      .000000      .000000      0      0

Xpress-SLP stopped after 3 iterations. 0 unconverged items
No unconverged values in active constraints

```

The final iteration summary contains the following fields:

Itr: The iteration number.

LPS: The LP status of the linearization, which can take the following values:

- O Linearization is optimal
- I Linearization is infeasible
- U Linearization is unbounded
- X Solving the linearization was interrupted

NetObj: The net objective of the SLP iteration.

ErrorSum: Sum of the error delta variables. A measure of infeasibility.

ErrorCost: The value of the weighted error delta variables in the objective. A measure of the effort needed to push the model towards feasibility.

Unconv: The number of SLP variables that are not converged.

Extended: The number of SLP variables that are converged, but only by extended criteria

Action: The special actions that happened in the iteration. These can be

- * A new incumbent solution was found
- O Failed line search (non-improving)
- A Adaptive iterations were enabled
- B Enforcing step bounds
- C Variable clamping was applied
- D The determining column filter was applied
- E Some infeasible rows were enforced
- F Function evaluation error
- G Discrete variables were fixed
- I At least one working problem was unexpectedly infeasible
- K Optimality validation induces further iterations
- P The solution needed polishing, postsolve instability
- P! Solution polishing failed
- R Penalty error vectors were removed
- s Switching to primal simplex
- S Step bound induced infeasibility was repaired
- V Feasibility validation induces further iterations

The presence of a **P !** suggests that the problem is particularly hard to solve without postsolve, and the model might benefit from setting `XSLP_NOLPPOLISHING` on `XSLP_ALGORITHM` (please note, that this should only be considered if the solution polishing features is very slow or fails, as the numerical inaccuracies it aims to remove can cause other problems to the solution process).

APPENDIX B

Selecting the right algorithm for a nonlinear problem - when to use the XPRS library instead of XSLP

This chapter focuses on the nonlinear capabilities of the FICO Xpress Optimizer. While Xpress XSLP is able to efficiently solve most nonlinear problems to local optimality, there are two reasons to use other Xpress libraries:

- Subclasses of nonlinear problems for which the Xpress Optimizer features specialized algorithms that are able to solve those problems more efficiently and in larger sizes. These are notably convex quadratic programming and convex quadratically constrained problems and their mixed integer counterparts.
- Situations in which a global optimum of a nonlinear problem is required. The functionality of [FICO Xpress Global](#) allows to also solve nonconvex QPs, MIQPs, NLPs and MINLPs to proven global optimality.

It is also possible to separate the convex quadratic information from the rest of XSLP, and let the Xpress XPRS optimizer handle those directly. Doing so is good modelling practice, but emphasis must be placed on that the Optimizer can only handle convex quadratic constraints, unless a license for FICO Xpress Global is present.

B.0.1 Convex Quadratic Programs (QPs)

Convex Quadratic Programming (QP) problems are an extension of Linear Programming (LP) problems where the objective function may include a second order polynomial. The FICO Xpress Optimizer can be used directly for solving QP problems (and the Mixed Integer version MIQP).

If there are no other nonlinearities in the problem, the XPRS library provides specialized algorithms for the solution of convex QP (and MIQP) problems, that are much more efficient than solving the problem as a general nonlinear problem with XSLP.

B.0.2 Convex 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:

$$\begin{aligned}
 &\text{minimize:} && c_1x_1 + \dots + c_nx_n + x^T Q_0 x \\
 &\text{subject to:} && a_{11}x_1 + \dots + a_{1n}x_n + x^T Q_1 x \leq b_1 \\
 &&& \dots \\
 &&& a_{m1}x_1 + \dots + a_{mn}x_n + x^T Q_m x \leq b_m \\
 &&& l_1 \leq x_1 \leq u_1, \dots, l_n \leq x_n \leq u_n
 \end{aligned}$$

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

If there are no other nonlinearities in the problem, the XPRS library provides specialized algorithms for the solution of convex QCQP (and the integer counterpart MIQCQP) problems, that are much more efficient than solving the problem as a general nonlinear problem with XSLP.

B.0.3 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 optimums, they may even introduce non-connected feasible regions as shown in Figure B.1.

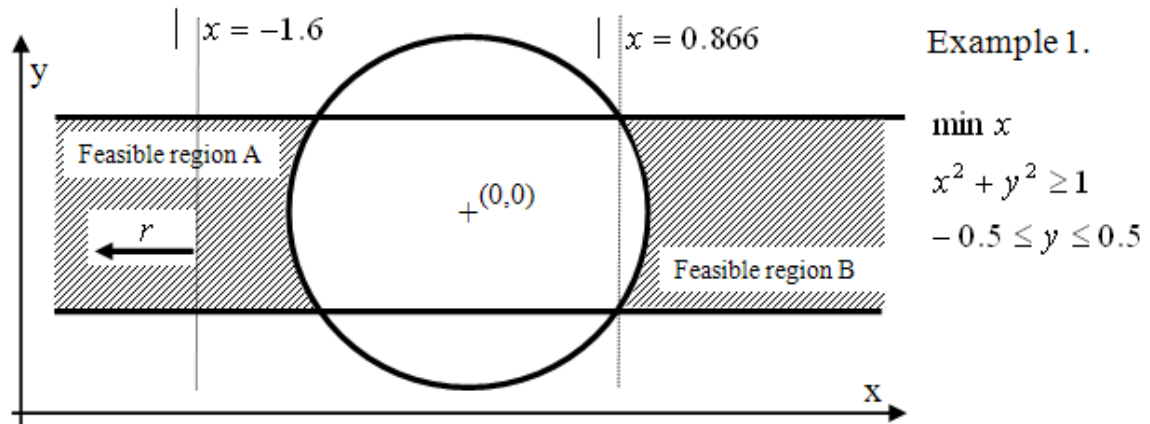


Figure B.1: Non-connected feasible regions

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

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

B.0.4 Characterizing Convexity in Quadratic Constraints

A quadratic constraint of form

$$a_1x_1 + \dots + a_nx_n + x^T Qx \leq b$$

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

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

It follows that for greater or equal constraints

$$a_1 x_1 + \dots + a_n x_n - x^T Q x \geq b$$

the negative of Q shall be PSD.

A nontrivial quadratic equality constraint (one for which not every coefficient is zero) always defines a nonconvex region, therefore those must be modelled as XSLP structures.

There is no straightforward way of checking if a matrix is PSD or not. An intuitive way of checking this property, is that the quadratic part shall always only make a constraint harder to satisfy (i.e. taking the quadratic part away shall always be a relaxation of the original problem).

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 defined;
2. having $-x^2$ in any quadratic expression in a less or equal, or having x^2 in any greater or equal row.

As a general rule, a convex quadratic objective and convex quadratic constraints are best handled by the XPRS library; while all nonconvex counterparts should be modelled as XSLP structures. It then depends on the application whether FICO Xpress Nonlinear is required (e.g., because of user functions) or [FICO Xpress Global](#) is the necessary solver (e.g., because a bound on the optimal objective is needed). Both use the same modelling API and, license and required features permitting, users can seamlessly switch between one solver or the other.

APPENDIX C

Files used by Xpress NonLinear

Most of the data used by Xpress NonLinear is held in memory. However, there are a few files which are written, either automatically or on demand, in addition to those created by the Xpress Optimizer.

<i>LOGFILE</i>	Created by: XSLPsetlogfile The file name and location are user-defined.
<i>NAME.mps</i>	Created by: XSLPwriteprob This is the matrix file in extended MPS format. The name is user-defined. The extension <i>.mps</i> is appended automatically.
<i>NAME.lp</i>	Created by: XSLPwriteprob This is the matrix file in human-readable "text". The name is user-defined. The extension <i>.lp</i> is appended automatically.

APPENDIX D

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

Symbols

= column, 33

A

ABS, 434

Absolute tolerance record

Tx, 35

ARCCOS, 427

ARCSIN, 428

ARCTAN, 429

Attributes, Problem, 137

Augmentation, 57

B

BOUNDS section in file, 33

C

Callbacks and user functions, 68

Callbacks in MISLP, 80

Cascading, 45

Closure convergence tolerance, 52

Coefficients

and terms, 27

COLUMNS section in file, 32

Control parameters, 163

Convergence

closure, 52

delta, 52

extended convergence continuation, 56

impact, 53

matrix, 52

slack impact, 53

static objective (1), 54

static objective (2), 55

static objective (3), 55

user-defined, 54

Convergence criteria, 47

convex region, 480

COS, 430

Counting, 268

D

Delta convergence tolerance, 52

Derivatives

returning from user function, 70

user function, 71

Determining Row record, 33

DR record in SLPDATA, 33

E

E-12001, 445

E-12002, 445

E-12003, 445

E-12004, 445

E-12005, 445

E-12006, 445

E-12007, 446

E-12008, 446

E-12009, 446

E-12010, 446

E-12011, 446

E-12012, 446

E-12013, 446

E-12014, 446

E-12015, 446

E-12016, 446

E-12017, 446

E-12018, 446

E-12019, 447

E-12020, 447

E-12022, 447

E-12025, 447

E-12026, 447

E-12027, 447

E-12029, 447

E-12030, 447

E-12031, 448

E-12032, 448

E-12033, 448

E-12037, 448

E-12038, 448

E-12084, 448

E-12085, 448

E-12105, 448

E-12107, 448

E-12110, 448

E-12111, 448

E-12112, 448

E-12113, 449

E-12114, 449

E-12121, 449

E-12124, 449

E-12125, 449

E-12147, 449

E-12150, 449

E-12151, 449

E-12156, 449

E-12157, 449

E-12158, 449

E-12159, 449

E-12160, 449

E-12161, 449

E-12163, 450

E-12164, 450

E-12165, 450

E-12166, 450
 E-12167, 450
 E-12168, 450
 E-12169, 450
 E-12170, 450
 E-12171, 450
 E-12173, 450
 E-12174, 450
 E-12192, 450
 E-12193, 450
 E-12194, 450
 E-12195, 450
 E-12196, 451
 E-12197, 451
 E-12198, 451
 E-12199, 451
 E-12200, 451
 E-12201, 451
 E-12202, 451
 E-12203, 451
 E-12204, 451
 E-12205, 451
 E-12206, 451
 E-12207, 451
 E-12208, 451
 E-12209, 451
 E-12210, 451
 EC record in SLPDATA, 34
 Enforced Constraint record, 34
 Equals column, 33
 ERF, 435
 ERFC, 436
 Error Messages, 445
 Error vectors, penalty, 61
 EXP, 437
 Extended convergence continuation tolerance, 56
 Extended MPS file format, 31

F

files
 .ini, 24
 Files used by Xpress NonLinear, 482
 Fixing values of SLP variables in MISLP, 78
 Formula
 Initial Value record, 34
 Formulae, 31, 65
 Functions, internal, 425
 Functions, library, 268
 Functions, user, 68

G

getslack, 11
 getsol, 11

H

Handling Infeasibilities, 42
 History, 64

I

Impact convergence tolerance, 53
 Implicit variable, 32

Infeasibilities, handling, 42
 Initial Value formula, 34
 Initial Value record, 34
 Instance
 user function, 71
 Internal Functions, 425
 Iterating at each node in MISLP, 79
 IV record in SLPDATA, 34

L

Library functions, 268, 269
 LN, 438
 loadprob, 10
 LOG10, 439

M

Matrix convergence tolerance, 52
 Matrix Name Generation, 61
 Matrix Structures, 57
 MAX, 440
 maximise, 11
 MIN, 441
 minimise, 11
 MINLP, 77
 MISLP
 Callbacks, 80
 Fixing or relaxing values of SLP variables, 78
 Iterating at each node, 79
 Termination criteria at each node, 79
 Mixed Integer Non-Linear Programming, 77
 mnxnlp, 9
 mutlistart, 82

N

Name Generation, 61
 nlctr, 9
 Nonlinear objectives, 76
 Nonlinear problems, 27

O

Objectives, nonlinear, 76
 Objectives, quadratic, 76

P

Parsed formula format, 65
 Penalty error vectors, 61
 Pointer (reference) attribute, 161
 positive semi-definite matrix, 481
 Problem attributes, 137
 Problem pointer, 268
 PWL, 442

Q

Quadratic objectives, 76

R

Relative tolerance record Rx, 35
 Relaxing values of SLP variables in MISLP, 78
 Row weight
 Extended MPS record, 36
 Rx record in SLPDATA, 35

S

SB record in SLPDATA, 35
 Sequential Linear Programming, *see* Successive Linear Programming
 setinitval, 10
 SIGN, 443
 SIN, 431
 Slack impact convergence tolerance, 53
 SLP problem pointer, 268
 SLP variable, 28
 SLPDATA
 DR record, 33
 EC record, 34
 IV record, 34
 Rx record, 35
 SB record, 35
 Tx record, 35
 UF record, 36
 WT record, 36
 SLPDATA section in file, 33
 solution, 479
 Solution Process, 38
 Special Types of Problem, *see* Problem, special types
 Mixed Integer Non-Linear Programming, 77
 Nonlinear objectives, 76
 Quadratic objectives, 76
 SQRT, 444
 Static objective (1) convergence tolerance, 54
 Static objective (2) convergence tolerance, 55
 Static objective (3) convergence tolerance, 55
 Statistics, Xpress-SLP, 62
 Step Bound record, 35
 Structures, SLP matrix, 57
 Successive Linear Programming, 27

T

TAN, 432
 Termination criteria at each node in MISLP, 79
 Terms
 and coefficients, 27
 Tolerance record
 Rx, 35
 Tolerance record Tx, 35
 Tolerances, convergence, 47
 Tx record in SLPDATA, 35

U

UF record in SLPDATA, 36
 Unparsed formula format, 65
 User function, 68
 declaration in native languages, 69
 Deltas, 70
 instance, 71
 programming techniques, 70
 ReturnArray, 70
 returning derivatives, 70
 User function Derivatives, 71
 User function interface, 68
 User Function record, 36
 User Functions, 68

User-defined convergence, 54

V

Values of SLP variables in MISLP, fixing or relaxing, 78
 Variable
 implicit, 32
 SLP, 28

W

W-12023, 447
 W-12024, 447
 W-12028, 447
 W-12034, 448
 W-12142, 449
 WT record in SLPDATA, 36

X

KNITRO_PARAM_ALGORITHM, 461
 XKTR_PARAM_ALGORITHM, 461
 KNITRO_PARAM_BAR_DIRECTINTERVAL, 461
 XKTR_PARAM_BAR_DIRECTINTERVAL, 461
 KNITRO_PARAM_BAR_FEASIBLE, 461
 XKTR_PARAM_BAR_FEASIBLE, 461
 KNITRO_PARAM_BAR_FEASMODETOL, 457
 XKTR_PARAM_BAR_FEASMODETOL, 457
 KNITRO_PARAM_BAR_INITMU, 457
 XKTR_PARAM_BAR_INITMU, 457
 KNITRO_PARAM_BAR_INITPT, 462
 XKTR_PARAM_BAR_INITPT, 462
 KNITRO_PARAM_BAR_MAXBACKTRACK, 462
 XKTR_PARAM_BAR_MAXBACKTRACK, 462
 KNITRO_PARAM_BAR_MAXCROSSIT, 463
 XKTR_PARAM_BAR_MAXCROSSIT, 463
 KNITRO_PARAM_BAR_MAXREFACTOR, 463
 XKTR_PARAM_BAR_MAXREFACTOR, 463
 KNITRO_PARAM_BAR_MURULE, 463
 XKTR_PARAM_BAR_MURULE, 463
 KNITRO_PARAM_BAR_PENCONS, 464
 XKTR_PARAM_BAR_PENCONS, 464
 KNITRO_PARAM_BAR_PENRULE, 464
 XKTR_PARAM_BAR_PENRULE, 464
 KNITRO_PARAM_BAR_SWITCHRULE, 465
 XKTR_PARAM_BAR_SWITCHRULE, 465
 KNITRO_PARAM_DELTA, 457
 XKTR_PARAM_DELTA, 457
 KNITRO_PARAM_FEAStOL, 457
 XKTR_PARAM_FEAStOL, 457
 KNITRO_PARAM_FEAStOLABS, 458
 XKTR_PARAM_FEAStOLABS, 458
 KNITRO_PARAM_GRADOPT, 465
 XKTR_PARAM_GRADOPT, 465
 KNITRO_PARAM_HESSOPT, 465
 XKTR_PARAM_HESSOPT, 465
 KNITRO_PARAM_HONORBNDS, 466
 XKTR_PARAM_HONORBNDS, 466
 KNITRO_PARAM_INFEASTOL, 458
 XKTR_PARAM_INFEASTOL, 458
 KNITRO_PARAM_LMSIZE, 466
 XKTR_PARAM_LMSIZE, 466
 KNITRO_PARAM_MAXCGIT, 467

XKTR_PARAM_MAXCGIT, 467
 KNITRO_PARAM_MAXIT, 467
 XKTR_PARAM_MAXIT, 467
 KNITRO_PARAM_MIP_BRANCHRULE, 467
 XKTR_PARAM_MIP_BRANCHRULE, 467
 KNITRO_PARAM_MIP_GUB_BRANCH, 468
 XKTR_PARAM_MIP_GUB_BRANCH, 468
 KNITRO_PARAM_MIP_HEURISTIC, 468
 XKTR_PARAM_MIP_HEURISTIC, 468
 KNITRO_PARAM_MIP_HEURISTIC_MAXIT, 468
 XKTR_PARAM_MIP_HEURISTIC_MAXIT, 468
 KNITRO_PARAM_MIP_IMPLICATNS, 468
 XKTR_PARAM_MIP_IMPLICATNS, 468
 KNITRO_PARAM_INTEGERTOL, 458
 XKTR_PARAM_MIP_INTEGERTOL, 458
 KNITRO_PARAM_INTGAPABS, 458
 XKTR_PARAM_MIP_INTGAPABS, 458
 KNITRO_PARAM_MIP_INTGAPREL, 459
 XKTR_PARAM_MIP_INTGAPREL, 459
 KNITRO_PARAM_MIP_KNAPSACK, 469
 XKTR_PARAM_MIP_KNAPSACK, 469
 KNITRO_PARAM_MIP_LPALG, 469
 XKTR_PARAM_MIP_LPALG, 469
 KNITRO_PARAM_MIP_MAXNODES, 469
 XKTR_PARAM_MIP_MAXNODES, 469
 KNITRO_PARAM_MIP_MAXSOLVES, 470
 XKTR_PARAM_MIP_MAXSOLVES, 470
 KNITRO_PARAM_MIP_METHOD, 470
 XKTR_PARAM_MIP_METHOD, 470
 KNITRO_PARAM_MIP_OUTINTERVAL, 470
 XKTR_PARAM_MIP_OUTINTERVAL, 470
 KNITRO_PARAM_MIP_OUTLEVEL, 470
 XKTR_PARAM_MIP_OUTLEVEL, 470
 KNITRO_PARAM_MIP_PSEUDOINIT, 471
 XKTR_PARAM_MIP_PSEUDOINIT, 471
 KNITRO_PARAM_MIP_ROOTALG, 471
 XKTR_PARAM_MIP_ROOTALG, 471
 KNITRO_PARAM_MIP_ROUNDING, 471
 XKTR_PARAM_MIP_ROUNDING, 471
 KNITRO_PARAM_MIP_SELECTRULE, 471
 XKTR_PARAM_MIP_SELECTRULE, 471
 KNITRO_PARAM_MIP_STRONG_CANDLIM, 472
 XKTR_PARAM_MIP_STRONG_CANDLIM, 472
 KNITRO_PARAM_MIP_STRONG_LEVEL, 472
 XKTR_PARAM_MIP_STRONG_LEVEL, 472
 KNITRO_PARAM_MIP_STRONG_MAXIT, 472
 XKTR_PARAM_MIP_STRONG_MAXIT, 472
 KNITRO_PARAM_MIP_TERMINATE, 472
 XKTR_PARAM_MIP_TERMINATE, 472
 KNITRO_PARAM_OBRANGE, 459
 XKTR_PARAM_OBRANGE, 459
 KNITRO_PARAM_OPTTOL, 459
 XKTR_PARAM_OPTTOL, 459
 KNITRO_PARAM_OPTTOLABS, 459
 XKTR_PARAM_OPTTOLABS, 459
 KNITRO_PARAM_OUTLEV, 473
 XKTR_PARAM_OUTLEV, 473
 KNITRO_PARAM_PRESOLVE, 473
 XKTR_PARAM_PRESOLVE, 473
 KNITRO_PARAM_PRESOLVE_TOL, 460
 XKTR_PARAM_PRESOLVE_TOL, 460
 KNITRO_PARAM_SCALE, 473
 XKTR_PARAM_SCALE, 473
 KNITRO_PARAM_SOC, 474
 XKTR_PARAM_SOC, 474
 KNITRO_PARAM_XTOL, 460
 XKTR_PARAM_XTOL, 460
 xnlp_verbose, 11
 Xpress NonLinear problem pointer, 268
 Xpress-SLP Statistics, 62
 xprs_verbose, 11
 XPRSdestroyprob, 21
 XPRSfree, 21
 XPRSgetlpcol, 21
 XPRSgetnlpsol, 342
 XPRSmsaddcustompreset, 362
 XPRSmsaddjob, 363
 XPRSmsaddpreset, 364
 XPRSmsclear, 365
 XPRSnlpaddformulas, 279
 XPRSnlpadduserfunction, 283
 XPRSnlpcalcslacks, 287
 XPRSnlpchgformula, 299
 XPRSnlpchgformulastr, 298
 XPRSnlpchgformulastring, 297
 XPRSnlpdelformulas, 312
 XPRSnlpdeluserfunction, 314
 XPRSnlpevaluateformula, 318
 XPRSnlpgetformula, 330
 XPRSnlpgetformularows, 333
 XPRSnlpgetformulastr, 331
 XPRSnlpgetformulastring, 332
 XPRSnlpimportlibfunc, 348
 XPRSnlploadformulas, 355
 XPRSnlpoptimize, 366
 XPRSnlppostsolve, 367
 XPRSnlpprintevalinfo, 370
 XPRSnlpsetcurrentiv, 405
 XPRSnlpsetfunctionerror, 410
 XPRSnlpsetinitval, 411
 XPRSnlpvalidate, 418
 XPRSnlpvalidatekkt, 419
 XPRSnlpvalidateprob, 420
 XPRSnlpvalidaterow, 421
 XPRSnlpvalidatevector, 422
 XPRSssetcbmsjobend, 394
 XPRSssetcbmsjobstart, 395
 XPRSssetcbmswinner, 396
 XPRSssetcbnlpcofevalerror, 383
 XPRSssetcbnlpascadeend, 379
 XPRSssetcbnlpascadestart, 380
 XPRSssetcbnlpascadevar, 381
 XPRSssetcbnlpascadevarfail, 382
 XPRSssetcbnlpconstruct, 384
 XPRSssetcbnlpdrcol, 387
 XPRSssetcbnlpintsol, 388
 XPRSssetcbnlpiterend, 389
 XPRSssetcbnlpiterstart, 390
 XPRSssetcbnlpitervar, 391
 XPRSssetcbnlpupdatelinearization, 399

XPRSSlpaddcoefs, 277
 XPRSSlpcascade, 288
 XPRSSlpcascadeorder, 289
 XPRSSlpchgascadenlimit, 290
 XPRSSlpchgcoef, 292
 XPRSSlpchgcoef, 293
 XPRSSlpchgcoefstr, 291
 XPRSSlpchgdelatetype, 295
 XPRSSlpchggrowstatus, 300
 XPRSSlpchggrowwt, 301
 XPRSSlpconstruct, 306
 XPRSSlpdelcoefs, 311
 XPRSSlpevaluatecoef, 317
 XPRSSlpfixpenalties, 319
 XPRSSlpgetcccoef, 322
 XPRSSlpgetcoeffformula, 324
 XPRSSlpgetcoefs, 325
 XPRSSlpgetcoefstr, 323
 XPRSSlpgetcolinfo, 326
 XPRSSlpgetrowinfo, 339
 XPRSSlpgetrowstatus, 340
 XPRSSlpgetrowwt, 341
 XPRSSlploadcoefs, 352
 XPRSSlpinitialize, 375
 XPRSSlpsetdetrow, 409
 XPRSSlpunconstruct, 416
 XPRSSlpupdatelinearization, 417
 XPRSSwriteprtsol, 21
 SLPALGORITHM, 216
 XSLP_ALGORITHM, 216
 SLPANALYZE, 218
 XSLP_ANALYZE, 218
 SLPATOL_A, 176
 XSLP_ATOL_A, 176
 SLPATOL_R, 176
 XSLP_ATOL_R, 176
 SLPAUGMENTATION, 219
 XSLP_AUGMENTATION, 219
 SLPAUTOSAVE, 220
 XSLP_AUTOSAVE, 220
 SLPBARCROSSOVERSTART, 221
 XSLP_BARCROSSOVERSTART, 221
 SLPBARLIMIT, 221
 XSLP_BARLIMIT, 221
 SLPBARSTALLINGLIMIT, 222
 XSLP_BARSTALLINGLIMIT, 222
 SLPBARSTALLINGOBJLIMIT, 222
 XSLP_BARSTALLINGOBJLIMIT, 222
 SLPBARSTALLINGTOL, 176
 XSLP_BARSTALLINGTOL, 176
 SLPBARSTARTOPS, 222
 XSLP_BARSTARTOPS, 222
 SLPBOUNDTHRESHOLD, 177
 XSLP_BOUNDTHRESHOLD, 177
 NLPCALCTHREADS, 223
 XSLP_CALCTHREADS, 223
 SLPCASCADE, 223
 XSLP_CASCADE, 223
 SLPCASCADENLIMIT, 224
 XSLP_CASCADENLIMIT, 224
 SLPCASCADETOL_PA, 177
 XSLP_CASCADETOL_PA, 177
 SLPCASCADETOL_PR, 177
 XSLP_CASCADETOL_PR, 177
 SLPCTOL_A, 178
 XSLP_CDTOL_A, 178
 SLPCTOL_R, 178
 XSLP_CDTOL_R, 178
 SLPCLAMP SHRINK, 179
 XSLP_CLAMP SHRINK, 179
 SLPCLAMPVALIDATIONTOL_A, 179
 XSLP_CLAMPVALIDATIONTOL_A, 179
 SLPCLAMPVALIDATIONTOL_R, 179
 XSLP_CLAMPVALIDATIONTOL_R, 179
 SLP COEFFICIENTS, 146
 XSLP_COEFFICIENTS, 146
 XSLP_COL, 18
 XSLP_CON, 18
 XSLP_CONTROL, 225
 SLP CONVERGENCEOPS, 225
 XSLP_CONVERGENCEOPS, 225
 SLPCTOL, 180
 XSLP_CTOL, 180
 SLP CURRENTDELTA COST, 142
 XSLP_CURRENTDELTA COST, 142
 SLP CURRENTERROR COST, 142
 XSLP_CURRENTERROR COST, 142
 SLP CUT STRATEGY, 226
 XSLP_CUT STRATEGY, 226
 SLPDAMP, 180
 XSLP_DAMP, 180
 SLPDAMP EXPAND, 180
 XSLP_DAMP EXPAND, 180
 SLPDAMP MAX, 181
 XSLP_DAMP MAX, 181
 SLPDAMP MIN, 181
 XSLP_DAMP MIN, 181
 SLPDAMP SHRINK, 182
 XSLP_DAMP SHRINK, 182
 SLPDAMP START, 227
 XSLP_DAMP START, 227
 NLP DEFAULTTIV, 182
 XSLP_DEFAULTTIV, 182
 SLP DEFAULTSTEPBOUND, 182
 XSLP_DEFAULTSTEPBOUND, 182
 SLP DELAYUPDATEROWS, 227
 XSLP_DELAYUPDATEROWS, 227
 SLP DELTA_A, 183
 XSLP_DELTA_A, 183
 SLP DELTA_INFINITY, 183
 XSLP_DELTA_INFINITY, 183
 SLP DELTA_R, 183
 XSLP_DELTA_R, 183
 SLP DELTA_X, 184
 XSLP_DELTA_X, 184
 SLP DELTA_Z, 184
 XSLP_DELTA_Z, 184
 SLP DELTA_ZERO, 184
 XSLP_DELTA_ZERO, 184
 SLP DELTA COST, 185

XSLP_DELTACOST, 185
 SLPDELTACOSTFACTOR, 185
 XSLP_DELTACOSTFACTOR, 185
 SLPDELTAFORMAT, 261
 XSLP_DELTAFORMAT, 261
 SLPDELTAMAXCOST, 185
 XSLP_DELTAMAXCOST, 185
 SLPDELTAOFFSET, 227
 XSLP_DELTAOFFSET, 227
 SLPDELTAS, 146
 XSLP_DELTAS, 146
 SLPDELTAZLIMIT, 228
 XSLP_DELTAZLIMIT, 228
 NLPDERIVATIVES, 228
 XSLP_DERIVATIVES, 228
 NLPDETERMINISTIC, 229
 XSLP_DETERMINISTIC, 229
 SLPDJTOL, 186
 XSLP_DJTOL, 186
 SLPDRCOLDJTOL, 186
 XSLP_DRCOLDJTOL, 186
 SLPDRCOLTOL, 186
 XSLP_DRCOLTOL, 186
 SLPDRFIXRANGE, 187
 XSLP_DRFIXRANGE, 187
 SLPECFCHECK, 229
 XSLP_ECFCHECK, 229
 SLPECFCOUNT, 146
 XSLP_ECFCOUNT, 146
 SLPECFTOL_A, 187
 XSLP_ECFTOL_A, 187
 SLPECFTOL_R, 187
 XSLP_ECFTOL_R, 187
 XSLP_ECHOXPRSMESSAGES, 230
 SLPENFORCECOSTSHRINK, 188
 XSLP_ENFORCECOSTSHRINK, 188
 SLPENFORCEMAXCOST, 188
 XSLP_ENFORCEMAXCOST, 188
 XSLP_EOF, 19
 NLPEQUALSCOLUMN, 147
 XSLP_EQUALSCOLUMN, 147
 SLPERRORCOST, 189
 XSLP_ERRORCOST, 189
 SLPERRORCOSTFACTOR, 189
 XSLP_ERRORCOSTFACTOR, 189
 SLPERRORCOSTS, 142
 XSLP_ERRORCOSTS, 142
 SLPERRORMAXCOST, 189
 XSLP_ERRORMAXCOST, 189
 SLPERROROFFSET, 230
 XSLP_ERROROFFSET, 230
 SLPERRORTOL_A, 190
 XSLP_ERRORTOL_A, 190
 SLPERRORTOL_P, 190
 XSLP_ERRORTOL_P, 190
 SLPESCALATION, 190
 XSLP_ESCALATION, 190
 SLPETOL_A, 191
 XSLP_ETOL_A, 191
 SLPETOL_R, 191
 XSLP_ETOL_R, 191
 NLPEVALUATE, 230
 XSLP_EVALUATE, 230
 SLPEVTOL_A, 191
 XSLP_EVTOL_A, 191
 SLPEVTOL_R, 192
 XSLP_EVTOL_R, 192
 SLPEXPAND, 192
 XSLP_EXPAND, 192
 SLPEXPLOREDELTAS, 146
 XSLP_EXPLOREDELTAS, 146
 SLPFEASTOLTARGET, 193
 XSLP_FEASTOLTARGET, 193
 SLPFILTER, 231
 XSLP_FILTER, 231
 NLPFINDIV, 231
 XSLP_FINDIV, 231
 XSLP_FUN, 19
 NLPFUNCEVAL, 232
 XSLP_FUNCEVAL, 232
 SLPGRANULARITY, 193
 XSLP_GRANULARITY, 193
 SLPGRIDHEURSELECT, 233
 XSLP_GRIDHEURSELECT, 233
 NLPHESSIAN, 233
 XSLP_HESSIAN, 233
 SLPHEURSTRATEGY, 233
 XSLP_HEURSTRATEGY, 233
 NLPIFS, 147
 XSLP_IFS, 147
 XSLP_IFUN, 19
 NLPIMPLICITVARIABLES, 147
 XSLP_IMPLICITVARIABLES, 147
 SLPINFEASLIMIT, 234
 XSLP_INFEASLIMIT, 234
 NLPINFINITY, 193
 XSLP_INFINITY, 193
 SLPINTEGERDELTAS, 147
 XSLP_INTEGERDELTAS, 147
 SLPITER, 148
 XSLP_ITER, 148
 SLPITERFALLBACKOPS, 261
 XSLP_ITERFALLBACKOPS, 261
 SLPITERLIMIT, 234
 XSLP_ITERLIMIT, 234
 SLPITOL_A, 194
 XSLP_ITOL_A, 194
 SLPITOL_R, 194
 XSLP_ITOL_R, 194
 NLPIVNAME, 262
 XSLP_IVNAME, 262
 NLPJACOBIAN, 234
 XSLP_JACOBIAN, 234
 NLPJOBID, 148
 XSLP_JOBID, 148
 NLPKEEPBESTITER, 148
 XSLP_KEEPBESTITER, 148
 NLPKEEPEQUALSCOLUMN, 235
 XSLP_KEEPEQUALSCOLUMN, 235
 NLPLINQUADBR, 235

XSLP_LINQUADBR, 235
 NLPLOG, 236
 XSLP_LOG, 236
 xslp_log, 11
 SLPLSITERLIMIT, 236
 XSLP_LSITERLIMIT, 236
 SLPLSPATTERNLIMIT, 236
 XSLP_LSPATTERNLIMIT, 236
 SLPLSSTART, 237
 XSLP_LSSTART, 237
 SLPLSZEROLIMIT, 237
 XSLP_LSZEROLIMIT, 237
 SLPMATRIXTOL, 195
 XSLP_MATRIXTOL, 195
 NLPMAXTIME, 237
 XSLP_MAXTIME, 237
 SLPMAXWEIGHT, 195
 XSLP_MAXWEIGHT, 195
 XSLP_MEMORYFACTOR, 196
 NLPMERITLAMBDA, 196
 XSLP_MERITLAMBDA, 196
 SLPMINSBFACTOR, 197
 XSLP_MINSBFACTOR, 197
 SLPMINUSDELTAFORMAT, 262
 XSLP_MINUSDELTAFORMAT, 262
 SLPMINUSERERRORFORMAT, 262
 XSLP_MINUSERERRORFORMAT, 262
 SLPMINUSPENALTYERRORS, 148
 XSLP_MINUSPENALTYERRORS, 148
 SLPMINWEIGHT, 197
 XSLP_MINWEIGHT, 197
 SLPMIPALGORITHM, 238
 XSLP_MIPALGORITHM, 238
 SLPMIPCUTOFF_A, 197
 XSLP_MIPCUTOFF_A, 197
 SLPMIPCUTOFF_R, 198
 XSLP_MIPCUTOFF_R, 198
 SLPMIPCUTOFFCOUNT, 239
 XSLP_MIPCUTOFFCOUNT, 239
 SLPMIPCUTOFFLIMIT, 239
 XSLP_MIPCUTOFFLIMIT, 239
 SLPMIPDEFAULTALGORITHM, 240
 XSLP_MIPDEFAULTALGORITHM, 240
 SLPMIPERRORTOL_A, 198
 XSLP_MIPERRORTOL_A, 198
 SLPMIPERRORTOL_R, 199
 XSLP_MIPERRORTOL_R, 199
 SLPMIPFIXSTEPBOUNDS, 240
 XSLP_MIPFIXSTEPBOUNDS, 240
 SLPMIPITER, 149
 XSLP_MIPITER, 149
 SLPMIPITERLIMIT, 241
 XSLP_MIPITERLIMIT, 241
 SLPMIPLOG, 241
 XSLP_MIPLOG, 241
 SLPMIPNODES, 149
 XSLP_MIPNODES, 149
 SLPMIPOCOUNT, 241
 XSLP_MIPOCOUNT, 241
 SLPMIPOTOL_A, 199
 XSLP_MIPOTOL_A, 199
 SLPMIPOTOL_R, 199
 XSLP_MIPOTOL_R, 199
 XSLP_MIPPROBLEM, 161
 SLPMIPRELAXSTEPBOUNDS, 242
 XSLP_MIPRELAXSTEPBOUNDS, 242
 SLPMIPSOLS, 149
 XSLP_MIPSOLS, 149
 NLPMODELCOLS, 149
 XSLP_MODELCOLS, 149
 NLPMODELROWS, 150
 XSLP_MODELROWS, 150
 MSMAXBOUNDRANGE, 200
 XSLP_MSMAXBOUNDRANGE, 200
 XSLP_MSSTATUS, 150
 SLPMTOL_A, 200
 XSLP_MTOL_A, 200
 SLPMTOL_R, 201
 XSLP_MTOL_R, 201
 MULTISTART, 242
 XSLP_MULTISTART, 242
 MULTISTART_LOG, 243
 XSLP_MULTISTART_LOG, 243
 MULTISTART_MAXSOLVES, 243
 XSLP_MULTISTART_MAXSOLVES, 243
 MULTISTART_MAXTIME, 243
 XSLP_MULTISTART_MAXTIME, 243
 MULTISTART_POOLSIZE, 244
 XSLP_MULTISTART_POOLSIZE, 244
 MULTISTART_SEED, 244
 XSLP_MULTISTART_SEED, 244
 MULTISTART_THREADS, 245
 XSLP_MULTISTART_THREADS, 245
 SLPMTOL, 202
 XSLP_MVTOL, 202
 NLPSOLVER, 244
 XSLP_NLPSOLVER, 244
 NLPSTATUS, 150
 XSLP_NLPSTATUS, 150
 SLPNONCONSTANTCOEFFS, 151
 XSLP_NONCONSTANTCOEFFS, 151
 NONLINEARCONSTRAINTS, 151
 XSLP_NONLINEARCONSTRAINTS, 151
 XSLP_OBJSENSE, 202
 SLPOBJTHRESHOLD, 202
 XSLP_OBJTHRESHOLD, 202
 SLPOBJTOPENALTYCOST, 203
 XSLP_OBJTOPENALTYCOST, 203
 NLPOBJVAL, 143
 XSLP_OBJVAL, 143
 SLPOCOUNT, 245
 XSLP_OCOUNT, 245
 XSLP_OP, 19
 SLPOPTIMALITYTOLTARGET, 203
 XSLP_OPTIMALITYTOLTARGET, 203
 NLP OPTTIME, 142
 XSLP_OPTTIME, 142
 NLPORIGINALCOLS, 151
 XSLP_ORIGINALCOLS, 151
 NLPORIGINALROWS, 151

XSLP_ORIGINALROWS, 151
 SLPOTOL_A, 204
 XSLP_OTOL_A, 204
 SLPOTOL_R, 204
 XSLP_OTOL_R, 204
 SLPPENALTYCOLFORMAT, 263
 XSLP_PENALTYCOLFORMAT, 263
 SLPPENALTYDELTACOLUMN, 152
 XSLP_PENALTYDELTACOLUMN, 152
 SLPPENALTYDELTAROW, 152
 XSLP_PENALTYDELTAROW, 152
 SLPPENALTYDELTAS, 152
 XSLP_PENALTYDELTAS, 152
 SLPPENALTYDELTATOTAL, 143
 XSLP_PENALTYDELTATOTAL, 143
 SLPPENALTYDELTAVALUE, 143
 XSLP_PENALTYDELTAVALUE, 143
 SLPPENALTYERRORCOLUMN, 152
 XSLP_PENALTYERRORCOLUMN, 152
 SLPPENALTYERRORROW, 153
 XSLP_PENALTYERRORROW, 153
 SLPPENALTYERRORS, 153
 XSLP_PENALTYERRORS, 153
 SLPPENALTYERRORTOTAL, 143
 XSLP_PENALTYERRORTOTAL, 143
 SLPPENALTYERRORVALUE, 143
 XSLP_PENALTYERRORVALUE, 143
 SLPPENALTYINFOSTART, 246
 XSLP_PENALTYINFOSTART, 246
 SLPPENALTYROWFORMAT, 263
 XSLP_PENALTYROWFORMAT, 263
 SLPPPLUSDELTAFORMAT, 263
 XSLP_PLUSDELTAFORMAT, 263
 SLPPPLUSERRORFORMAT, 264
 XSLP_PLUSERRORFORMAT, 264
 SLPPPLUSPENALTYERRORS, 153
 XSLP_PLUSPENALTYERRORS, 153
 NLPPPOSTSOLVE, 246
 XSLP_POSTSOLVE, 246
 NLPPRESOLVE, 246
 XSLP_PRESOLVE, 246
 NLPPRESOLVE_ELIMTOL, 205
 XSLP_PRESOLVE_ELIMTOL, 205
 NLPPRESOLVEELIMINATIONS, 153
 XSLP_PRESOLVEELIMINATIONS, 153
 NLPPRESOLVELEVEL, 247
 XSLP_PRESOLVELEVEL, 247
 NLPPRESOLVEOPS, 247
 XSLP_PRESOLVEOPS, 247
 XSLP_PRESOLVESTATE, 154
 NLPPRESOLVEZERO, 205
 XSLP_PRESOLVEZERO, 205
 NLPPRIMALINTEGRAL, 144
 XSLP_PRIMALINTEGRAL, 144
 NLPPRIMALINTEGRALALPHA, 205
 XSLP_PRIMALINTEGRALALPHA, 205
 NLPPRIMALINTEGRALREF, 206
 XSLP_PRIMALINTEGRALREF, 206
 NLPPROBING, 248
 XSLP_PROBING, 248
 NLPREFORMULATE, 248
 XSLP_REFORMULATE, 248
 SLPSAMECOUNT, 249
 XSLP_SAMECOUNT, 249
 SLPSAMEDAMP, 250
 XSLP_SAMEDAMP, 250
 SLPSBLOROWFORMAT, 264
 XSLP_SBLOROWFORMAT, 264
 SLPSBNAME, 265
 XSLP_SBNAME, 265
 SLPSBROWOFFSET, 250
 XSLP_SBROWOFFSET, 250
 SLPSBSTART, 250
 XSLP_SBSTART, 250
 SLPSBUPROWFORMAT, 265
 XSLP_SBUPROWFORMAT, 265
 SLPSBXCONVERGED, 154
 XSLP_SBXCONVERGED, 154
 XSLP_SCALE, 251
 XSLP_SCALECOUNT, 251
 SLPSEMICONTDELTAS, 154
 XSLP_SEMICONTDELTAS, 154
 SLP SHRINK, 206
 XSLP_SHRINK, 206
 SLP SHRINKBIAS, 206
 XSLP_SHRINKBIAS, 206
 XSLP_SLPLLOG, 252
 xslp_slplog, 11
 NLP SOLSTATUS, 155
 XSLP_SOLSTATUS, 155
 LOCALSOLVER, 252
 XSLP_SOLVER, 252
 LOCALSOLVERSELECTED, 154
 XSLP_SOLVERSELECTED, 154
 SLPSTATUS, 155
 XSLP_STATUS, 155
 SLPSTOL_A, 207
 XSLP_STOL_A, 207
 SLPSTOL_R, 207
 XSLP_STOL_R, 207
 NLPSTOPOUTOFRANGE, 252
 XSLP_STOPOUTOFRANGE, 252
 NLPSTOPSTATUS, 157
 XSLP_STOPSTATUS, 157
 NLP THREADS, 253
 XSLP_THREADS, 253
 NLP THREADSAFEUSERFUNC, 253
 XSLP_THREADSAFEUSERFUNC, 253
 SLPTOLNAME, 265
 XSLP_TOLNAME, 265
 SLPTOLSETS, 157
 XSLP_TOLSETS, 157
 NLPTOTALEVALUATIONERRORS, 157
 XSLP_TOTALEVALUATIONERRORS, 157
 SLPTRACEMASK, 266
 XSLP_TRACEMASK, 266
 SLPTRACEMASKOPS, 253
 XSLP_TRACEMASKOPS, 253
 SLPUCCONSTRAINEDCOUNT, 158
 XSLP_UCCONSTRAINEDCOUNT, 158

XSLP_UFINSTANCES, 158
 NLPUFS, 158
 XSLP_UFS, 158
 SLPUNCONVERGED, 158
 XSLP_UNCONVERGED, 158
 SLPUNFINISHEDLIMIT, 254
 XSLP_UNFINISHEDLIMIT, 254
 SLPUPDATEFORMAT, 266
 XSLP_UPDATEFORMAT, 266
 SLPUPDATEOFFSET, 255
 XSLP_UPDATEOFFSET, 255
 NLPUSEDERIVATIVES, 159
 XSLP_USEDERIVATIVES, 159
 NLPUSERFUNCCALLS, 159
 XSLP_USERFUNCCALLS, 159
 NLPVALIDATIONFACTOR, 208
 XSLP_VALIDATIONFACTOR, 208
 NLPVALIDATIONINDEX_A, 144
 XSLP_VALIDATIONINDEX_A, 144
 NLPVALIDATIONINDEX_K, 144
 XSLP_VALIDATIONINDEX_K, 144
 NLPVALIDATIONINDEX_R, 144
 XSLP_VALIDATIONINDEX_R, 144
 NLPVALIDATIONNETOBJ, 144
 XSLP_VALIDATIONNETOBJ, 144
 NLPVALIDATIONSTATUS, 159
 XSLP_VALIDATIONSTATUS, 159
 NLPVALIDATIONTARGET_K, 208
 XSLP_VALIDATIONTARGET_K, 208
 NLPVALIDATIONTARGET_R, 208
 XSLP_VALIDATIONTARGET_R, 208
 NLPVALIDATIONTOL_A, 209
 XSLP_VALIDATIONTOL_A, 209
 NLPVALIDATIONTOL_K, 209
 XSLP_VALIDATIONTOL_K, 209
 NLPVALIDATIONTOL_R, 209
 XSLP_VALIDATIONTOL_R, 209
 NLPVARIABLES, 159
 XSLP_VARIABLES, 159
 SLPVCOUNT, 255
 XSLP_VCOUNT, 255
 XSLP_VERSIONDATE, 162
 SLPVLIMIT, 256
 XSLP_VLIMIT, 256
 XSLP_VSOLINDEX, 145
 SLPVTOL_A, 210
 XSLP_VTOL_A, 210
 SLPVTOL_R, 211
 XSLP_VTOL_R, 211
 SLPWCOUNT, 256
 XSLP_WCOUNT, 256
 SLPWTOL_A, 211
 XSLP_WTOL_A, 211
 SLPWTOL_R, 212
 XSLP_WTOL_R, 212
 SLPXCOUNT, 257
 XSLP_XCOUNT, 257
 SLPXLIMIT, 258
 XSLP_XLIMIT, 258
 XSLP_XPRSPROBLEM, 161
 XSLP_XSLPPROBLEM, 161
 SLPXTOL_A, 213
 XSLP_XTOL_A, 213
 SLPXTOL_R, 214
 XSLP_XTOL_R, 214
 NLPZERO, 215
 XSLP_ZERO, 215
 SLPZEROCRITERION, 259
 XSLP_ZEROCRITERION, 259
 SLPZEROCRITERIONCOUNT, 259
 XSLP_ZEROCRITERIONCOUNT, 259
 SLPZEROCRITERIONSTART, 260
 XSLP_ZEROCRITERIONSTART, 260
 SLPZERORESRESET, 160
 XSLP_ZERORESRESET, 160
 SLPZERORESRETAINED, 160
 XSLP_ZERORESRETAINED, 160
 SLPZERORESTOTAL, 160
 XSLP_ZERORESTOTAL, 160
 XSLP_ATOL, 52
 XSLP_CTOL, 52
 XSLP_ITOL, 53
 XSLP_MTOL, 53
 XSLP_OCOUNT, 55
 XSLP_OTOL, 55
 XSLP_STOL, 54
 XSLP_VCOUNT, 54
 XSLP_VLIMIT, 55
 XSLP_VTOL, 55
 XSLP_WCOUNT, 56
 XSLP_WTOL, 56
 XSLP_XCOUNT, 56
 XSLP_XLIMIT, 56
 XSLP_XTOL, 56
 XSLPaddcoefs, 277
 XSLPadddfs, 281
 XSLPaddformulas, 279
 XSLPaddtolsets, 282
 XSLPadduserfunction, 283
 XSLPaddvars, 285
 XSLPcalcslack, 287
 XSLPcascade, 288
 XSLPcascadeorder, 289
 XSLPchgcascadenlimit, 290
 XSLPchgccoef, 292
 XSLPchgcoef, 293
 XSLPchgdelatetype, 295
 XSLPchgdf, 296
 XSLPchgformula, 19, 299
 XSLPchgformulastring, 297
 XSLPchgrowstatus, 300
 XSLPchgrowwt, 301
 XSLPchgtolset, 302
 XSLPchgvar, 304
 XSLPconstruct, 306
 XSLPcopycallbacks, 307
 XSLPcopycontrols, 308
 XSLPcopyprob, 309
 XSLPcreateprob, 310
 XSLPdelcoefs, 311

XSLPdelformulas, 312
 XSLPdeltolsets, 313
 XSLPdeluserfunction, 314
 XSLPdelvars, 315
 XSLPdestroyprob, 21, 316
 XSLPevaluatecoef, 317
 XSLPevaluateformula, 318
 XSLPfixpenalties, 319
 XSLPfree, 21, 320
 XSLPgetbanner, 321
 XSLPgetccoef, 322
 XSLPgetcoeffformula, 324
 XSLPgetcoefs, 325
 XSLPgetcolinfo, 326
 XSLPgetdblattrib, 327
 XSLPgetdblcontrol, 328
 XSLPgetdf, 329
 XSLPgetformula, 330
 XSLPgetformularows, 333
 XSLPgetformulastring, 332
 XSLPgetindex, 334
 XSLPgetintattrib, 335
 XSLPgetintcontrol, 336
 XSLPgetlasterror, 337
 XSLPgetptrattrib, 338
 XSLPgetrowinfo, 339
 XSLPgetrowstatus, 340
 XSLPgetrowwt, 341
 XSLPgetslpsol, 342
 XSLPgetstrattrib, 343
 XSLPgetstrcontrol, 344
 XSLPgettolset, 345
 XSLPgetvar, 21, 346
 XSLPimportlibfunc, 348
 XSLPinit, 349
 XSLPinterrupt, 350
 XSLPitemname, 351
 XSLPloadcoefs, 352
 XSLPloaddfs, 354
 XSLPloadformulas, 19, 355
 XSLPloadtolsets, 357
 XSLPloadvars, 358
 XSLPmaxim, 21, 360
 XSLPminim, 21, 361
 XSLPmsaddcustompreset, 362
 XSLPmsaddjob, 363
 XSLPmsaddpreset, 364
 XSLPmsclear, 365
 XSLPnlpoptimize, 366
 XSLPpostsolve, 367
 XSLPpresolve, 368
 XSLPprintevalinfo, 370
 XSLPprintmemory, 369
 XSLPprob, 268
 XSLPreadprob, 371
 XSLPreinitialize, 375
 XSLPmaxim, 372
 XSLPreminim, 373
 XSLPrestore, 374
 XSLPsave, 376
 XSLPsaveas, 377
 XSLPscaling, 378
 XSLPsetcbcascadeend, 379
 XSLPsetcbcascadestart, 380
 XSLPsetcbcascadevar, 381
 XSLPsetcbcascadevarfail, 382
 XSLPsetcbccoefevalerror, 383
 XSLPsetcbconstruct, 384
 XSLPsetcbdestroy, 386
 XSLPsetcbdrcol, 387
 XSLPsetcbintsol, 388
 XSLPsetcbiterend, 389
 XSLPsetcbiterstart, 390
 XSLPsetcbitervar, 391
 XSLPsetcbmessage, 14, 392
 XSLPsetcbmsjobend, 394
 XSLPsetcbmsjobstart, 395
 XSLPsetcbmswinner, 396
 XSLPsetcboptnode, 397
 XSLPsetcbprenode, 398
 XSLPsetcbpresolved, 401
 XSLPsetcbpreupdatelinearization, 399
 XSLPsetcbslpend, 402
 XSLPsetcbslpnode, 403
 XSLPsetcbslpstart, 404
 XSLPsetcurrentiv, 405
 XSLPsetdblcontrol, 406
 XSLPsetdefaultcontrol, 407
 XSLPsetdefaults, 408
 XSLPsetdetrow, 409
 XSLPsetfunctionerror, 410
 XSLPsetinitval, 411
 XSLPsetintcontrol, 412
 XSLPsetlogfile, 413
 XSLPsetparam, 414
 XSLPsetstrcontrol, 415
 XSLPunconstruct, 416
 XSLPupdatelinearization, 417
 XSLPvalidate, 418
 XSLPvalidatekkt, 419
 XSLPvalidateprob, 420
 XSLPvalidaterow, 421
 XSLPvalidatevector, 422
 XSLPwriteprob, 21, 423
 XSLPwriteslxsol, 424