

Xpress-SLP

Program Reference Manual

Release version 1.41

Last update 4 June, 2008

Published by Fair Isaac Corporation

©Copyright Fair Isaac Corporation 2009. All rights reserved.

All trademarks referenced in this manual that are not the property of Fair Isaac are acknowledged.

All companies, products, names and data contained within this book are completely fictitious and are used solely to illustrate the use of Xpress. Any similarity between these names or data and reality is purely coincidental.

How to Contact the Xpress Team

Information, Sales and Licensing

USA, CANADA AND ALL AMERICAS

Email: XpressSalesUS@fico.com

WORLDWIDE

Email: XpressSalesUK@fico.com

Tel: +44 1926 315862

Fax: +44 1926 315854

FICO, Xpress team
Leam House, 64 Trinity Street
Leamington Spa
Warwickshire CV32 5YN
UK

Product Support

Email: Support@fico.com

(Please include 'Xpress' in the subject line)

Telephone:

NORTH AMERICA

Tel (toll free): +1 (877) 4FI-SUPP

Fax: +1 (402) 496-2224

EUROPE, MIDDLE EAST, AFRICA

Tel: +44 (0) 870-420-3777

UK (toll free): 0800-0152-153

South Africa (toll free): 0800-996-153

Fax: +44 (0) 870-420-3778

ASIA-PACIFIC, LATIN AMERICA, CARIBBEAN

Tel: +1 (415) 446-6185

Brazil (toll free): 0800-891-6146

For the latest news and Xpress software and documentation updates, please visit the Xpress website at <http://www.fico.com/xpress> or subscribe to our mailing list.

Contents

1	Nonlinear Problems	1
1.1	Coefficients and terms	1
1.2	SLP variables	2
2	Extended MPS file format	3
2.1	Formulae	3
2.2	COLUMNS	4
2.3	BOUNDS	5
2.4	SLPDATA	5
2.4.1	CV (Character variable)	5
2.4.2	DC (Delayed constraint)	6
2.4.3	DR (Determining row)	6
2.4.4	EC (Enforced constraint)	7
2.4.5	FR (Free variable)	7
2.4.6	FX (Fixed variable)	7
2.4.7	IV (Initial value)	7
2.4.8	LO (Lower bounded variable)	8
2.4.9	Rx, Tx (Relative and absolute convergence tolerances)	8
2.4.10	SB (Initial step bound)	9
2.4.11	UF (User function)	9
2.4.12	UP (Free variable)	9
2.4.13	WT (Explicit row weight)	9
2.4.14	XV (Extended variable array)	10
3	Xpress-SLP Solution Process	12
4	Handling Infeasibilities	14
5	Cascading	15
6	Convergence criteria	17
6.1	Convergence criteria	17
6.1.1	Closure tolerance (CTOL)	19
6.1.2	Delta tolerance (ATOL)	20
6.1.3	Matrix tolerance (MTOL)	20
6.1.4	Impact tolerance (ITOL)	21
6.1.5	Slack impact tolerance (STOL)	21
6.1.6	User-defined convergence	22
6.1.7	Static objective function (1) tolerance (VTOL)	22
6.1.8	Static objective function (2) tolerance (OTOL)	22
6.1.9	Static objective function (3) tolerance (XTOL)	23
6.1.10	Extended convergence continuation tolerance (WTOL)	24

7	Xpress-SLP Structures	25
7.1	SLP Matrix Structures	25
7.1.1	Augmentation of a nonlinear coefficient	26
7.1.2	Augmentation of a nonlinear term	27
7.1.3	Augmentation of a user-defined SLP variable	28
7.1.4	SLP penalty error vectors	29
7.2	Xpress-SLP Matrix Name Generation	30
7.3	Xpress-SLP Statistics	31
7.4	SLP Variable History	33
8	Problem Attributes	34
8.1	Double problem attributes	37
	XSLP_CURRENTDELTACOST	37
	XSLP_CURRENTERRORCOST	37
	XSLP_ERRORCOSTS	37
	XSLP_OBJSENSE	37
	XSLP_OBJVAL	38
	XSLP_PENALTYDELTATOTAL	38
	XSLP_PENALTYDELTAVALUE	38
	XSLP_PENALTYERRORTOTAL	38
	XSLP_PENALTYERRORVALUE	38
	XSLP_VALIDATIONINDEX_A	39
	XSLP_VALIDATIONINDEX_R	39
	XSLP_VSOLINDEX	39
8.2	Integer problem attributes	40
	XSLP_COEFFICIENTS	40
	XSLP_CVS	40
	XSLP_DELTAS	40
	XSLP_ECFCOUNT	40
	XSLP_EQUALS COLUMN	40
	XSLP_IFS	41
	XSLP_IMPLICITVARIABLES	41
	XSLP_INTERNALFUNCCALLS	41
	XSLP_ITER	41
	XSLP_MINORVERSION	41
	XSLP_MINUSPENALTYERRORS	42
	XSLP_MIPITER	42
	XSLP_MIPNODES	42
	XSLP_NONLINEARCONSTRAINTS	42
	XSLP_PENALTYDELTACOLUMN	42
	XSLP_PENALTYDELTAROW	43
	XSLP_PENALTYDELTAS	43
	XSLP_PENALTYERRORCOLUMN	43
	XSLP_PENALTYERRORROW	43
	XSLP_PENALTYERRORS	43
	XSLP_PLUSPENALTYERRORS	44
	XSLP_PREOLVEDELETEDDELTA	44
	XSLP_PREOLVEFIXEDCOEF	44
	XSLP_PREOLVEFIXEDDR	44
	XSLP_PREOLVEFIXEDNZCOL	45
	XSLP_PREOLVEFIXEDSLPVAR	45
	XSLP_PREOLVEFIXEDZCOL	45
	XSLP_PREOLVEPASSES	45
	XSLP_PREOLVETIGHTENED	46
	XSLP_SBXCONVERGED	46

XSLP_STATUS	46
XSLP_TOLSETS	47
XSLP_UCCONSTRAINEDCOUNT	47
XSLP_UFINSTANCES	47
XSLP_UFS	47
XSLP_UNCONVERGED	47
XSLP_USEDERIVATIVES	48
XSLP_USERFUNCCALLS	48
XSLP_VARIABLES	48
XSLP_VERSION	48
XSLP_XVS	48
XSLP_ZEROESRESET	49
XSLP_ZEROESRETAINED	49
XSLP_ZEROESTOTAL	49
8.3 Reference (pointer) problem attributes	50
XSLP_MIPPROBLEM	50
XSLP_XPRSPROBLEM	50
XSLP_XSLPPROBLEM	50
XSLP_GLOBALFUNCOBJECT	50
8.4 String problem attributes	51
XSLP_UNIQUEPREFIX	51
XSLP_VERSIONDATE	51
9 Control Parameters	52
9.1 Double control parameters	59
XSLP_ATOL_A	59
XSLP_ATOL_R	59
XSLP_CASCADETOL_PA	59
XSLP_CASCADETOL_PR	60
XSLP_CDTOL_A	60
XSLP_CDTOL_R	60
XSLP_CTOL	61
XSLP_DAMP	61
XSLP_DAMPEXPAND	61
XSLP_DAMPMAX	62
XSLP_DAMPMIN	62
XSLP_DAMPSHRINK	63
XSLP_DEFAULTIV	63
XSLP_DEFAULTSTEPBOUND	63
XSLP_DELTA_A	64
XSLP_DELTA_R	64
XSLP_DELTA_X	64
XSLP_DELTA_Z	65
XSLP_DELTACOST	65
XSLP_DELTACOSTFACTOR	65
XSLP_DELTAMAXCOST	66
XSLP_DJTOL	66
XSLP_ECFTOL_A	66
XSLP_ECFTOL_R	67
XSLP_EQTOL_A	67
XSLP_EQTOL_R	67
XSLP_ERRORCOST	68
XSLP_ERRORCOSTFACTOR	68
XSLP_ERRORMAXCOST	68
XSLP_ERRORTOL_A	69

XSLP_ERRORTOL_P	69
XSLP_ESCALATION	69
XSLP_ETOL_A	70
XSLP_ETOL_R	70
XSLP_EVTOL_A	70
XSLP_EVTOL_R	71
XSLP_EXPAND	71
XSLP_GRANULARITY	71
XSLP_INFINITY	72
XSLP_ITOL_A	72
XSLP_ITOL_R	73
XSLP_MAXWEIGHT	73
XSLP_MEMORYFACTOR	74
XSLP_MINWEIGHT	74
XSLP_MIPCUTOFF_A	74
XSLP_MIPCUTOFF_R	75
XSLP_MIPERRORTOL_A	75
XSLP_MIPERRORTOL_R	76
XSLP_MIPOTOL_A	76
XSLP_MIPOTOL_R	76
XSLP_MTOL_A	77
XSLP_MTOL_R	77
XSLP_MVTOL	78
XSLP_OBJSENSE	79
XSLP_OBJTOPENALTYCOST	79
XSLP_OTOL_A	80
XSLP_OTOL_R	80
XSLP_PRESOLVEZERO	81
XSLP_SHRINK	81
XSLP_STOL_A	81
XSLP_STOL_R	82
XSLP_VALIDATIONTOL_A	82
XSLP_VALIDATIONTOL_R	83
XSLP_VTOL_A	83
XSLP_VTOL_R	84
XSLP_WTOL_A	84
XSLP_WTOL_R	85
XSLP_XTOL_A	86
XSLP_XTOL_R	87
XSLP_ZERO	87
9.2 Integer control parameters	89
XSLP_ALGORITHM	89
XSLP_AUGMENTATION	90
XSLP_AUTOSAVE	92
XSLP_BARLIMIT	92
XSLP_CASCADE	92
XSLP_CASCADENLIMIT	93
XSLP_CONTROL	93
XSLP_DAMPSTART	94
XSLP_DCLIMIT	94
XSLP_DCLOG	95
XSLP_DELAYUPDATEROWS	95
XSLP_DECOMPOSE	95
XSLP_DECOMPOSEPASSLIMIT	96
XSLP_DELTAOFFSET	96

XSLP_DELTAZLIMIT	97
XSLP_DERIVATIVES	97
XSLP_ECFCHECK	98
XSLP_ERROROFFSET	98
XSLP_EVALUATE	99
XSLP_EXCELVISIBLE	99
XSLP_EXTRACVS	99
XSLP_EXTRAUFS	100
XSLP_EXTRAXVITEMS	100
XSLP_EXTRAXVS	101
XSLP_FUNCEVAL	101
XSLP_INFEASLIMIT	102
XSLP_ITERLIMIT	102
XSLP_LOG	102
XSLP_MAXTIME	103
XSLP_MIPALGORITHM	103
XSLP_MIPCUTOFFCOUNT	104
XSLP_MIPCUTOFFLIMIT	104
XSLP_MIPDEFAULTALGORITHM	105
XSLP_MIPFIXSTEPBOUNDS	105
XSLP_MIPITERLIMIT	106
XSLP_MIPLOG	106
XSLP_MIPOCOUNT	106
XSLP_MIPRELAXSTEPBOUNDS	107
XSLP_OCOUNT	107
XSLP_PENALTYINFOSTART	108
XSLP_PRESOLVE	108
XSLP_PRESOLVEPASSLIMIT	108
XSLP_SAMECOUNT	109
XSLP_SAMEDAMP	109
XSLP_SBROWOFFSET	110
XSLP_SBSTART	110
XSLP_SCALE	110
XSLP_SCALECOUNT	111
XSLP_SLPLOG	111
XSLP_STOPOUTOFRANGE	112
XSLP_TIMEPRINT	112
XSLP_UNFINISHEDLIMIT	112
XSLP_UPDATEOFFSET	113
XSLP_VCOUNT	113
XSLP_VLIMIT	114
XSLP_WCOUNT	114
XSLP_XCOUNT	115
XSLP_XLIMIT	116
XSLP_ZEROCRITERION	117
XSLP_ZEROCRITERIONCOUNT	117
XSLP_ZEROCRITERIONSTART	118
9.3 Memory control parameters	119
XSLP_MEM_CALCSTACK	119
XSLP_MEM_COEF	120
XSLP_MEM_COL	120
XSLP_MEM_CVAR	120
XSLP_MEM_DERIVATIVES	120
XSLP_MEM_EXCELDDOUBLE	120
XSLP_MEM_FORMULA	120

XSLP_MEM_FORMULAHASH	121
XSLP_MEM_FORMULAVALUE	121
XSLP_MEM_ITERLOG	121
XSLP_MEM_RETURNARRAY	121
XSLP_MEM_ROW	121
XSLP_MEM_STACK	121
XSLP_MEM_STRING	122
XSLP_MEM_TOL	122
XSLP_MEM_UF	122
XSLP_MEM_VALSTACK	122
XSLP_MEM_VAR	122
XSLP_MEM_XF	122
XSLP_MEM_XFNAMES	123
XSLP_MEM_XFVALUE	123
XSLP_MEM_XROW	123
XSLP_MEM_XV	123
XSLP_MEM_XVITEM	123
9.4 String control parameters	124
XSLP_CVNAME	124
XSLP_DELTAFORMAT	124
XSLP_IVNAME	124
XSLP_MINUSDELTAFORMAT	125
XSLP_MINUSERERRORFORMAT	125
XSLP_PENALTYCOLFORMAT	125
XSLP_PENALTYROWFORMAT	126
XSLP_PLUSDELTAFORMAT	126
XSLP_PLUSERERRORFORMAT	127
XSLP_SBLOROWFORMAT	127
XSLP_SBNAME	127
XSLP_SBUPROWFORMAT	128
XSLP_TOLNAME	128
XSLP_UPDATEFORMAT	128
10 Library functions and the programming interface	130
10.1 Counting	130
10.2 The Xpress-SLP problem pointer	131
10.3 The XSLPload... functions	132
10.4 Library functions	132
XSLPaddcoefs	139
XSLPaddcvars	141
XSLPadddcs	142
XSLPadddfs	144
XSLPadddivfs	145
XSLPaddnames	147
XSLPaddtolsets	148
XSLPadduserfuncs	149
XSLPaddvars	151
XSLPaddxvs	153
XSLPcalluserfunc	155
XSLPcascade	156
XSLPcascadeorder	157
XSLPchgcoef	158
XSLPchgcoef	159
XSLPchgcvar	161
XSLPchgdc	162

XSLPchgdf	164
XSLPchgfuncobject	165
XSLPchgivf	166
XSLPchgrow	167
XSLPchgrowwt	168
XSLPchgtolset	169
XSLPchguserfunc	170
XSLPchguserfuncaddress	172
XSLPchguserfuncobject	173
XSLPchgvar	174
XSLPchg xv	176
XSLPchg xvitem	177
XSLPconstruct	179
XSLPcopycallbacks	180
XSLPcopycontrols	181
XSLPcopyprob	182
XSLPcreateprob	183
XSLPdecompose	184
XSLPdestroyprob	185
XSLPevaluatecoef	186
XSLPevaluateformula	187
XSLPformatvalue	188
XSLPfree	189
XSLPgetbanner	190
XSLPgetccoef	191
XSLPgetcoef	192
XSLPgetcvar	193
XSLPgetdblattrib	194
XSLPgetdblcontrol	195
XSLPgetdc	196
XSLPgetdf	197
XSLPgetdtime	198
XSLPgetfuncinfo	199
XSLPgetfuncinfoV	200
XSLPgetfuncobject	201
XSLPgetfuncobjectV	202
XSLPgetindex	203
XSLPgetintattrib	204
XSLPgetintcontrol	205
XSLPgetivf	206
XSLPgetlasterror	207
XSLPgetmessagetype	208
XSLPgetnames	209
XSLPgetparam	210
XSLPgetptrattrib	211
XSLPgetrow	212
XSLPgetrowwt	213
XSLPgetslpsol	214
XSLPgetstrattrib	215
XSLPgetstrcontrol	216
XSLPgetstring	217
XSLPgettime	218
XSLPgettolset	219
XSLPgetuserfunc	220
XSLPgetuserfuncaddress	222

XSLPgetuserfuncobject	223
XSLPgetvar	224
XSLPgetversion	226
XSLPgetxv	227
XSLPgetxvitem	228
XSLPglobal	230
XSLPinit	231
XSLPitemname	232
XSLPloadcoefs	233
XSLPloadcvars	235
XSLPloaddc	236
XSLPloaddfs	238
XSLPloaddivfs	239
XSLPloadtolsets	241
XSLPloaduserfuncs	242
XSLPloadvars	244
XSLPloadxvs	246
XSLPmaxim	248
XSLPminim	249
XSLPopt	250
XSLPparsecformula	251
XSLPparseformula	252
XSLPpreparseformula	253
XSLPpresolve	254
XSLPprintmsg	255
XSLPqparse	256
XSLPreadprob	257
XSLPmaxim	258
XSLPminim	259
XSLPstore	260
XSLPrevis	261
XSLProwinfo	262
XSLPsave	263
XSLPsaveas	264
XSLPscaling	265
XSLPsetbcascadeend	266
XSLPsetbcascadestart	267
XSLPsetbcascadevar	268
XSLPsetbcascadevarF	270
XSLPsetcbconstruct	272
XSLPsetcbdestroy	274
XSLPsetcbformula	275
XSLPsetcbintsol	277
XSLPsetcbiterend	278
XSLPsetcbiterstart	279
XSLPsetcbitervar	280
XSLPsetcbitervarF	282
XSLPsetcbmessage	284
XSLPsetcbmessageF	286
XSLPsetcboptnode	288
XSLPsetcbprenode	289
XSLPsetcbslpend	291
XSLPsetcbslpnode	292
XSLPsetcbslpstart	293
XSLPsetdblcontrol	294

XSLPsetdefaultcontrol	295
XSLPsetdefaults	296
XSLPsetfuncobject	297
XSLPsetfunctionerror	298
XSLPsetintcontrol	299
XSLPsetlogfile	300
XSLPsetparam	301
XSLPsetstrcontrol	302
XSLPsetstring	303
XSLPsetuniqueprefix	304
XSLPsetuserfuncaddress	305
XSLPsetuserfuncinfo	306
XSLPsetuserfuncobject	307
XSLPtime	308
XSLPtokencount	309
XSLPtoVBString	310
XSLPuprintmemory	311
XSLPuserfuncinfo	312
XSLPvalidformula	313
XSLPvalidate	315
XSLPwriteprob	316
11 Internal Functions	317
11.1 Trigonometric functions	319
ARCCOS	320
ARCSIN	321
ARCTAN	322
COS	323
SIN	324
TAN	325
11.2 Other mathematical functions	326
ABS	327
EXP	328
LN	329
LOG, LOG10	330
MAX	331
MIN	332
SQRT	333
11.3 Logical functions	334
EQ	335
GE	336
GT	337
IF	338
LE	339
LT	340
NE	341
NOT	342
11.4 Problem-related functions	343
ACT	344
DJ	345
LO	346
MATRIX	347
MV	348
PARAM	349
RHS	350

RHSRANGE	351
SLACK	352
UP	353
11.5 Specialized functions	354
IAC	355
INTERP	356
12 Xpress-SLP Formulae	357
12.1 Parsed and unparsed formulae	357
12.2 Example of an arithmetic formula	359
12.3 Example of a formula involving a simple function	360
12.4 Example of a formula involving a complicated function	361
12.5 Example of a formula defining a user function	362
12.6 Example of a formula defining an XV	362
12.7 Example of a formula defining a DC	363
12.8 Formula evaluation and derivatives	363
13 User Functions	364
13.1 Constant Derivatives	364
13.2 Callbacks and user functions	365
13.3 User function interface	365
13.4 Function Declaration in Xpress-SLP	366
13.4.1 Function declaration in Extended MPS format	366
13.4.2 Function declaration through XSLPloaduserfuncs and XSLPadduserfuncs	369
13.4.3 Function declaration through XSLPchguserfunc	371
13.4.4 Function declaration through SLPDATA in Mosel	371
13.5 User Function declaration in native languages	372
13.5.1 User function declaration in C	372
13.5.2 User function declaration in Fortran	373
13.5.3 User function declaration in Excel (spreadsheet)	373
13.5.4 User function declaration in VBA (Excel macro)	374
13.5.5 User function declaration in Visual Basic	374
13.5.6 User function declaration in COM	375
13.5.7 User function declaration in MOSEL	375
13.6 Simple functions and general functions	376
13.6.1 Simple user functions	376
13.6.2 General user functions returning an array of values through a reference	376
13.6.3 General user functions returning an array of values through an argument	377
13.7 Programming Techniques for User Functions	379
13.7.1 FunctionInfo	379
13.7.2 InputNames	379
13.7.3 ReturnNames	379
13.7.4 Deltas	380
13.7.5 Return values and ReturnArray	380
13.7.6 Returning Derivatives	380
13.7.7 Function Instances	381
13.7.8 Function Objects	382
13.7.9 Calling user functions	385
13.8 Function Derivatives	387
14 Management of zero placeholder entries	389
14.1 The augmented matrix structure	389
14.2 Derivatives and zero derivatives	389
14.3 Placeholder management	390
15 Special Types of Problem	392

15.1 Nonlinear objectives	392
15.2 Quadratic Programming	392
15.3 Mixed Integer Nonlinear Programming	393
15.3.1 Approaches to MSLP	393
15.3.2 Fixing or relaxing the values of the SLP variables	394
15.3.3 Iterating at each node	394
15.3.4 Termination criteria at each node	395
15.3.5 Callbacks	395
16 Error Messages	397
17 Files used by Xpress-SLP	403
Index	404

Chapter 1

Nonlinear Problems

Xpress-SLP 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-SLP will also solve linear problems — that is, if the problem presented to Xpress-SLP does not contain any nonlinear terms, then Xpress-SLP 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.

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.

1.1 Coefficients and terms

Later in this manual, it will be helpful to distinguish between formulae written as coefficients and those written as terms.

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

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

$$X - \sin(Y) = 0,$$

$\sin(Y)$ is a *term* and cannot be written as a coefficient.

1.2 SLP variables

A variable which appears in a nonlinear coefficient or term 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 terms are not treated as SLP variables.

Any variable, whether it is related to a nonlinear term 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.

Chapter 2

Extended MPS file format

One method of inputting a problem to Xpress-SLP is from a text file which is similar to the normal MPS format matrix file. The Xpress-SLP 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).

2.1 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.
- = `MyFunc2 ( C1 , - C2 , C3 : RVP )` `MyFunc1` must be a function which can take three arguments and which returns an array of results. This formula is asking for the item in the array which is named `RVP`.

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 in principle to the length of a formula. However, there is a limit on the length of a record read by `XSLPreadprob`, which is 31000 characters. Parsing very long records can be slow, and consideration should be given to pre-parsing them and passing the parsed formula to `Xpress-SLP` rather than asking it to parse the formula itself.

2.2 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. The maximum record length for files read by `XSLPreadprob` is 31000. Note that there are limits applied by the Optimizer to the lengths of the names of rows and columns.

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).

2.3 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.

Global 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.

2.4 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 global 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-SLP function calls.

2.4.1 CV (Character variable)

`CV setname variable value`

The `CV` record defines a character variable. This is only required for user functions which have character arguments (for example, file names). The value field begins with the first non-blank character after the variable name, and the value of the variable is made up of all the characters

from that point to the end of the record. The normal free-format rules do not apply in the value field, and all spacing will be retained exactly as in the original record.

Examples:

```
CV CVSET1 MyCV1 Program Files\MyLibs\MyLib1
```

This defines the character variable named `MyCV1`. It is required because there is an embedded space in the path name which it holds.

```
CV CVSET1 MyCV1 Program Files\MyLibs\MyLib1
```

```
CV CVSET2 MyCV1 Program Files\MyLibs\MyLib2
```

This defines the character variable named `MyCV1`. There are two definitions, and the appropriate one is selected by setting the string control variable `XSLP_CVNAME` before calling `XSLPreadprob` to load the problem.

2.4.2 DC (Delayed constraint)

DC rowname [value] [= formula]

The `DC` record defines a *delayed constraint*. This allows a constraint defined in the matrix to be made non-constraining for the first few SLP iterations, before reverting to its original type (L, G, E).

The *value* field is the number of SLP iterations by which the constraint will be delayed (i.e. the number of SLP iterations during which it will be non-constraining). If a formula is used as well, then the delay will start from the time that the formula becomes nonzero.

A formula can be included as well as or instead of the value. If a formula is provided, then the constraint will be delayed until the formula evaluates to non-zero. At this point, the constraint will be delayed further in accordance with the *value* field.

If *value* is zero or is omitted, then the value of `XSLP_DCLIMIT` will be used for the value; to start immediately after the formula evaluates to nonzero, set *value* to 1.

DCs are normally checked at the end of each SLP iteration, so it is possible that a solution will be converged but activation of additional DCs will force optimization to continue. A negative *value* may be given, in which case the absolute value is used but the DC is not checked at the end of the optimization.

Examples:

```
DC Row1 3 = MV ( Row99 )
```

This defines `Row1` as a delayed constraint. When the SLP optimization starts, it will not be constraining, even though it has been defined with a constraint type in the `ROWS` section. When the marginal value of `Row99` becomes nonzero, the countdown begins, and will last for 3 further iterations. After that, the row will revert to its original constraint type.

```
DC Row1 = GT ( MV ( Row99 ) , 5 )
```

This defines `Row1` as a delayed constraint. When the SLP optimization starts, it will not be constraining, even though it has been defined with a constraint type in the `ROWS` section. When the marginal value of `Row99` is greater than 5, the countdown begins, and will last for `XSLP_DCLIMIT` further iterations. After that, the row will revert to its original constraint type.

2.4.3 DR (Determining row)

DR variable rowname [weighting]

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-SLP 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.

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 $Y * Z + 6$.

2.4.4 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.

2.4.5 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.

2.4.6 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.

2.4.7 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.

2.4.8 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.

2.4.9 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.

2.4.10 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.

2.4.11 UF (User function)

UF funcname [= extname] (arguments) linkage [= [param1] [= [param2] [= [param3]]]]

A **UF** record defines a user function.

The definition includes the list of required arguments, and the linkage or calling mechanism. For details of the fields, see the section on "Function Declaration in Xpress-SLP".

Example:

```
UF MyFunc ( DOUBLE , INTEGER ) DLL = UserLib
```

This defines a user function called **MyFunc**. It takes two arguments (an array of type double precision and an array of type integer). The linkage is **DLL** (free-standing user library or DLL) and the function is in file **UserLib**.

2.4.12 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.

2.4.13 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.

2.4.14 XV (Extended variable array)

```
XV XVname [variable] [= [inputname] [= [value]]]
```

The `XV` record defines one item of an extended variable array. With the usual abuse of notation, we shall use `XV` as a shorthand for "extended variable array". XVs are typically used to provide a list of arguments to a function, but can be used in other ways.

The meanings of the fields are as follows:

<code>XVname</code>	The name of the XV. This must be unique and must not be the same as the name of a variable or a character variable (CV).
<code>variable</code>	The name of the variable. This can be any one of the following: • a variable in the COLUMNS section • a variable implied in the coefficients within the COLUMNS section • another XV The name must be omitted if the value is provided in the value field
<code>inputname</code>	This field is used when the XV is providing arguments to a function which takes its arguments by name rather than by position. In this case, the field holds the name of the variable as known to the function. If the function takes its arguments in a fixed order, this field is not required.
<code>value</code>	The value of the item. This is not used if variable has been provided, but must be provided in other cases. The value can be a constant or a formula. If it is a formula, then it must conform to the normal rules for formulae (starting with an equals sign, each token separated by spaces).

Example:

```
XV XV1 QN2ARFD
XV XV1 QSEVREF
XV XV2 QSULCCD = CI7
XV XV2 QCONCCD = CI8
XV XV2 = CI21 = 0.6
XV XV2 = CI47 = = QRVPCCD ^ 1.25
```

`XV1` contains two items. If used in a function call such as `MyFunc (XV1)` it is equivalent to `MyFunc (QN2ARFD, QSEVREF)`.

`XV2` contains four items. All are given input names, so that a user function can identify the inputs by name instead of by position (so the order is no longer important). The third item is a constant (0.6). The fourth item is a formula (`QRVPCCD ^ 1.25`).

The main purpose of an XV is to provide a list of arguments to a function where it is inappropriate simply to list the arguments themselves. It also provides a convenient method of recording a set of arguments which is used in different functions, or in a single function which returns multiple arguments. The XV also provides functionality which is not available in simple argument lists.

The following should be noted:

- Any XV record can have an input name (even if it is used in a function which does not use or does not accept named arguments).
- Every XV record must have either a `variable` or `value` field but not both.
It is incorrect to provide both the variable and value fields, because either the item is a variable (in which case the variable name is required) or it is not (in which case the value field is required).
It is incorrect to omit both the variable and value fields because there is then no way to obtain a value for the item.
- All the records for an XV must appear together.
- The order in which the records appear in an XV will be the order in which they are used.

Chapter 3

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.

```
[Call out to user callback if set by XSLPsetcbslpstart]
Augment the matrix (create the linearized structure) if not already done
DO
  [Call out to user callback if set by XSLPsetcbiterstart]
  Solve linearized problem using the Xpress Optimizer
  Recover SLP variable and delta solution values
  [Call out to user callback if set by XSLPsetcbcascadestart]
  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 XSLPsetcbcascadevar]
  [Call out to user callback if set by XSLPsetcbcascadeend]
  Test convergence against specified tolerances and other criteria
  For each variable:
    Test convergence against specified tolerances
    [Call out to user callback if set by XSLPsetcbitervar]
  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 XSLPsetcbiterend]
  Update coefficients, bounds and RHS in linearized matrix
  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 XSLPsetcbslpend]
```

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`.

Chapter 4

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).

If any penalty error vectors are used in the solution to an SLP iteration, the penalty is automatically increased, so that ultimately the penalty error vectors should be forced out of the solution if it is feasible. If the penalty costs reach the maximum (determined by the control parameter `XSLP_ERRORMAXCOST`) then the penalty error vectors will be fixed to zero. This may cause the final solution to be infeasible, implying that there is no feasible solution to the original problem in the neighborhood of the current values of the variables. This may be because the original solution was infeasible, or it may be that the initial penalty costs were too low, allowing the linear approximations to explore too far away from the feasible region of the nonlinear problem.

If the control parameter `XSLP_ALGORITHM` has bit 8 (escalate penalties) set, then the penalty costs of individual penalty vectors which are used in the solution will be increased, as well as the overall penalty increase described above. This allows the optimization to re-balance the penalty costs which can avoid subsequent SLP iterations becoming stuck on a particular penalty vector.

If the control parameter `XSLP_ALGORITHM` has bit 10 (max cost option) set, then the penalty vectors are not fixed when `XSLP_ERRORMAXCOST` is reached. Instead, optimization continues with the cost remaining at the same value. This option should only be selected if `XSLP_ITERLIMIT` has also been set to prevent the optimization continuing for ever, or if there is some other method (such as a callback) which can be used to terminate the optimization. Using this option is helpful when the problem takes many SLP iterations to converge and has active penalty vectors in most solutions; this situation has the danger that the penalty costs become too high, resulting in numerical instability, before convergence can be achieved.

Chapter 5

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-SLP function library, functions `XSLPaddvars`, `XSLPloadvars`, and `XSLPchgvar` allow 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-SLP 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		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

Chapter 6

Convergence criteria

6.1 Convergence criteria

In Xpress-SLP there are three types of test for convergence:

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

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.

6.1.1 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.

6.1.2 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.

6.1.3 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.

6.1.4 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.

6.1.5 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.

6.1.6 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
$S \geq 10$	mark variable as converged with status S

Values of S in the range 1 to 9 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.

6.1.7 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.

6.1.8 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:

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

The static objective function (2) test is applied only after at least `XSLP_OLIMIT` SLP iterations have taken place.

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

6.1.9 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

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

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

Absolute static objective function (1) test:

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

Relative static objective function (1) test:

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

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.

6.1.10 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:

$$\delta Obj = Obj - ConvergedObj$$

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

Absolute extended convergence continuation test:

$$\delta Obj > XSLP_WTOL_A$$

Relative extended convergence continuation test:

$$\delta Obj > ABS(ConvergedObj) * XSLP_WTOL_R$$

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 7

Xpress-SLP Structures

7.1 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 .

7.1.1 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*.

7.1.2 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.

7.1.3 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

$$\begin{array}{c|c|c} & = & X \\ \hline R & F(Y, Z) & -1 \end{array}$$

Matrix structure: all augmentations

$$\begin{array}{c|c|c|c|c|c|c|c|c|c} & = & X & Y & Z & dX & dY & dZ & & \\ \hline R & F(Y_0, Z_0) & -1 & & & & F'_Y(Y_0, Z_0) & F'_Z(Y_0, Z_0) & & \\ uX & & 1 & & & -1 & & & = & X_0 \\ uY & & & 1 & & & -1 & & = & Y_0 \\ uZ & & & & 1 & & & -1 & = & Z_0 \end{array}$$

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).

7.1.4 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

	X	R+	R-
R	$F(Y,Z)$	+1	-1
P_ERROR		+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 zero 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.

7.2 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_MINUSERRORFORMAT	pE-r
Penalty error (above RHS)	XSLP_PLUSERRORFORMAT	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_bcdefghi`).

7.3 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)
Global 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 SLP-presolved prior to augmentation, the following statistics are produced:

```
Xpress-SLP Presolve:
  3 presolve passes (10)
  247 SLP variables newly identified as fixed (11)
  425 determining rows fixed (12)
  32 coefficients identified as fixed (13)
  58 columns fixed to zero (56 SLP variables) (14)
  367 columns fixed to nonzero (360 SLP variables) (15)
  139 column deltas deleted (16)
  34 column bounds tightened (6 SLP variables) (17)
```

Notes:

10. Presolve is an iterative process. Each iteration refines the problem until no further progress is made. The number of iterations (*presolve passes*) can be limited by using `XSLP_PRESOLVEPASSES`
11. SLP variables which are deduced to be fixed by virtue of constraints in the model (over and above any which are fixed by bounds in the original problem)
12. Number of determining rows which have fixed variables and constant coefficients

13. Number of coefficients which are fixed because they are functions of constants and fixed variables
14. Total number of columns fixed to zero (number of fixed SLP variables shown in brackets)
15. Total number of columns fixed to nonzero values (number of fixed SLP variables shown in brackets)
16. Total number of deltas deleted because the SLP variable is fixed
17. Total number of bounds tightened by virtue of constraints in the model.

If any of these items is zero, it will be omitted. Unless specifically requested by setting additional bits of control `XSLP_PRESOLVE`, newly fixed variables and tightened bounds are not actually applied to the model. However, they are used in the initial augmentation and during cascading to ensure that the starting points for each iteration are within the tighter bounds.

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

```
Xpress-SLP Augmentation Statistics:
Columns:
    754 implicit SLP variables                (18)
    1010 delta vectors                        (19)
    2138 penalty error vectors (1177 positive, 961 negative) (20)
Rows:
    1370 nonlinear constraints                (21)
    1010 update rows                        (22)
    1 penalty error rows                    (23)
Coefficients:
    11862 non-constant coefficients          (24)
```

Notes:

18. SLP variables appearing only in coefficients and having no constant elements
 19. Number of delta vectors created
 20. Numbers of penalty error vectors
 21. Number of constraints containing nonlinear terms
 22. Number of update rows (equals number of delta vectors)
 23. Number of rows totaling penalty vectors (error or delta)
 24. Number of non-constant coefficients in the linear augmented matrix
- The total number of rows in the augmented matrix is (2) + (22) + (23)
 - The total number of columns in the augmented matrix is (3) + (18) + (19) + (20) + (23)
 - The total number of elements in the original matrix is (4) + (6)
 - The total number of elements in the augmented matrix is (4) + (24) + (19) + 2*(20) + 2*(23)

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-SLP integer attributes using the `XSLPgetintattrib` function.

7.4 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 8

Problem Attributes

During the optimization process, various properties of the problem being solved are stored and made available to users of the Xpress-SLP 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_`. It is possible to use the above functions with 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-SLP attributes:

<code>XSLP_COEFFICIENTS</code>	Number of nonlinear coefficients	p. 40
<code>XSLP_CURRENTDELTACOST</code>	Current value of penalty cost multiplier for penalty delta vectors	p. 37
<code>XSLP_CURRENTERRORCOST</code>	Current value of penalty cost multiplier for penalty error vectors	p. 37
<code>XSLP_CVS</code>	Number of character variables	p. 40
<code>XSLP_DELTAS</code>	Number of delta vectors created during augmentation	p. 40
<code>XSLP_ECFCOUNT</code>	Number of infeasible constraints found at the point of linearization	p. 40
<code>XSLP_EQUALSCOLUMN</code>	Index of the reserved "=" column	p. 40
<code>XSLP_ERRORCOSTS</code>	Total penalty costs in the solution	p. 37
<code>XSLP_GLOBALFUNCOBJECT</code>	The user-defined global function object	p. 50

XSLP_IFS	Number of internal functions	p. 41
XSLP_IMPLICITVARIABLES	Number of SLP variables appearing only in coefficients	p. 41
XSLP_INTERNALFUNCCALLS	Number of calls made to internal functions	p. 41
XSLP_ITER	SLP iteration count	p. 41
XSLP_MINORVERSION	Xpress-SLP minor version number	p. 41
XSLP_MINUSPENALTYERRORS	Number of negative penalty error vectors	p. 42
XSLP_MIPITER	Total number of SLP iterations in MISLP	p. 42
XSLP_MIPNODES	Number of nodes explored in MISLP	p. 42
XSLP_MIPPROBLEM	The underlying Optimizer MIP problem	p. 50
XSLP_NONLINEARCONSTRAINTS	Number of nonlinear constraints in the problem	p. 42
XSLP_OBJSENSE	Objective function sense	p. 79
XSLP_OBJVAL	Objective function value excluding any penalty costs	p. 38
XSLP_PENALTYDELTACOLUMN	Index of column costing the penalty delta row	p. 42
XSLP_PENALTYDELTAROW	Index of equality row holding the penalties for delta vectors	p. 43
XSLP_PENALTYDELTAS	Number of penalty delta vectors	p. 43
XSLP_PENALTYDELTATOTAL	Total activity of penalty delta vectors	p. 38
XSLP_PENALTYDELTAVALUE	Total penalty cost attributed to penalty delta vectors	p. 38
XSLP_PENALTYERRORCOLUMN	Index of column costing the penalty error row	p. 43
XSLP_PENALTYERRORROW	Index of equality row holding the penalties for penalty error vectors	p. 43
XSLP_PENALTYERRORS	Number of penalty error vectors	p. 43
XSLP_PENALTYERRORTOTAL	Total activity of penalty error vectors	p. 38
XSLP_PENALTYERRORVALUE	Total penalty cost attributed to penalty error vectors	p. 38
XSLP_PLUSPENALTYERRORS	Number of positive penalty error vectors	p. 44
XSLP_PRESOLVEDELETEDDELTA	Number of potential delta variables deleted by XSLPpresolve	p. 44
XSLP_PRESOLVEFIXEDCOEF	Number of SLP coefficients fixed by XSLPpresolve	p. 44
XSLP_PRESOLVEFIXEDDR	Number of determining rows fixed by XSLPpresolve	p. 44
XSLP_PRESOLVEFIXEDNZCOL	Number of variables fixed to a nonzero value by XSLPpresolve	p. 45
XSLP_PRESOLVEFIXEDSLPVAR	Number of SLP variables fixed by XSLPpresolve	p. 45
XSLP_PRESOLVEFIXEDZCOL	Number of variables fixed at zero by XSLPpresolve	p. 45
XSLP_PRESOLVEPASSES	Number of passes made by the SLP nonlinear presolve procedure	p. 45
XSLP_PRESOLVETIGHTENED	Number of bounds tightened by XSLPpresolve	p. 46

XSLP_SBXCONVERGED	Number of step-bounded variables converged only on extended criteria	p. 46
XSLP_STATUS	Bitmap holding the problem convergence status	p. 46
XSLP_TOLSETS	Number of tolerance sets	p. 47
XSLP_UCCONSTRAINEDCOUNT	Number of unconverged variables with coefficients in constraining rows	p. 47
XSLP_UFINSTANCES	Number of user function instances	p. 47
XSLP_UFS	Number of user functions	p. 47
XSLP_UNCONVERGED	Number of unconverged values	p. 47
XSLP_UNIQUEPREFIX	Unique prefix for generated names	p. 51
XSLP_USEDERIVATIVES	Indicates whether numeric or analytic derivatives were used to create the linear approximations and solve the problem	p. 48
XSLP_USERFUNCCALLS	Number of calls made to user functions	p. 48
XSLP_VALIDATIONINDEX_A	Absolute validation index	p. 39
XSLP_VALIDATIONINDEX_R	Relative validation index	p. 39
XSLP_VARIABLES	Number of SLP variables	p. 48
XSLP_VERSION	Xpress-SLP major version number	p. 48
XSLP_VERSIONDATE	Date of creation of Xpress-SLP	p. 51
XSLP_VSOLINDEX	Vertex solution index	p. 39
XSLP_XPRSPROBLEM	The underlying Optimizer problem	p. 50
XSLP_XSLPPROBLEM	The Xpress-SLP problem	p. 50
XSLP_XVS	Number of extended variable arrays	p. 48
XSLP_ZEROESRESET	Number of placeholder entries set to zero	p. 49
XSLP_ZEROESRETAINED	Number of potentially zero placeholders left untouched	p. 49
XSLP_ZEROESTOTAL	Number of potential zero placeholder entries	p. 49

8.1 Double problem attributes

XSLP_CURRENTDELTACOST

Description	Current value of penalty cost multiplier for penalty delta vectors
Type	Double
Set by routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_DELTACOST</code> , <code>XSLP_ERRORCOST</code> , <code>XSLP_CURRENTERRORCOST</code>

XSLP_CURRENTERRORCOST

Description	Current value of penalty cost multiplier for penalty error vectors
Type	Double
Set by routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_DELTACOST</code> , <code>XSLP_ERRORCOST</code> , <code>XSLP_CURRENTDELTACOST</code>

XSLP_ERRORCOSTS

Description	Total penalty costs in the solution
Type	Double
Set by routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_OBJSENSE

Description	Objective function sense
Type	Double
Values	-1 Maximize 1 Minimize
Set by routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_OBJVAL

Description	Objective function value excluding any penalty costs
Type	Double
Set by routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>

XSLP_PENALTYDELTATOTAL

Description	Total activity of penalty delta vectors
Type	Double
Set by routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>

XSLP_PENALTYDELTAVALUE

Description	Total penalty cost attributed to penalty delta vectors
Type	Double
Set by routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>

XSLP_PENALTYERRORTOTAL

Description	Total activity of penalty error vectors
Type	Double
Set by routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>

XSLP_PENALTYERRORVALUE

Description	Total penalty cost attributed to penalty error vectors
Type	Double
Set by routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>

XSLP_VALIDATIONINDEX_A

Description	Absolute validation index
Type	Double
Set by routines	XSLPvalidate

XSLP_VALIDATIONINDEX_R

Description	Relative validation index
Type	Double
Set by routines	XSLPvalidate

XSLP_VSOLINDEX

Description	Vertex solution index
Type	Double
Notes	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. 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.

8.2 Integer problem attributes

XSLP_COEFFICIENTS

Description	Number of nonlinear coefficients
Type	Integer
Set by routines	XSLPaddcoefs , XSLPchgcoef , XSLPloadcoefs , XSLPreadprob

XSLP_CVS

Description	Number of character variables
Type	Integer
Set by routines	XSLPaddcvars , XSLPchgcvvar , XSLPloadcvars , XSLPreadprob

XSLP_DELTAS

Description	Number of delta vectors created during augmentation
Type	Integer
Set by routines	XSLPconstruct

XSLP_ECFCOUNT

Description	Number of infeasible constraints found at the point of linearization
Type	Integer
Set by routines	XSLPmaxim , XSLPminim
See also	XSLP_ECFCHECK , XSLP_ECFTOL_A , XSLP_ECFTOL_R

XSLP_EQUALS COLUMN

Description	Index of the reserved "=" column
Type	Integer
Note	This index always counts from 1. It is zero if there is no "equals column".
Set by routines	XSLPconstruct , XSLPreadprob

XSLP_IFS

Description	Number of internal functions
Type	Integer
Set by routines	<code>XSLPcreateprob</code>

XSLP_IMPLICITVARIABLES

Description	Number of SLP variables appearing only in coefficients
Type	Integer
Set by routines	<code>XSLPconstruct</code>

XSLP_INTERNALFUNCCALLS

Description	Number of calls made to internal functions
Type	Integer
Set by routines	<code>XSLPcascade</code> , <code>XSLPconstruct</code> , <code>XSLPevaluatecoef</code> , <code>XSLPevaluateformula</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_ITER

Description	SLP iteration count
Type	Integer
Set by routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_MINORVERSION

Description	Xpress-SLP minor version number
Type	Integer
Set by routines	<code>XSLPinit</code>

XSLP_MINUSPENALTYERRORS

Description	Number of negative penalty error vectors
Type	Integer
Set by routines	<code>XSLPconstruct</code>

XSLP_MIPITER

Description	Total number of SLP iterations in MISLP
Type	Integer
Set by routines	<code>XSLPglobal</code>

XSLP_MIPNODES

Description	Number of nodes explored in MISLP
Type	Integer
Set by routines	<code>XSLPglobal</code>

XSLP_NONLINEARCONSTRAINTS

Description	Number of nonlinear constraints in the problem
Type	Integer
Set by routines	<code>XSLPconstruct</code>

XSLP_PENALTYDELTACOLUMN

Description	Index of column costing the penalty delta row
Type	Integer
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

Description	Index of equality row holding the penalties for delta vectors
Type	Integer
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

Description	Number of penalty delta vectors
Type	Integer
Set by routines	<code>XSLPconstruct</code>

XSLP_PENALTYERRORCOLUMN

Description	Index of column costing the penalty error row
Type	Integer
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

Description	Index of equality row holding the penalties for penalty error vectors
Type	Integer
Note	This index always counts from 1. It is zero if there are no penalty error vectors.
Set by routines	<code>XSLPconstruct</code>

XSLP_PENALTYERRORS

Description	Number of penalty error vectors
Type	Integer
Set by routines	<code>XSLPconstruct</code>

XSLP_PLUSPENALTYERRORS

Description	Number of positive penalty error vectors
Type	Integer
Set by routines	<code>XSLPconstruct</code>

XSLP_PRESOLVEDELETEDDELTA

Description	Number of potential delta variables deleted by <code>XSLPpresolve</code>
Type	Integer
Note	A potential delta variable is deleted when an SLP variable is identified as not interacting in a nonlinear way with any constraints (that is, it appears only in linear constraints, or is fixed).
Set by routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVEFIXEDCOEF</code> , <code>XSLP_PRESOLVEFIXEDDDR</code> , <code>XSLP_PRESOLVEFIXEDNZCOL</code> , <code>XSLP_PRESOLVEFIXEDSLPVAR</code> , <code>XSLP_PRESOLVEFIXEDZCOL</code> , <code>XSLP_PRESOLVETIGHTENED</code>

XSLP_PRESOLVEFIXEDCOEF

Description	Number of SLP coefficients fixed by <code>XSLPpresolve</code>
Type	Integer
Set by routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVEDELETEDDELTA</code> , <code>XSLP_PRESOLVEFIXEDDDR</code> , <code>XSLP_PRESOLVEFIXEDNZCOL</code> , <code>XSLP_PRESOLVEFIXEDSLPVAR</code> , <code>XSLP_PRESOLVEFIXEDZCOL</code> , <code>XSLP_PRESOLVETIGHTENED</code>

XSLP_PRESOLVEFIXEDDDR

Description	Number of determining rows fixed by <code>XSLPpresolve</code>
Type	Integer
Set by routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVEDELETEDDELTA</code> , <code>XSLP_PRESOLVEFIXEDCOEF</code> , <code>XSLP_PRESOLVEFIXEDNZCOL</code> , <code>XSLP_PRESOLVEFIXEDSLPVAR</code> , <code>XSLP_PRESOLVEFIXEDZCOL</code> , <code>XSLP_PRESOLVETIGHTENED</code>

XSLP_PRESOLVEFIXEDNZCOL

Description	Number of variables fixed to a nonzero value by <code>XSLPpresolve</code>
Type	Integer
Set by routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVEDELETEDDELTA</code> , <code>XSLP_PRESOLVEFIXEDCOEF</code> , <code>XSLP_PRESOLVEFIXEDDDR</code> , <code>XSLP_PRESOLVEFIXEDSLPVAR</code> , <code>XSLP_PRESOLVEFIXEDZCOL</code> , <code>XSLP_PRESOLVETIGHTENED</code>

XSLP_PRESOLVEFIXEDSLPVAR

Description	Number of SLP variables fixed by <code>XSLPpresolve</code>
Type	Integer
Set by routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVEDELETEDDELTA</code> , <code>XSLP_PRESOLVEFIXEDCOEF</code> , <code>XSLP_PRESOLVEFIXEDDDR</code> , <code>XSLP_PRESOLVEFIXEDNZCOL</code> , <code>XSLP_PRESOLVEFIXEDZCOL</code> , <code>XSLP_PRESOLVETIGHTENED</code>

XSLP_PRESOLVEFIXEDZCOL

Description	Number of variables fixed at zero by <code>XSLPpresolve</code>
Type	Integer
Set by routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVEDELETEDDELTA</code> , <code>XSLP_PRESOLVEFIXEDCOEF</code> , <code>XSLP_PRESOLVEFIXEDDDR</code> , <code>XSLP_PRESOLVEFIXEDNZCOL</code> , <code>XSLP_PRESOLVEFIXEDSLPVAR</code> , <code>XSLP_PRESOLVETIGHTENED</code>

XSLP_PRESOLVEPASSES

Description	Number of passes made by the SLP nonlinear presolve procedure
Type	Integer
Set by routines	<code>XSLPpresolve</code>

XSLP_PRESOLVETIGHTENED

Description	Number of bounds tightened by <code>XSLPpresolve</code>
Type	Integer
Set by routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVEDELETEDDELTA</code> , <code>XSLP_PRESOLVEFIXEDCOEF</code> , <code>XSLP_PRESOLVEFIXEDDDR</code> , <code>XSLP_PRESOLVEFIXEDNZCOL</code> , <code>XSLP_PRESOLVEFIXEDSLPVAR</code> , <code>XSLP_PRESOLVEFIXEDZCOL</code>

XSLP_SBXCONVERGED

Description	Number of step-bounded variables converged only on extended criteria
Type	Integer
Set by routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_STATUS

Description	Bitmap holding the problem convergence status	
Type	Integer	
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 function with respect to <code>XSLP_VTOL_A</code> , <code>XSLP_VTOL_R</code> and <code>XSLP_VCOUNT</code> .
	9	SLP terminated on max time.
	10	SLP terminated by user.
Note	A value of zero after SLP optimization means that the solution is fully converged.	
Set by routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>	

XSLP_TOLSETS

Description Number of tolerance sets

Type Integer

Set by routines `XSLPaddtolsets`, `XSLPchgtolset`, `XSLPloadtolsets`, `XSLPreadprob`

XSLP_UCCONSTRAINEDCOUNT

Description Number of unconverged variables with coefficients in constraining rows

Type Integer

Set by routines `XSLPmaxim`, `XSLPminim`

XSLP_UFINSTANCES

Description Number of user function instances

Type Integer

Set by routines `XSLPconstruct`

XSLP_UFS

Description Number of user functions

Type Integer

Set by routines `XSLPadduserfuncs`, `XSLPchguserfunc`, `XSLPloaduserfuncs`, `XSLPreadprob`

XSLP_UNCONVERGED

Description Number of unconverged values

Type Integer

Note Prior to the first iteration this will return -1.

Set by routines `XSLPmaxim`, `XSLPminim`

XSLP_USEDERIVATIVES

Description	Indicates whether numeric or analytic derivatives were used to create the linear approximations and solve the problem	
Type	Integer	
Values	0	numeric derivatives.
	1	analytic derivatives for all formulae unless otherwise specified.
Set by routines	XSLPconstruct	

XSLP_USERFUNCCALLS

Description	Number of calls made to user functions	
Type	Integer	
Set by routines	XSLPccascade, XSLPconstruct, XSLPevaluatecoef, XSLPevaluateformula, XSLPmaxim, XSLPminim	

XSLP_VARIABLES

Description	Number of SLP variables	
Type	Integer	
Set by routines	XSLPconstruct	

XSLP_VERSION

Description	Xpress-SLP major version number	
Type	Integer	
Set by routines	XSLPinit	

XSLP_XVS

Description	Number of extended variable arrays	
Type	Integer	
Set by routines	XSLPaddxvs, XSLPchg xv, XSLPloadxvs, XSLPreadprob	

XSLP_ZEROESRESET

Description	Number of placeholder entries set to zero
Type	Integer
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

Description	Number of potentially zero placeholders left untouched
Type	Integer
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

Description	Number of potential zero placeholder entries
Type	Integer
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>

8.3 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
Type	Reference
Set by routines	<code>XSLPglobal</code>

XSLP_XPRSPROBLEM

Description	The underlying Optimizer problem
Type	Reference
Set by routines	<code>XSLPcreateprob</code>

XSLP_XSLPPROBLEM

Description	The Xpress-SLP problem
Type	Reference
Set by routines	<code>XSLPcreateprob</code>

XSLP_GLOBALFUNCOBJECT

Description	The user-defined global function object
Type	Reference
Set by routines	<code>XSLPchgfuncobject</code> , <code>XSLPchguserfuncobject</code> , <code>XSLPsetfuncobject</code> , <code>XSLPsetuserfuncobject</code>

8.4 String problem attributes

XSLP_UNIQUEPREFIX

Description Unique prefix for generated names

Type String

Set by routines `XSLPsetuniqueprefix`

XSLP_VERSIONDATE

Description Date of creation of Xpress-SLP

Type String

Note The format of the date is dd mmm yyyy.

Set by routines `XSLPinit`

Chapter 9

Control Parameters

Various controls exist within Xpress-SLP 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-SLP 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_`. It is possible to use the above functions with 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-SLP controls:

<code>XSLP_ALGORITHM</code>	Bit map describing the SLP algorithm(s) to be used	p. 89
<code>XSLP_ATOL_A</code>	Absolute delta convergence tolerance	p. 59
<code>XSLP_ATOL_R</code>	Relative delta convergence tolerance	p. 59
<code>XSLP_AUGMENTATION</code>	Bit map describing the SLP augmentation method(s) to be used	p. 90
<code>XSLP_AUTOSAVE</code>	Frequency with which to save the model	p. 92
<code>XSLP_BARLIMIT</code>	Number of initial SLP iterations using the barrier method	p. 92
<code>XSLP_CASCADE</code>	Bit map describing the cascading to be used	p. 92
<code>XSLP_CASCADENLIMIT</code>	Maximum number of iterations for cascading with non-linear determining rows	p. 93

XSLP_CASCADETOL_PA	Absolute cascading print tolerance	p. 59
XSLP_CASCADETOL_PR	Relative cascading print tolerance	p. 60
XSLP_CDTOL_A	Absolute tolerance for deducing constant derivatives	p. 60
XSLP_CDTOL_R	Relative tolerance for deducing constant derivatives	p. 60
XSLP_CONTROL	Bit map describing which Xpress-SLP functions also activate the corresponding Optimizer Library function	p. 93
XSLP_CTOI	Closure convergence tolerance	p. 61
XSLP_CVNAME	Name of the set of character variables to be used	p. 124
XSLP_DAMP	Damping factor for updating values of variables	p. 61
XSLP_DAMPEXPAND	Multiplier to increase damping factor during dynamic damping	p. 61
XSLP_DAMPMAX	Maximum value for the damping factor of a variable during dynamic damping	p. 62
XSLP_DAMPMIN	Minimum value for the damping factor of a variable during dynamic damping	p. 62
XSLP_DAMP SHRINK	Multiplier to decrease damping factor during dynamic damping	p. 63
XSLP_DAMPSTART	SLP iteration at which damping is activated	p. 94
XSLP_DCLIMIT	Default iteration delay for delayed constraints	p. 94
XSLP_DCLOG	Amount of logging information for activation of delayed constraints	p. 95
XSLP_DECOMPOSE	Bitmap controlling the action of function XSLPdecompose	p. 95
XSLP_DECOMPOSEPASSLIMIT	Maximum number of repeats of presolve+decompose	p. 96
XSLP_DEFAULTIV	Default initial value for an SLP variable if none is explicitly given	p. 63
XSLP_DEFAULTSTEPBOUND	Minimum initial value for the step bound of an SLP variable if none is explicitly given	p. 63
XSLP_DELAYUPDATEROWS	Number of SLP iterations before update rows are fully activated	p. 95
XSLP_DELTA_A	Absolute perturbation of values for calculating numerical derivatives	p. 64
XSLP_DELTA_R	Relative perturbation of values for calculating numerical derivatives	p. 64
XSLP_DELTA_X	Minimum absolute value of delta coefficients to be retained	p. 64
XSLP_DELTA_Z	Zero tolerance used when calculating numerical derivatives	p. 65
XSLP_DELTACOST	Initial penalty cost multiplier for penalty delta vectors	p. 65
XSLP_DELTACOSTFACTOR	Factor for increasing cost multiplier on total penalty delta vectors	p. 65
XSLP_DELTAFORMAT	Formatting string for creation of names for SLP delta vectors	p. 124
XSLP_DELTAMAXCOST	Maximum penalty cost multiplier for penalty delta vectors	p. 66

XSLP_DELTAOFFSET	Position of first character of SLP variable name used to create name of delta vector	p. 96
XSLP_DELTAZLIMIT	Number of SLP iterations during which to apply XSLP_DELTA_Z	p. 97
XSLP_DERIVATIVES	Method of calculating derivatives	p. 97
XSLP_DJTOL	Tolerance on DJ value for determining if a variable is at its step bound	p. 66
XSLP_ECFCHECK	Check feasibility at the point of linearization for extended convergence criteria	p. 98
XSLP_ECFTOL_A	Absolute tolerance on testing feasibility at the point of linearization	p. 66
XSLP_ECFTOL_R	Relative tolerance on testing feasibility at the point of linearization	p. 67
XSLP_EQTOL_A	Absolute tolerance on equality testing in logical functions	p. 67
XSLP_EQTOL_R	Relative tolerance on equality testing in logical functions	p. 67
XSLP_ERRORCOST	Initial penalty cost multiplier for penalty error vectors	p. 68
XSLP_ERRORCOSTFACTOR	Factor for increasing cost multiplier on total penalty error vectors	p. 68
XSLP_ERRORMAXCOST	Maximum penalty cost multiplier for penalty error vectors	p. 68
XSLP_ERROROFFSET	Position of first character of constraint name used to create name of penalty error vectors	p. 98
XSLP_ERRORTOL_A	Absolute tolerance for error vectors	p. 69
XSLP_ERRORTOL_P	Absolute tolerance for printing error vectors	p. 69
XSLP_ESCALATION	Factor for increasing cost multiplier on individual penalty error vectors	p. 69
XSLP_ETOL_A	Absolute tolerance on penalty vectors	p. 70
XSLP_ETOL_R	Relative tolerance on penalty vectors	p. 70
XSLP_EVALUATE	Evaluation strategy for user functions	p. 99
XSLP_EVTOL_A	Absolute tolerance on total penalty costs	p. 70
XSLP_EVTOL_R	Relative tolerance on total penalty costs	p. 71
XSLP_EXCELVISIBLE	Display of Excel when evaluating user functions written in Excel	p. 99
XSLP_EXPAND	Multiplier to increase a step bound	p. 71
XSLP_EXTRACVS	Expansion number for character variables	p. 99
XSLP_EXTRAUFS	Expansion number for user functions	p. 100
XSLP_EXTRAXVITEMS	Expansion number for XV items	p. 100
XSLP_EXTRAXVS	Expansion number for XVs	p. 101
XSLP_FUNCEVAL	Bit map for determining the method of evaluating user functions and their derivatives	p. 101

XSLP_GRANULARITY	Base for calculating penalty costs	p. 71
XSLP_INFEASLIMIT	The maximum number of consecutive infeasible SLP iterations which can occur before Xpress-SLP terminates	p. 102
XSLP_INFINITY	Value returned by a divide-by-zero in a formula	p. 72
XSLP_ITERLIMIT	The maximum number of SLP iterations	p. 102
XSLP_ITOL_A	Absolute impact convergence tolerance	p. 72
XSLP_ITOL_R	Relative impact convergence tolerance	p. 73
XSLP_IVNAME	Name of the set of initial values to be used	p. 124
XSLP_LOG	Level of printing during SLP iterations	p. 102
XSLP_MAXTIME	The maximum time in seconds that the SLP optimization will run before it terminates	p. 103
XSLP_MAXWEIGHT	Maximum penalty weight for delta or error vectors	p. 73
XSLP_MEM_CALCSTACK	Memory allocation for formula calculations	p. 119
XSLP_MEM_COEF	Memory allocation for nonlinear coefficients	p. 120
XSLP_MEM_COL	Memory allocation for additional information on matrix columns	p. 120
XSLP_MEM_CVAR	Memory allocation for character variables	p. 120
XSLP_MEM_DERIVATIVES	Memory allocation for analytic derivatives	p. 120
XSLP_MEM_EXCELDDOUBLE	Memory allocation for return values from Excel user functions	p. 120
XSLP_MEM_FORMULA	Memory allocation for formulae	p. 120
XSLP_MEM_FORMULAHASH	Memory allocation for internal formula array	p. 121
XSLP_MEM_FORMULAVALUE	Memory allocation for formula values and derivatives	p. 121
XSLP_MEM_ITERLOG	Memory allocation for SLP iteration summary	p. 121
XSLP_MEM_RETURNARRAY	Memory allocation for return values from multi-valued user function	p. 121
XSLP_MEM_ROW	Memory allocation for additional information on matrix rows	p. 121
XSLP_MEM_STACK	Memory allocation for parsed formulae, analytic derivatives	p. 121
XSLP_MEM_STRING	Memory allocation for strings of all types	p. 122
XSLP_MEM_TOL	Memory allocation for tolerance sets	p. 122
XSLP_MEM_UF	Memory allocation for user functions	p. 122
XSLP_MEM_VALSTACK	Memory allocation for intermediate values for analytic derivatives	p. 122
XSLP_MEM_VAR	Memory allocation for SLP variables	p. 122
XSLP_MEM_XF	Memory allocation for complicated functions	p. 122
XSLP_MEM_XFNAMES	Memory allocation for complicated function input and return names	p. 123

XSLP_MEM_XFVALUE	Memory allocation for complicated function values	p. 123
XSLP_MEM_XROW	Memory allocation for extended row information	p. 123
XSLP_MEM_XV	Memory allocation for XVs	p. 123
XSLP_MEM_XVITEM	Memory allocation for individual XV entries	p. 123
XSLP_MEMORYFACTOR	Factor for expanding size of dynamic arrays in memory	p. 74
XSLP_MINUSDELTAFORMAT	Formatting string for creation of names for SLP negative penalty delta vectors	p. 125
XSLP_MINUSERERRORFORMAT	Formatting string for creation of names for SLP negative penalty error vectors	p. 125
XSLP_MINWEIGHT	Minimum penalty weight for delta or error vectors	p. 74
XSLP_MIPALGORITHM	Bitmap describing the MISLP algorithms to be used	p. 103
XSLP_MIPCUTOFF_A	Absolute objective function cutoff for MIP termination	p. 74
XSLP_MIPCUTOFF_R	Absolute objective function cutoff for MIP termination	p. 75
XSLP_MIPCUTOFFCOUNT	Number of SLP iterations to check when considering a node for cutting off	p. 104
XSLP_MIPCUTOFFLIMIT	Number of SLP iterations to check when considering a node for cutting off	p. 104
XSLP_MIPDEFAULTALGORITHM	Default algorithm to be used during the global search in MISLP	p. 105
XSLP_MIPERRORTOL_A	Absolute penalty error cost tolerance for MIP cut-off	p. 75
XSLP_MIPERRORTOL_R	Relative penalty error cost tolerance for MIP cut-off	p. 76
XSLP_MIPFIXSTEPBOUNDS	Bitmap describing the step-bound fixing strategy during MISLP	p. 105
XSLP_MIPITERLIMIT	Maximum number of SLP iterations at each node	p. 106
XSLP_MIPLOG	Frequency with which MIP status is printed	p. 106
XSLP_MIPOCOUNT	Number of SLP iterations at each node over which to measure objective function variation	p. 106
XSLP_MIPOTOL_A	Absolute objective function tolerance for MIP termination	p. 76
XSLP_MIPOTOL_R	Relative objective function tolerance for MIP termination	p. 76
XSLP_MIPRELAXSTEPBOUNDS	Bitmap describing the step-bound relaxation strategy during MISLP	p. 107
XSLP_MTOL_A	Absolute effective matrix element convergence tolerance	p. 77
XSLP_MTOL_R	Relative effective matrix element convergence tolerance	p. 77
XSLP_MVTOL	Marginal value tolerance for determining if a constraint is slack	p. 78
XSLP_OBJSENSE	Objective function sense	p. 79
XSLP_OBJTOPENALTYCOST	Factor to estimate initial penalty costs from objective function	p. 79

XSLP_OCOUNT	Number of SLP iterations over which to measure objective function variation for static objective (2) convergence criterion	p. 107
XSLP_OTOL_A	Absolute static objective (2) convergence tolerance	p. 80
XSLP_OTOL_R	Relative static objective (2) convergence tolerance	p. 80
XSLP_PENALTYCOLFORMAT	Formatting string for creation of the names of the SLP penalty transfer vectors	p. 125
XSLP_PENALTYINFOSTART	Iteration from which to record row penalty information	p. 108
XSLP_PENALTYROWFORMAT	Formatting string for creation of the names of the SLP penalty rows	p. 126
XSLP_PLUSDELTAFORMAT	Formatting string for creation of names for SLP positive penalty delta vectors	p. 126
XSLP_PLUSERRORFORMAT	Formatting string for creation of names for SLP positive penalty error vectors	p. 127
XSLP_PRESOLVE	Bitmap indicating the SLP presolve actions to be taken	p. 108
XSLP_PRESOLVEPASSLIMIT	Maximum number of passes through the problem to improve SLP bounds	p. 108
XSLP_PRESOLVEZERO	Minimum absolute value for a variable which is identified as nonzero during SLP presolve	p. 81
XSLP_SAMECOUNT	Number of steps reaching the step bound in the same direction before step bounds are increased	p. 109
XSLP_SAMEDAMP	Number of steps in same direction before damping factor is increased	p. 109
XSLP_SBLOROWFORMAT	Formatting string for creation of names for SLP lower step bound rows	p. 127
XSLP_SBNAME	Name of the set of initial step bounds to be used	p. 127
XSLP_SBROWOFFSET	Position of first character of SLP variable name used to create name of SLP lower and upper step bound rows	p. 110
XSLP_SBSTART	SLP iteration after which step bounds are first applied	p. 110
XSLP_SBUPROWFORMAT	Formatting string for creation of names for SLP upper step bound rows	p. 128
XSLP_SCALE	When to re-scale the SLP problem	p. 110
XSLP_SCALECOUNT	Iteration limit used in determining when to re-scale the SLP matrix	p. 111
XSLP_SHRINK	Multiplier to reduce a step bound	p. 81
XSLP_SLPLOG	Frequency with which SLP status is printed	p. 111
XSLP_STOL_A	Absolute slack convergence tolerance	p. 81
XSLP_STOL_R	Relative slack convergence tolerance	p. 82
XSLP_STOPOUTOFRANGE	Stop optimization and return error code if internal function argument is out of range	p. 112

XSLP__TIMEPRINT	Print additional timings during SLP optimization	p. 112
XSLP__TOLNAME	Name of the set of tolerance sets to be used	p. 128
XSLP__UNFINISHEDLIMIT	Number of times within one SLP iteration that an unfinished LP optimization will be continued	p. 112
XSLP__UPDATEFORMAT	Formatting string for creation of names for SLP update rows	p. 128
XSLP__UPDATEOFFSET	Position of first character of SLP variable name used to create name of SLP update row	p. 113
XSLP__VALIDATIONTOL_A	Absolute tolerance for the XSLPvalidate procedure	p. 82
XSLP__VALIDATIONTOL_R	Relative tolerance for the XSLPvalidate procedure	p. 83
XSLP__VCOUNT	Number of SLP iterations over which to measure static objective (3) convergence	p. 113
XSLP__VLIMIT	Number of SLP iterations after which static objective (3) convergence testing starts	p. 114
XSLP__VTOL_A	Absolute static objective (3) convergence tolerance	p. 83
XSLP__VTOL_R	Relative static objective (3) convergence tolerance	p. 84
XSLP__WCOUNT	Number of SLP iterations over which to measure the objective for the extended convergence continuation criterion	p. 114
XSLP__WTOL_A	Absolute extended convergence continuation tolerance	p. 84
XSLP__WTOL_R	Relative extended convergence continuation tolerance	p. 85
XSLP__XCOUNT	Number of SLP iterations over which to measure static objective (1) convergence	p. 115
XSLP__XLIMIT	Number of SLP iterations up to which static objective (1) convergence testing starts	p. 116
XSLP__XTOL_A	Absolute static objective function (1) tolerance	p. 86
XSLP__XTOL_R	Relative static objective function (1) tolerance	p. 87
XSLP__ZERO	Absolute zero tolerance	p. 87
XSLP__ZEROCRITERION	Bitmap determining the behavior of the placeholder deletion procedure	p. 117
XSLP__ZEROCRITERIONCOUNT	Number of consecutive times a placeholder entry is zero before being considered for deletion	p. 117
XSLP__ZEROCRITERIONSTART	SLP iteration at which criteria for deletion of placeholder entries are first activated.	p. 118

9.1 Double control parameters

XSLP_ATOL_A

Description	Absolute delta convergence tolerance
Type	Double
Note	The absolute delta convergence criterion assesses the change in value of a variable (δX) against the absolute delta convergence tolerance. If $\delta X < XSLP_ATOL_A$ then the variable has converged on the absolute delta convergence criterion.
Default value	0.001
See also	Convergence Criteria, XSLP_ATOL_R

XSLP_ATOL_R

Description	Relative delta convergence tolerance
Type	Double
Note	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 $\delta X < X * XSLP_ATOL_R$ then the variable has converged on the relative delta convergence criterion.
Default value	0.001
See also	Convergence Criteria, XSLP_ATOL_A

XSLP_CASCADETOL_PA

Description	Absolute cascading print tolerance
Type	Double
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 $\delta X < XSLP_CASCADETOL_PA$ then the change is within the absolute cascading print tolerance and will not be printed.
Default value	0.01
See also	Cascading, XSLP_CASCADETOL_PR
Affects routines	XSLPcascade

XSLP_CASCADETOL_PR

Description	Relative cascading print tolerance
Type	Double
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 $\delta X < X * XSLP_CASCADETOL_PR$ then the change is within the relative cascading print tolerance and will not be printed.
Default value	0.01
See also	Cascading, XSLP_CASCADETOL_PA
Affects routines	XSLPcascade

XSLP_CDTOL_A

Description	Absolute tolerance for deducing constant derivatives
Type	Double
Note	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 $abs(D_d - D_u) \leq XSLP_CDTOL_A$ then the derivative is regarded as constant.
Default value	1.0e-08
See also	XSLP_CDTOL_R

XSLP_CDTOL_R

Description	Relative tolerance for deducing constant derivatives
Type	Double
Note	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.

Default value 1.0e-08
See also [XSLP_CDTOL_A](#)

XSLP_CTOL

Description Closure convergence tolerance
Type Double
Notes 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 $\delta X < B * \text{XSLP_CTOL}$ then the variable has converged on the closure convergence criterion. 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.
Default value 0.001
See also Convergence Criteria, [XSLP_ATOL_A](#), [XSLP_ATOL_R](#)

XSLP_DAMP

Description Damping factor for updating values of variables
Type Double
Note The damping factor sets the next *assumed value* for a variable based on the previous assumed value (X_0) and the *actual value* (X_1). The new assumed value is given by $X_1 * \text{XSLP_DAMP} + X_0 * (1 - \text{XSLP_DAMP})$
Default value 1
See also Xpress-SLP Solution Process, [XSLP_DAMPEXPAND](#), [XSLP_DAMPMAX](#), [XSLP_DAMPMIN](#), [XSLP_DAMP SHRINK](#), [XSLP_DAMPSTART](#)
Affects routines [XSLPmaxim](#), [XSLPminim](#)

XSLP_DAMPEXPAND

Description Multiplier to increase damping factor during dynamic damping
Type Double

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 <code>XSLP_SAMEDAMP</code> successive changes in the same direction for a variable, then the damping factor (<i>D</i>) for the variable will be reset to $D * XSLP_DAMPEXPAND + XSLP_DAMPMAX * (1 - XSLP_DAMPEXPAND)$
Default value	1
See also	Xpress-SLP Solution Process, <code>XSLP_ALGORITHM</code> , <code>XSLP_DAMP</code> , <code>XSLP_DAMPMAX</code> , <code>XSLP_DAMPMIN</code> , <code>XSLP_DAMPSHRINK</code> , <code>XSLP_DAMPSTART</code> , <code>XSLP_SAMEDAMP</code>
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_DAMPMAX

Description	Maximum value for the damping factor of a variable during dynamic damping
Type	Double
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 <code>XSLP_SAMEDAMP</code> successive changes in the same direction for a variable, then the damping factor (<i>D</i>) for the variable will be reset to $D * XSLP_DAMPEXPAND + XSLP_DAMPMAX * (1 - XSLP_DAMPEXPAND)$
Default value	1
See also	Xpress-SLP Solution Process, <code>XSLP_ALGORITHM</code> , <code>XSLP_DAMP</code> , <code>XSLP_DAMPEXPAND</code> , <code>XSLP_DAMPMIN</code> , <code>XSLP_DAMPSHRINK</code> , <code>XSLP_DAMPSTART</code> , <code>XSLP_SAMEDAMP</code>
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_DAMPMIN

Description	Minimum value for the damping factor of a variable during dynamic damping
Type	Double
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 (<i>D</i>) for the variable will be reset to $D * XSLP_DAMPSHRINK + XSLP_DAMPMIN * (1 - XSLP_DAMPEXPAND)$
Default value	1
See also	Xpress-SLP Solution Process, <code>XSLP_ALGORITHM</code> , <code>XSLP_DAMP</code> , <code>XSLP_DAMPEXPAND</code> , <code>XSLP_DAMPMAX</code> , <code>XSLP_DAMPSHRINK</code> , <code>XSLP_DAMPSTART</code>
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_DAMPSHRINK

Description	Multiplier to decrease damping factor during dynamic damping
Type	Double
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

Description	Default initial value for an SLP variable if none is explicitly given
Type	Double
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	1000
Affects routines	XSLPconstruct

XSLP_DEFAULTSTEPBOUND

Description	Minimum initial value for the step bound of an SLP variable if none is explicitly given
Type	Double
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

Description	Absolute perturbation of values for calculating numerical derivatives
Type	Double
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: $XSLP_DELTA_A + abs(X) * XSLP_DELTA_R$ 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_R

Description	Relative perturbation of values for calculating numerical derivatives
Type	Double
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: $XSLP_DELTA_A + abs(X) * XSLP_DELTA_R$ 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

Description	Minimum absolute value of delta coefficients to be retained
Type	Double
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

Description Zero tolerance used when calculating numerical derivatives

Type Double

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 numerical 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_DELTACOST

Description Initial penalty cost multiplier for penalty delta vectors

Type Double

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

Description Factor for increasing cost multiplier on total penalty delta vectors

Type Double

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

Description Maximum penalty cost multiplier for penalty delta vectors

Type Double

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 infinite

Affects routines XSLPmaxim, XSLPminim

See also XSLP_AUGMENTATION, XSLP_DELTACOST, XSLP_DELTACOSTFACTOR, XSLP_ERRORCOST

XSLP_DJTOL

Description Tolerance on DJ value for determining if a variable is at its step bound

Type Double

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_ECFTOL_A

Description Absolute tolerance on testing feasibility at the point of linearization

Type Double

Notes 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 $\max(abs(Positive), abs(Negative)) * XSLP_{ECFTOL_R}$. 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.

Default value	.01
Affects routines	XSLPmaxim, XSLPminim
See also	Convergence criteria, XSLP_ECFCHECK, XSLP_ECFCOUNT, XSLP_ECFTOL_R

XSLP_ECFTOL_R

Description	Relative tolerance on testing feasibility at the point of linearization
Type	Double
Notes	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 $\max(abs(Positive), abs(Negative)) * XSLP_{ECFTOL_R}$. 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.
Default value	.01
Affects routines	XSLPmaxim, XSLPminim
See also	Convergence criteria, XSLP_ECFCHECK, XSLP_ECFCOUNT, XSLP_ECFTOL_A

XSLP_EQTOL_A

Description	Absolute tolerance on equality testing in logical functions
Type	Double
Note	If two values A and B are within $XSLP_{EQTOL_A}$ and $abs(A) * XSLP_{EQTOL_R}$ then they are regarded as equal by the logical functions.
Default value	0.00001
Affects routines	EQ, GE, GT, NE, LE, LT
See also	XSLP_EQTOL_R

XSLP_EQTOL_R

Description	Relative tolerance on equality testing in logical functions
Type	Double
Note	If two values A and B are within $XSLP_{EQTOL_A}$ and $abs(A) * XSLP_{EQTOL_R}$ then they are regarded as equal by the logical functions.

Default value	0.00001
Affects routines	EQ, GE, GT, NE, LE, LT
See also	XSLP_EQTOL_A

XSLP_ERRORCOST

Description	Initial penalty cost multiplier for penalty error vectors
Type	Double
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

Description	Factor for increasing cost multiplier on total penalty error vectors
Type	Double
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

Description	Maximum penalty cost multiplier for penalty error vectors
Type	Double
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	infinite
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_AUGMENTATION, XSLP_DELTACOST, XSLP_ERRORCOST, XSLP_ERRORCOSTFACTOR

XSLP_ERRORTOL_A

Description	Absolute tolerance for error vectors
Type	Double
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 <i>XSLP_ERRORTOL_A</i> ; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than <i>XSLP_EVTOL_A</i> ; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than <i>XSLP_EVTOL_R * Obj</i> where <i>Obj</i> is the current objective function value.
Default value	0.001
Affects routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>
See also	<i>XSLP_EVTOL_A</i> , <i>XSLP_EVTOL_R</i>

XSLP_ERRORTOL_P

Description	Absolute tolerance for printing error vectors
Type	Double
Note	The solution log includes a print of penalty delta and penalty error vectors with an activity greater than <i>XSLP_ERRORTOL_P</i> .
Default value	0.0001
Affects routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>

XSLP_ESCALATION

Description	Factor for increasing cost multiplier on individual penalty error vectors
Type	Double
Note	If penalty cost escalation is activated in <i>XSLP_ALGORITHM</i> then the penalty cost multiplier will be increased by a factor of <i>XSLP_ESCALATION</i> for any active error vector up to a maximum of <i>XSLP_MAXWEIGHT</i> .
Default value	1.25
Affects routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>
See also	<i>XSLP_ALGORITHM</i> , <i>XSLP_DELTACOST</i> , <i>XSLP_ERRORCOST</i> , <i>XSLP_MAXWEIGHT</i>

XSLP_ETOL_A

Description	Absolute tolerance on penalty vectors
Type	Double
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 <i>XSLP_ETOL_A</i> or less than <i>Positive * XSLP_ETOL_R</i> or less than <i>abs(Negative) * XSLP_ETOL_R</i> then it will be regarded as insignificant and will not have its penalty increased.
Default value	0.0001
Affects routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>
See also	<i>XSLP_ETOL_R</i> <i>XSLP_DELTACOST</i> , <i>XSLP_ERRORCOST</i> , <i>XSLP_ESCALATION</i>

XSLP_ETOL_R

Description	Relative tolerance on penalty vectors
Type	Double
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 <i>XSLP_ETOL_A</i> or less than <i>Positive * XSLP_ETOL_R</i> or less than <i>abs(Negative) * XSLP_ETOL_R</i> then it will be regarded as insignificant and will not have its penalty increased.
Default value	0.0001
Affects routines	<i>XSLPmaxim</i> , <i>XSLPminim</i>
See also	<i>XSLP_ETOL_A</i> <i>XSLP_DELTACOST</i> , <i>XSLP_ERRORCOST</i> , <i>XSLP_ESCALATION</i>

XSLP_EVTOL_A

Description	Absolute tolerance on total penalty costs
Type	Double
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 <i>XSLP_ERRORTOL_A</i> ; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than <i>XSLP_EVTOL_A</i> ; the sum of the cost contributions from all the penalty error and penalty delta vectors is less than <i>XSLP_EVTOL_R * Obj</i> where <i>Obj</i> is the current objective function value.

Default value	0.001
Affects routines	XSLPmaxim , XSLPminim
See also	XSLP_ERRORTOL_A , XSLP_EVTOL_R

XSLP_EVTOL_R

Description	Relative tolerance on total penalty costs
Type	Double
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 * <i>Obj</i> where <i>Obj</i> is the current objective function value.
Default value	0.001
Affects routines	XSLPmaxim , XSLPminim
See also	XSLP_ERRORTOL_A , XSLP_EVTOL_A

XSLP_EXPAND

Description	Multiplier to increase a step bound
Type	Double
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 (<i>B</i>) for the variable will be reset to $B * XSLP_EXPAND$.
Default value	2
Affects routines	XSLPmaxim , XSLPminim
See also	XSLP_SHRINK , XSLP_SAMECOUNT

XSLP_GRANULARITY

Description	Base for calculating penalty costs
Type	Double

Note	If <code>XSLP_GRANULARITY > 1</code> , then initial penalty costs will be powers of <code>XSLP_GRANULARITY</code> .
Default value	4
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_MAXWEIGHT</code> , <code>XSLP_MINWEIGHT</code>

XSLP_INFINITY

Description	Value returned by a divide-by-zero in a formula
Type	Double
Default value	1.0e+10

XSLP_ITOL_A

Description	Absolute impact convergence tolerance
Type	Double
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 $\delta E < XSLP_ITOL_A$ 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.

Default value	0.001
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_ITOL_R</code> , <code>XSLP_MTOL_A</code> , <code>XSLP_MTOL_R</code> , <code>XSLP_STOL_A</code> , <code>XSLP_STOL_R</code>

XSLP_ITOL_R

Description	Relative impact convergence tolerance
Type	Double
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.

Default value	0.001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_ITOL_A, XSLP_MTOL_A, XSLP_MTOL_R, XSLP_STOL_A, XSLP_STOL_R

XSLP_MAXWEIGHT

Description	Maximum penalty weight for delta or error vectors
--------------------	---

Type	Double
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
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	1.5
See also	Memory control variables <code>XSLP_MEM*</code>

XSLP_MINWEIGHT

Description	Minimum penalty weight for delta or error vectors
Type	Double
Note	When penalty vectors are created, the minimum weight that will be used is given by <code>XSLP_MINWEIGHT</code> .
Default value	0.01
Affects routines	<code>XSLPconstruct</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_AUGMENTATION</code> , <code>XSLP_MAXWEIGHT</code>

XSLP_MIPCUTOFF_A

Description	Absolute objective function cutoff for MIP termination
Type	Double

Note	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 <code>XSLP_MIPCUTOFFCOUNT</code> SLP iterations are all worse than the best obtained so far, and the difference is greater than <code>XSLP_MIPCUTOFF_A</code> and $OBJ * XSLP_MIPCUTOFF_R$ where <i>OBJ</i> is the best integer solution obtained so far. The MIP cutoff tests are only applied after <code>XSLP_MIPCUTOFF_LIMIT</code> SLP iterations at the current node.
Default value	0.0001
Affects routines	<code>XSLPglocal</code>
See also	<code>XSLP_MIPCUTOFF_COUNT</code> , <code>XSLP_MIPCUTOFF_LIMIT</code> , <code>XSLP_MIPCUTOFF_R</code>

XSLP_MIPCUTOFF_R

Description	Absolute objective function cutoff for MIP termination
Type	Double
Note	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 <code>XSLP_MIPCUTOFFCOUNT</code> SLP iterations are all worse than the best obtained so far, and the difference is greater than <code>XSLP_MIPCUTOFF_A</code> and $OBJ * XSLP_MIPCUTOFF_R$ where <i>OBJ</i> is the best integer solution obtained so far. The MIP cutoff tests are only applied after <code>XSLP_MIPCUTOFF_LIMIT</code> SLP iterations at the current node.
Default value	0.0001
Affects routines	<code>XSLPglocal</code>
See also	<code>XSLP_MIPCUTOFF_COUNT</code> , <code>XSLP_MIPCUTOFF_LIMIT</code> , <code>XSLP_MIPCUTOFF_A</code>

XSLP_MIPERRORTOL_A

Description	Absolute penalty error cost tolerance for MIP cut-off
Type	Double
Note	The penalty error cost test is applied at each node where there are active penalties in the solution. If <code>XSLP_MIPERRORTOL_A</code> is nonzero and the absolute value of the penalty costs is greater than <code>XSLP_MIPERRORTOL_A</code>, the node will be declared infeasible. If <code>XSLP_MIPERRORTOL_A</code> is zero then no test is made and the node will not be declared infeasible on this criterion.
Default value	0 (inactive)
Affects routines	<code>XSLPglocal</code>
See also	<code>XSLP_MIPERRORTOL_R</code>

XSLP_MIPERRORTOL_R

Description	Relative penalty error cost tolerance for MIP cut-off
Type	Double
Note	The penalty error cost test is applied at each node where there are active penalties in the solution. If <code>XSLP_MIPERRORTOL_R</code> is nonzero and the absolute value of the penalty costs is greater than $XSLP_MIPERRORTOL_R * abs(Obj)$ where <i>Obj</i> is the value of the objective function, then the node will be declared infeasible. If <code>XSLP_MIPERRORTOL_R</code> is zero then no test is made and the node will not be declared infeasible on this criterion.
Default value	0 (inactive)
Affects routines	<code>XSLPglobal</code>
See also	<code>XSLP_MIPERRORTOL_A</code>

XSLP_MIPOTOL_A

Description	Absolute objective function tolerance for MIP termination
Type	Double
Note	The objective function test for MIP termination is applied only when step bounding has been applied (or <code>XSLP_SBSTART</code> 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 <code>XSLP_MIPOCOUNT</code> SLP iterations is within <code>XSLP_MIPOTOL_A</code> or within $OBJ * XSLP_MIPOTOL_R$ where <i>OBJ</i> is the average value of the objective function over those iterations.
Default value	0.0001
Affects routines	<code>XSLPglobal</code>
See also	<code>XSLP_MIPOCOUNT</code> <code>XSLP_MIPOTOL_R</code> <code>XSLP_SBSTART</code>

XSLP_MIPOTOL_R

Description	Relative objective function tolerance for MIP termination
Type	Double
Note	The objective function test for MIP termination is applied only when step bounding has been applied (or <code>XSLP_SBSTART</code> 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 <code>XSLP_MIPOCOUNT</code> SLP iterations is within <code>XSLP_MIPOTOL_A</code> or within $OBJ * XSLP_MIPOTOL_R$ where <i>OBJ</i> is the average value of the objective function over those iterations.
Default value	0.0001

Affects routines XSLPglobal
See also XSLP_MIPOCOUNT XSLP_MIPOTOL_A XSLP_SBSTART

XSLP_MTOL_A

Description Absolute effective matrix element convergence tolerance
Type Double
Note The absolute effective matrix element convergence criterion assesses the change in the effect of a coefficient in a constraint. 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

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

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.

Default value 0.001
Affects routines XSLPmaxim, XSLPminim
See also XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_R, XSLP_STOL_A, XSLP_STOL_R

XSLP_MTOL_R

Description Relative effective matrix element convergence tolerance
Type Double

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.

Default value 0.001

Affects routines XSLPmaxim, XSLPminim

See also XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_A, XSLP_STOL_A, XSLP_STOL_R

XSLP_MVTOL

Description Marginal value tolerance for determining if a constraint is slack

Type Double

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 *active constraint* when identifying unconverged variables in active constraints.

Default value 0.001

Affects routines XSLPmaxim, XSLPminim

See also XSLP_STOL_A, XSLP_STOL_R

XSLP_OBJSENSE

Description	Objective function sense
Type	Double
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 XSLPopt.
Set by routines	XSLPmaxim, XSLPminim
Default value	+1
Affects routines	XSLPmaxim, XSLPminim, XSLPopt

XSLP_OBJTOPENALTYCOST

Description	Factor to estimate initial penalty costs from objective function
Type	Double
Notes	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. 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. 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.
Default value	0
Affects routines	XSLPmaxim, XSLPminim

XSLP_OTOL_A

Description	Absolute static objective (2) convergence tolerance
Type	Double
Note	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 <code>XSLP_MVTOL</code>) 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 <i>Iter</i> is the <code>XSLP_OCOUNT</code> most recent SLP iterations and <i>Obj</i> is the corresponding objective function value. If $ABS(\delta Obj) \leq XSLP_OTOL_A$ then the problem has converged on the absolute static objective (2) convergence criterion. The static objective function (2) test is applied only if <code>XSLP_OCOUNT</code> is at least 2.
Default value	0.001
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_OCOUNT</code> , <code>XSLP_OTOL_R</code>

XSLP_OTOL_R

Description	Relative static objective (2) convergence tolerance
Type	Double
Note	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 <code>XSLP_MVTOL</code>) 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 <i>Iter</i> is the <code>XSLP_OCOUNT</code> most recent SLP iterations and <i>Obj</i> 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 <code>XSLP_OCOUNT</code> is at least 2.

Default value	0.00001
Affects routines	XSLPmaxim , XSLPminim
See also	XSLP_OCOUNT , XSLP_OTOL_A

XSLP_PRESOLVEZERO

Description	Minimum absolute value for a variable which is identified as nonzero during SLP presolve
Type	Double
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_SHRINK

Description	Multiplier to reduce a step bound
Type	Double
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 (<i>B</i>) for the variable will be reset to $B * \text{XSLP_SHRINK}$. 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_STOL_A

Description	Absolute slack convergence tolerance
Type	Double
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.
Default value	0.001

Affects routines XSLPmaxim, XSLPminim

See also XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_A, XSLP_MTOL_R, XSLP_STOL_R

XSLP_STOL_R

Description Relative slack convergence tolerance

Type Double

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.

Default value 0.001

Affects routines XSLPmaxim, XSLPminim

See also XSLP_ITOL_A, XSLP_ITOL_R, XSLP_MTOL_A, XSLP_MTOL_R, XSLP_STOL_A

XSLP_VALIDATIONTOL_A

Description Absolute tolerance for the XSLPvalidate procedure

Type Double

Note 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 $D < XSLP_VALIDATIONTOL_A$
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 $D < MAX(ABS(TPos), ABS(TNeg)) * XSLP_VALIDATIONTOL_A$
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.
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.

Default value 0.0001

Affects routines XSLPvalidate

See also XSLP_VALIDATIONINDEX_A, XSLP_VALIDATIONINDEX_R, XSLP_VALIDATIONTOL_R

XSLP_VALIDATIONTOL_R

Description	Relative tolerance for the XSLPvalidate procedure
Type	Double
Note	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 $D < \text{MAX}(\text{ABS}(TPos), \text{ABS}(TNeg)) * \text{XSLP_VALIDATIONTOL_R}$ 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 smaller factor is printed in the validation report. 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.
Default value	0.0001
Affects routines	XSLPvalidate
See also	XSLP_VALIDATIONINDEX_A, XSLP_VALIDATIONINDEX_R, XSLP_VALIDATIONTOL_A

XSLP_VTOL_A

Description	Absolute static objective (3) convergence tolerance
Type	Double
Note	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. The variation in the objective function is defined as $\delta Obj = \text{MAX}_{Iter}(Obj) - \text{MIN}_{Iter}(Obj)$ where $Iter$ is the XSLP_VCOUNT most recent SLP iterations and Obj is the corresponding objective function value. If $\text{ABS}(\delta Obj) \leq \text{XSLP_VTOL_A}$ then the problem has converged on the absolute static objective function (3) criterion. The static objective function (3) test is applied only 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.

Default value 0.001

Affects routines XSLPmaxim, XSLPminim

See also XSLP_SBSTART, XSLP_VCOUNT, XSLP_VLIMIT, XSLP_VTOL_R

XSLP_VTOL_R

Description Relative static objective (3) convergence tolerance

Type Double

Note 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. 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.

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

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.

Default value 0.001

Affects routines XSLPmaxim, XSLPminim

See also XSLP_SBSTART, XSLP_VCOUNT, XSLP_VLIMIT, XSLP_VTOL_A

XSLP_WTOL_A

Description Absolute extended convergence continuation tolerance

Type Double

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:

$$\delta Obj = Obj - LastConvergedObj$$

For a minimization problem, the sign is reversed.

If $\delta Obj > XSLP_WTOL_A$ and

$\delta Obj > ABS(ConvergedObj) * XSLP_WTOL_R$ 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.

Default value 0.0001

Affects routines `XSLPmaxim`, `XSLPminim`

See also `XSLP_WCOUNT`, `XSLP_WTOL_R`

XSLP_WTOL_R

Description Relative extended convergence continuation tolerance

Type Double

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:

$$\delta Obj = Obj - LastConvergedObj$$

For a minimization problem, the sign is reversed.

If $\delta Obj > XSLP_WTOL_A$ and

$\delta Obj > ABS(ConvergedObj) * XSLP_WTOL_R$ 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.

Default value 0.0001

Affects routines `XSLPmaxim`, `XSLPminim`

See also `XSLP_WCOUNT`, `XSLP_WTOL_A`

XSLP_XTOL_A

Description Absolute static objective function (1) tolerance

Type Double

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	0.001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_XCOUNT, XSLP_XLIMIT, XSLP_XTOL_R

XSLP_XTOL_R

Description	Relative static objective function (1) tolerance
Type	Double
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 <i>Iter</i> is the XSLP_XCOUNT most recent SLP iterations and <i>Obj</i> 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	0.001
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_XCOUNT, XSLP_XLIMIT, XSLP_XTOL_A

XSLP_ZERO

Description	Absolute zero tolerance
--------------------	-------------------------

Type	Double
Note	If a value is below <code>XSLP_ZERO</code> in magnitude, then it will be regarded as zero in certain formula calculations: - 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-10
Affects routines	<code>XSLPevaluatecoef</code> , <code>XSLPevaluateformula</code>

9.2 Integer control parameters

XSLP_ALGORITHM

Description	Bit map describing the SLP algorithm(s) to be used	
Type	Integer	
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.

Notes	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. 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. 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. 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. 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. 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 <code>XSLP_DELTA_Z</code> 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.

Recommended setting: Bits 1, 2, 5, 7 and usually bits 8 and 9.

Default value	166 (sets bits 1, 2, 5, 7)
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_DELTA_Z</code> , <code>XSLP_ERRORMAXCOST</code> , <code>XSLP_ESCALATION</code>

XSLP_AUGMENTATION

Description	Bit map describing the SLP augmentation method(s) to be used
Type	Integer

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.

Notes

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.

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.

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.

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.

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.

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.

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

Description Frequency with which to save the model

Type Integer

Note A value of zero means that the model will not automatically be saved. A positive value of *n* will save the model every *n* SLP iterations.

Default value 0

Affects routines `XSLPmaxim`, `XSLPminim`

XSLP_BARLIMIT

Description Number of initial SLP iterations using the barrier method

Type Integer

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 2

Affects routines `XSLPmaxim`, `XSLPminim`

XSLP_CASCADE

Description Bit map describing the cascading to be used

Type Integer

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 <code>XPRS_FEASTOL</code> .
	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 <code>XPRS_FEASTOL</code> .
	4	Apply cascading to all SLP variables which are structural.
Note	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.	
Default value	1	
Affects routines	<code>XSLPcascade</code>	

XSLP_CASCADENLIMIT

Description	Maximum number of iterations for cascading with non-linear determining rows
Type	Integer
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 <code>XSLP_CASCADENLIMIT</code>. 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	<code>XSLPcascade</code>
See also	<code>XSLP_CASCADE</code>

XSLP_CONTROL

Description	Bit map describing which Xpress-SLP functions also activate the corresponding Optimizer Library function
Type	Integer

Values	<table> <tr> <th>Bit</th><th>Meaning</th></tr> <tr> <td>0</td><td>Xpress-SLP problem management functions do NOT invoke the corresponding Optimizer Library function for the underlying linear problem.</td></tr> <tr> <td>1</td><td><code>XSLPcopycontrols</code> does NOT invoke <code>XPRScopycontrols</code>.</td></tr> <tr> <td>2</td><td><code>XSLPcopycallbacks</code> does NOT invoke <code>XPRScopycallbacks</code>.</td></tr> <tr> <td>3</td><td><code>XSLPcopyprob</code> does NOT invoke <code>XPRScopyprob</code>.</td></tr> <tr> <td>4</td><td><code>XSLPsetdefaults</code> does NOT invoke <code>XPRSsetdefaults</code>.</td></tr> <tr> <td>5</td><td><code>XSLPsave</code> does NOT invoke <code>XPRSsave</code>.</td></tr> <tr> <td>6</td><td><code>XSLPrestore</code> does NOT invoke <code>XPRSrestore</code>.</td></tr> </table>	Bit	Meaning	0	Xpress-SLP problem management functions do NOT invoke the corresponding Optimizer Library function for the underlying linear problem.	1	<code>XSLPcopycontrols</code> does NOT invoke <code>XPRScopycontrols</code> .	2	<code>XSLPcopycallbacks</code> does NOT invoke <code>XPRScopycallbacks</code> .	3	<code>XSLPcopyprob</code> does NOT invoke <code>XPRScopyprob</code> .	4	<code>XSLPsetdefaults</code> does NOT invoke <code>XPRSsetdefaults</code> .	5	<code>XSLPsave</code> does NOT invoke <code>XPRSsave</code> .	6	<code>XSLPrestore</code> does NOT invoke <code>XPRSrestore</code> .
Bit	Meaning																
0	Xpress-SLP problem management functions do NOT invoke the corresponding Optimizer Library function for the underlying linear problem.																
1	<code>XSLPcopycontrols</code> does NOT invoke <code>XPRScopycontrols</code> .																
2	<code>XSLPcopycallbacks</code> does NOT invoke <code>XPRScopycallbacks</code> .																
3	<code>XSLPcopyprob</code> does NOT invoke <code>XPRScopyprob</code> .																
4	<code>XSLPsetdefaults</code> does NOT invoke <code>XPRSsetdefaults</code> .																
5	<code>XSLPsave</code> does NOT invoke <code>XPRSsave</code> .																
6	<code>XSLPrestore</code> does NOT invoke <code>XPRSrestore</code> .																
Note	The problem management functions are: <code>XSLPcopyprob</code> to copy from an existing problem; <code>XSLPcopycontrols</code> and <code>XSLPcopycallbacks</code> to copy the current controls and callbacks from an existing problem; <code>XSLPsetdefaults</code> to reset the controls to their default values; <code>XSLPsave</code> and <code>XSLPrestore</code> for saving and restoring a problem.																
Default value	0 (no bits set)																
Affects routines	<code>XSLPcopycontrols</code> , <code>XSLPcopycallbacks</code> , <code>XSLPcopyprob</code> , <code>XSLPrestore</code> , <code>XSLPsave</code> , <code>XSLPsetdefaults</code>																

XSLP_DAMPSTART

Description	SLP iteration at which damping is activated
Type	Integer
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	<code>XSLPmaxim</code> , <code>XSLPmaxim</code>
See also	<code>XSLP_ALGORITHM</code> , <code>XSLP_DAMPEXPAND</code> , <code>XSLP_DAMPMAX</code> , <code>XSLP_DAMPMIN</code> , <code>XSLP_DAMP SHRINK</code>

XSLP_DCLIMIT

Description	Default iteration delay for delayed constraints
Type	Integer
Note	If a delayed constraint does not have an explicit delay, then the value of <code>XSLP_DCLIMIT</code> will be used.
Default value	5
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_DCLOG

Description	Amount of logging information for activation of delayed constraints
Type	Integer
Note	If XSLP_DCLOG is set to 1, then a message will be produced for each DC as it is activated.
Default value	0
Affects routines	XSLPmaxim, XSLPminim

XSLP_DELAYUPDATEROWS

Description	Number of SLP iterations before update rows are fully activated
Type	Integer
Notes	Update rows are an integral part of the augmented matrix used to create linear approximations of the nonlinear problem. However, if determining rows are present, then it is possible for some updated values to be calculated during cascading, and the corresponding update rows are then not required. When SLP variables have explicit bounds, and particularly when step bounding is enforced, update rows become important to the solutions actually obtained. It is therefore normal practice to delay update rows for only a few initial SLP iterations. Update rows can only be delayed for variables which are not structural (that is, they do not have explicit coefficients in the original problem) and for which determining rows are provided.
Default value	2
Affects routines	XSLPmaxim, XSLPminim

XSLP_DECOMPOSE

Description	Bitmap controlling the action of function XSLPdecompose
Type	Integer

Values	<table> <tr> <th>Bit</th><th>Meaning</th></tr> <tr> <td>0</td><td>(=1) Set to 1 to activate automatic decomposition during problem augmentation</td></tr> <tr> <td>1</td><td>(=2) Only decompose formulae which are entirely linear (default is to extract any linear constituents)</td></tr> <tr> <td>2</td><td>(=4) Decompose formulae in any fixed column (default is to decompose only formulae in the "equals column")</td></tr> <tr> <td>3</td><td>(=8) Only extract structural columns – that is, columns which already have coefficients in the problem (default is to extract any column which appears linearly)</td></tr> <tr> <td>4</td><td>(=16) Treat fixed variables as constants when deciding linearity (default is to treat all variables as non-constant)</td></tr> <tr> <td>5</td><td>(=32) Do not decompose coefficients in columns which are fixed to zero (default is to decompose coefficients in all eligible columns)</td></tr> </table>	Bit	Meaning	0	(=1) Set to 1 to activate automatic decomposition during problem augmentation	1	(=2) Only decompose formulae which are entirely linear (default is to extract any linear constituents)	2	(=4) Decompose formulae in any fixed column (default is to decompose only formulae in the "equals column")	3	(=8) Only extract structural columns – that is, columns which already have coefficients in the problem (default is to extract any column which appears linearly)	4	(=16) Treat fixed variables as constants when deciding linearity (default is to treat all variables as non-constant)	5	(=32) Do not decompose coefficients in columns which are fixed to zero (default is to decompose coefficients in all eligible columns)
Bit	Meaning														
0	(=1) Set to 1 to activate automatic decomposition during problem augmentation														
1	(=2) Only decompose formulae which are entirely linear (default is to extract any linear constituents)														
2	(=4) Decompose formulae in any fixed column (default is to decompose only formulae in the "equals column")														
3	(=8) Only extract structural columns – that is, columns which already have coefficients in the problem (default is to extract any column which appears linearly)														
4	(=16) Treat fixed variables as constants when deciding linearity (default is to treat all variables as non-constant)														
5	(=32) Do not decompose coefficients in columns which are fixed to zero (default is to decompose coefficients in all eligible columns)														
Notes	Bit 0 of <code>XSLP_DECOMPOSE</code> must be set for automatic decomposition during problem augmentation (<code>XSLPconstruct</code>). This decomposition happens after SLP presolving (<code>XSLPpresolve</code>). <code>XSLP_PRESOLVE</code> can be set to fix any variables that it finds, which may allow more decomposition to take place. The remaining bits of <code>XSLP_DECOMPOSE</code> apply whether decomposition is automatic or called explicitly through <code>XSLPdecompose</code>.														
Default value	0														
Affects routines	<code>XSLPconstruct</code> , <code>XSLPdecompose</code>														

XSLP_DECOMPOSEPASSLIMIT

Description	Maximum number of repeats of presolve+decompose
Type	Integer
Notes	If <code>XSLP_DECOMPOSEPASSLIMIT</code> is set to a positive integer, and formula decomposition is activated (either by setting <code>XSLP_DECOMPOSE</code> or by calling <code>XSLPdecompose</code> directly), then the SLP presolve procedure will be activated after decomposition is completed. If any changes are made to the problem as a result of presolving, then decomposition + presolve will be repeated (up to <code>XSLP_DECOMPOSEPASSLIMIT</code> times) as long as the problem continues to be changed. <code>XSLP_PRESOLVE</code> should be set to 15 to ensure that bounds are changed for any columns identified by <code>XSLPpresolve</code> as fixed
Default value	0
Affects routines	<code>XSLPdecompose</code>
See also	<code>XSLP_DECOMPOSE</code> , <code>XSLP_PRESOLVE</code> ,

XSLP_DELTAOFFSET

Description	Position of first character of SLP variable name used to create name of delta vector
Type	Integer

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

Description	Number of SLP iterations during which to apply <code>XSLP_DELTA_Z</code>
Type	Integer
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.
Default value	0
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_DELTA_Z</code>

XSLP_DERIVATIVES

Description	Method of calculating derivatives
Type	Integer
Values	0 numerical approximation (finite differences) 1 analytic derivatives where possible
Notes	Partial derivatives are used to create the linear approximation to the nonlinear problem. The default method of calculating derivatives is by perturbing the value of each variable in turn and calculating derivatives from the changes in the values of the coefficients. The alternative is to use analytic derivatives, in which formulae are derived for the derivatives of each non-constant coefficient. These formulae are then evaluated directly to obtain the derivatives. Analytic derivatives cannot be used for formulae involving discontinuous functions (such as the logical functions <code>EQ</code>, <code>LT</code>, etc). 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>LOG</code> of small values).
Default value	0
Affects routines	<code>XSLPconstruct</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_ECFCHECK

Description	Check feasibility at the point of linearization for extended convergence criteria
Type	Integer
Values	0 no check (extended criteria are always used); 1 check until one infeasible constraint is found; 2 check all constraints.
Notes	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. 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. 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. The number of infeasible constraints found at the point of linearization is returned in XSLP_ECFCOUNT.
Default value	1
Affects routines	Convergence criteria, XSLPmaxim, XSLPminim
See also	XSLP_ECFCOUNT, XSLP_ECFTOL_A, XSLP_ECFTOL_R

XSLP_ERROROFFSET

Description	Position of first character of constraint name used to create name of penalty error vectors
Type	Integer
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

Description	Evaluation strategy for user functions
Type	Integer
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-SLP to estimate the value of the function from its derivatives if the new point of evaluation is sufficiently close to the original. Setting <code>XSLP_EVALUATE</code> 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	<code>XSLPevaluatecoef</code> , <code>XSLPevaluateformula</code>
See also	<code>XSLP_FUNCEVAL</code>

XSLP_EXCELVISIBLE

Description	Display of Excel when evaluating user functions written in Excel
Type	Integer
Values	0 do not display; 1 display.
Notes	Normally, Excel is hidden when used as the source of user functions. This is generally more efficient because (for example) no screen updating is required. During model development, or if Excel is being used for visualization, it may be appropriate to have Excel displayed. <code>XSLP_EXCELVISIBLE</code> must be set before any user function written in Excel is called.
Default value	0
Affects routines	<code>XSLPevaluatecoef</code> <code>XSLPevaluateformula</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_EXTRACVS

Description	Expansion number for character variables
Type	Integer
Note	The expansion number is the number of additional items for which space is provided in memory. Before any items are loaded, it is the initial space available. After any items have been loaded, it is the amount by which the space will be expanded if required. The expansion number may be increased by the system beyond any value set by the user. Setting the expansion number is one way of increasing efficiency during loading or adding character variables.

Set by routines	XSLPaddcvars , XSLPchgcvvar , XSLPloadcvars
Default value	10
Affects routines	XSLPaddcvars , XSLPchgcvvar , XSLPloadcvars , XSLPreadprob
See also	XSLP_MEM_CVAR , XSLP_MEMORYFACTOR

XSLP_EXTRAUFS

Description	Expansion number for user functions
Type	Integer
Note	The expansion number is the number of additional items for which space is provided in memory. Before any items are loaded, it is the initial space available. After any items have been loaded, it is the amount by which the space will be expanded if required. The expansion number may be increased by the system beyond any value set by the user. Setting the expansion number is one way of increasing efficiency during loading or adding user function definitions.
Set by routines	XSLPadduserfuncs , XSLPchguserfunc , XSLPloaduserfuncs
Default value	10
Affects routines	XSLPadduserfuncs , XSLPchguserfunc , XSLPloaduserfuncs XSLPreadprob
See also	XSLP_MEM_UF , XSLP_MEMORYFACTOR

XSLP_EXTRAXVITEMS

Description	Expansion number for XV items
Type	Integer
Note	The expansion number is the number of additional items for which space is provided in memory. Before any items are loaded, it is the initial space available. After any items have been loaded, it is the amount by which the space will be expanded if required. The expansion number may be increased by the system beyond any value set by the user. Setting the expansion number is one way of increasing efficiency during loading or adding XVs or XV items.
Set by routines	XSLPaddxvs , XSLPchg xvitem , XSLPloadxvs
Default value	100
Affects routines	XSLPaddxvs , XSLPchg xvitem , XSLPloadxvs XSLPreadprob
See also	XSLP_MEM_XVITEM , XSLP_MEMORYFACTOR

XSLP_EXTRAXVS

Description	Expansion number for XVs
Type	Integer
Note	The expansion number is the number of additional items for which space is provided in memory. Before any items are loaded, it is the initial space available. After any items have been loaded, it is the amount by which the space will be expanded if required. The expansion number may be increased by the system beyond any value set by the user. Setting the expansion number is one way of increasing efficiency during loading or adding XVs.
Set by routines	XSLPaddxvs , XSLPchg xv , XSLPloadxvs
Default value	100
Affects routines	XSLPaddxvs , XSLPchg xv , XSLPloadxvs XSLPreadprob
See also	XSLP_MEM_XV , XSLP_MEMORYFACTOR

XSLP_FUNCEVAL

Description	Bit map for determining the method of evaluating user functions and their derivatives	
Type	Integer	
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.	
Default value	0	
Affects routines	XSLPevaluatecoef , XSLPevaluateformula	
See also	XSLP_EVALUATE	

XSLP_INFEASLIMIT

Description	The maximum number of consecutive infeasible SLP iterations which can occur before Xpress-SLP terminates
Type	Integer
Note	An infeasible solution to an SLP iteration means that it 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. <code>XSLP_INFEASLIMIT</code> 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	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_ITERLIMIT

Description	The maximum number of SLP iterations
Type	Integer
Note	If Xpress-SLP reaches <code>XSLP_ITERLIMIT</code> 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	<code>XSLPglocal</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_LOG

Description	Level of printing during SLP iterations												
Type	Integer												
Values	<table><tr><td>-1</td><td>none</td></tr><tr><td>0</td><td>minimal</td></tr><tr><td>1</td><td>normal: iteration, penalty vectors</td></tr><tr><td>2</td><td>omit from convergence log any variables which have converged</td></tr><tr><td>3</td><td>omit from convergence log any variables which have already converged</td></tr><tr><td>4</td><td>include all variables in convergence log</td></tr></table>	-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	4	include all variables in convergence log
-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												
4	include all variables in convergence log												
Default value	1												
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>												

XSLP_MAXTIME

Description	The maximum time in seconds that the SLP optimization will run before it terminates
Type	Integer
Notes	The (elapsed) time is measured from the beginning of the first SLP optimization. If <code>XSLP_MAXTIME</code> is negative, Xpress-SLP will terminate after $(-XSLP_MAXTIME)$ seconds. If it is positive, Xpress-SLP will terminate in MISLP after <code>XSLP_MAXTIME</code> seconds or as soon as an integer solution has been found thereafter.
Default value	0
Affects routines	<code>XSLPglocal</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_MIPALGORITHM

Description	Bitmap describing the MISLP algorithms to be used	
Type	Integer	
Values	Bit	Meaning
	0	Solve initial SLP to convergence.
	1	Re-solve final SLP to convergence.
	2	Relax step bounds according to <code>XSLP_MIPRELAXSTEPBOUNDS</code> after initial node.
	3	Fix step bounds according to <code>XSLP_MIPFIXSTEPBOUNDS</code> after initial node.
	4	Relax step bounds according to <code>XSLP_MIPRELAXSTEPBOUNDS</code> at each node.
	5	Fix step bounds according to <code>XSLP_MIPFIXSTEPBOUNDS</code> at each node.
	6	Limit iterations at each node to <code>XSLP_MIPITERLIMIT</code> .
	7	Relax step bounds according to <code>XSLP_MIPRELAXSTEPBOUNDS</code> after MIP solution is found.
	8	Fix step bounds according to <code>XSLP_MIPFIXSTEPBOUNDS</code> 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 `XSLP_MIPALGORITHM` determines the strategy of `XSLPglocal` for solving MINLP problems. The recommended approach is to solve the problem first without reference to the global 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.

Some versions of the optimizer re-run the initial node as part of the global 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.

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). When the optimal node has been found, it is possible to fix the global variables and then re-optimize with SLP. Step bounds can be relaxed or fixed for this optimization as well (use bits 7 and 8). 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"). Another strategy is to solve the problem to convergence ignoring the nature of the global variables. Then, fixing the linearization, use MIP to find the optimal setting of the global variables. Then, fixing the global variables, but varying the linearization, solve to convergence. Set bit 10 to implement this strategy ("SLP then MIP"). For mode details about MISLP algorithms and strategies, see the separate section.

Default value 17 (bits 0 and 4)

Affects routines `XSLPglobal`

See also `XSLP_ALGORITHM`, `XSLP_MIPFIXSTEPBOUNDS`, `XSLP_MIPITERLIMIT`, `XSLP_MIPRELAXSTEPBOUNDS`

XSLP_MIPCUTOFFCOUNT

Description Number of SLP iterations to check when considering a node for cutting off

Type Integer

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 `XSLPglobal`

See also `XSLP_MIPCUTOFF_A`, `XSLP_MIPCUTOFF_R`, `XSLP_MIPCUTOFFLIMIT`

XSLP_MIPCUTOFFLIMIT

Description Number of SLP iterations to check when considering a node for cutting off

Type Integer

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 <code>XSLP_MIPCUTOFFCOUNT</code> SLP iterations are all worse than the best obtained so far, and the difference is greater than <code>XSLP_MIPCUTOFF_A</code> and <code>OBJ * XSLP_MIPCUTOFF_R</code> where <code>OBJ</code> is the best integer solution obtained so far. The test is not applied until at least <code>XSLP_MIPCUTOFFLIMIT</code> SLP iterations have been carried out at the current node.
Default value	5
Affects routines	<code>XSLPglobal</code>
See also	<code>XSLP_MIPCUTOFF_A</code> , <code>XSLP_MIPCUTOFF_R</code> , <code>XSLP_MIPCUTOFFCOUNT</code>

XSLP_MIPDEFAULTALGORITHM

Description	Default algorithm to be used during the global search in MISLP
Type	Integer
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 (primal simplex)
Affects routines	<code>XSLPglobal</code>
See also	<code>XPRS_DEFAULTALG</code> , <code>XSLP_MIPALGORITHM</code>

XSLP_MIPFIXSTEPBOUNDS

Description	Bitmap describing the step-bound fixing strategy during MISLP										
Type	Integer										
Values	<table> <tr> <th>Bit</th><th>Meaning</th></tr> <tr> <td>0</td><td>Fix step bounds on structural SLP variables which are not in coefficients.</td></tr> <tr> <td>1</td><td>Fix step bounds on all structural SLP variables.</td></tr> <tr> <td>2</td><td>Fix step bounds on SLP variables appearing only in coefficients.</td></tr> <tr> <td>3</td><td>Fix step bounds on SLP variables appearing in coefficients.</td></tr> </table>	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.
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>XSLPglobal</code>										
See also	<code>XSLP_MIPALGORITHM</code> , <code>XSLP_MIPRELAXSTEPBOUNDS</code>										

XSLP_MIPITERLIMIT

Description	Maximum number of SLP iterations at each node
Type	Integer
Note	If bit 6 of <code>XSLP_MIPALGORITHM</code> is set, then the number of iterations at each node will be limited to <code>XSLP_MIPITERLIMIT</code> .
Default value	0
Affects routines	<code>XSLPglobal</code>
See also	<code>XSLP_ITERLIMIT</code> , <code>XSLP_MIPALGORITHM</code>

XSLP_MIPLOG

Description	Frequency with which MIP status is printed
Type	Integer
Note	If <code>XSLP_MIPLOG</code> is set to a positive integer, then the current MIP status (node number, best value, best bound) is printed every <code>XSLP_MIPLOG</code> nodes.
Default value	0 (no printing)
Affects routines	<code>XSLPglobal</code>
See also	<code>XSLP_LOG</code> , <code>XSLP_SLPLOG</code>

XSLP_MIPOCOUNT

Description	Number of SLP iterations at each node over which to measure objective function variation
Type	Integer
Note	The objective function test for MIP termination is applied only when step bounding has been applied (or <code>XSLP_SBSTART</code> 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 <code>XSLP_MIPOCOUNT</code> SLP iterations is within <code>XSLP_MIPOTOL_A</code> or within $OBJ * XSLP_MIPOTOL_R$ where <i>OBJ</i> is the average value of the objective function over those iterations.
Default value	5
Affects routines	<code>XSLPglobal</code>
See also	<code>XSLP_MIPOTOL_A</code> <code>XSLP_MIPOTOL_R</code> <code>XSLP_SBSTART</code>

XSLP_MIPRELAXSTEPBOUNDS

Description	Bitmap describing the step-bound relaxation strategy during MISLP	
Type	Integer	
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 [XSLPglobal](#)

See also [XSLP_MIPALGORITHM](#), [XSLP_MIPFIXSTEPBOUNDS](#)

XSLP_OCOUNT

Description Number of SLP iterations over which to measure objective function variation for static objective (2) convergence criterion

Type Integer

Note 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. 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.

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

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.

Default value 5

Affects routines `XSLPmaxim`, `XSLPminim`
See also `XSLP_OTOL_A` `XSLP_OTOL_R`

XSLP_PENALTYINFOSTART

Description Iteration from which to record row penalty information
Type Integer
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
Affects routines `XSLProwinfo`

XSLP_PRESOLVE

Description Bitmap indicating the SLP presolve actions to be taken
Type Integer
Values Bit Meaning
 0 Activate SLP presolve.
 1 Explicitly fix columns identified as fixed to zero.
 2 Explicitly fix columns all columns identified as fixed.
 3 SLP bound tightening4 MISLP bound tightening.
 8 Do not presolve coefficients.
 9 Do not remove delta variables.
Note The Xpress-SLP nonlinear presolve (which is carried out once, before augmentation) is independent of the Optimizer presolve (which is carried out during each SLP iteration).
Default value 0 (no presolve)
Affects routines `XSLPconstruct`, `XSLPpresolve`
See also `XSLP_PRESOLVEPASSLIMIT`

XSLP_PRESOLVEPASSLIMIT

Description Maximum number of passes through the problem to improve SLP bounds
Type Integer

Note	The Xpress-SLP nonlinear presolve (which is carried out once, before augmentation) is independent of the Optimizer presolve (which is carried out during each SLP iteration). The procedure carries out a number of passes through the SLP problem, seeking to tighten implied bounds or to identify fixed values. <code>XSLP_PRESOLVEPASSLIMIT</code> can be used to change the maximum number of passes carried out.
Default value	20
Affects routines	<code>XSLPpresolve</code>
See also	<code>XSLP_PRESOLVE</code>

XSLP_SAMECOUNT

Description	Number of steps reaching the step bound in the same direction before step bounds are increased
Type	Integer
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 <code>XSLP_SAMECOUNT</code> successive changes reaching the step bound and in the same direction for a variable, then the step bound (<i>B</i>) for the variable will be reset to $B * XSLP_EXPAND$.
Default value	3
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_EXPAND</code>

XSLP_SAMEDAMP

Description	Number of steps in same direction before damping factor is increased
Type	Integer
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 <code>XSLP_SAMEDAMP</code> successive changes in the same direction for a variable, then the damping factor (<i>D</i>) for the variable will be reset to $D * XSLP_DAMPEXPAND + XSLP_DAMPMAX * (1 - XSLP_DAMPEXPAND)$
Default value	3
See also	Xpress-SLP Solution Process, <code>XSLP_ALGORITHM</code> , <code>XSLP_DAMP</code> , <code>XSLP_DAMPMAX</code>
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_SBROWOFFSET

Description	Position of first character of SLP variable name used to create name of SLP lower and upper step bound rows
Type	Integer
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

Description	SLP iteration after which step bounds are first applied
Type	Integer
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

Values	0	No re-scaling.
	1	Re-scale every SLP iteration up to <code>XSLP_SCALECOUNT</code> iterations after the end of barrier optimization.
	2	Re-scale every SLP iteration up to <code>XSLP_SCALECOUNT</code> 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 <code>XSLP_SCALECOUNT</code> SLP iterations.
	6	Re-scale every <code>XSLP_SCALECOUNT</code> 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	<code>XSLPmaxim</code> , <code>XSLPminim</code>	
See also	<code>XSLP_SCALECOUNT</code>	

XSLP_SCALECOUNT

Description	Iteration limit used in determining when to re-scale the SLP matrix
Type	Integer
Notes	If <code>XSLP_SCALE</code> is set to 1 or 2, then <code>XSLP_SCALECOUNT</code> 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	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_SCALE</code>

XSLP_SLPLOG

Description	Frequency with which SLP status is printed
Type	Integer
Note	If <code>XSLP_LOG</code> is set to zero (minimal logging) then a nonzero value for <code>XSLP_SLPLOG</code> defines the frequency (in SLP iterations) when summary information is printed out.
Default value	0
Affects routines	<code>XSLPglobal</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_LOG</code> , <code>XSLP_MIPLOG</code>

XSLP_STOPOUTOFRANGE

Description	Stop optimization and return error code if internal function argument is out of range
Type	Integer
Note	If <code>XSLP_STOPOUTOFRANGE</code> is set to 1, then if an internal function receives an argument which is out of its allowable range (for example, <i>LOG</i> of a negative number), an error code is set and the optimization is terminated.
Default value	0
Affects routines	<code>XSLPevaluatecoef</code> , <code>XSLPevaluateformula</code> <code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_TIMEPRINT

Description	Print additional timings during SLP optimization
Type	Integer
Note	Date and time printing can be useful for identifying slow procedures during the SLP optimization. Setting <code>XSLP_TIMEPRINT</code> to 1 prints times at additional points during the optimization.
Default value	0
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_UNFINISHEDLIMIT

Description	Number of times within one SLP iteration that an unfinished LP optimization will be continued
Type	Integer
Note	If the optimization of the current linear approximation terminates with an "unfinished" status (for example, because it has reached maximum iterations), Xpress-SLP will attempt to continue using the primal simplex algorithm. This process will be repeated for up to <code>XSLP_UNFINISHEDLIMIT</code> successive LP optimizations within any one SLP iteration. If the limit is reached, Xpress-SLP will terminate with <code>XSLP_STATUS</code> set to <code>XSLP_LPUNFINISHED</code>
Default value	3
Affects routines	<code>XSLPglocal</code> , <code>XSLPmaxim</code> , <code>XSLPminim</code>

XSLP_UPDATEOFFSET

Description	Position of first character of SLP variable name used to create name of SLP update row
Type	Integer
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 <code>XSLP_UPDATEFORMAT</code> to define a format and <code>XSLP_UPDATEOFFSET</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_UPDATEFORMAT</code>

XSLP_VCOUNT

Description	Number of SLP iterations over which to measure static objective (3) convergence
Type	Integer
Note	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. The variation in the objective function is defined as $\delta Obj = MAX_{Iter}(Obj) - MIN_{Iter}(Obj)$ where <i>Iter</i> is the <code>XSLP_VCOUNT</code> most recent SLP iterations and <i>Obj</i> is the corresponding objective function value. If $ABS(\delta Obj) \leq XSLP_VTOL_A$ then the problem has converged on the absolute static objective function (3) criterion. The static objective function (3) test is applied only if after at least <code>XSLP_VLIMIT</code> + <code>XSLP_SBSTART</code> SLP iterations have taken place and only if <code>XSLP_VCOUNT</code> 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.
Default value	0
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_SBSTART</code> , <code>XSLP_VLIMIT</code> , <code>XSLP_VTOL_A</code> , <code>XSLP_VTOL_R</code>

XSLP_VLIMIT

Description	Number of SLP iterations after which static objective (3) convergence testing starts
Type	Integer
Note	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. 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.

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

then the problem has converged on the absolute static objective function (3) criterion. The static objective function (3) test is applied only 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.

Default value	0
Affects routines	XSLPmaxim, XSLPminim
See also	XSLP_SBSTART, XSLP_VCOUNT, XSLP_VTOL_A, XSLP_VTOL_R

XSLP_WCOUNT

Description	Number of SLP iterations over which to measure the objective for the extended convergence continuation criterion
Type	Integer
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: $\delta Obj = Obj - LastConvergedObj$

For a minimization problem, the sign is reversed.

If $\delta Obj > XSLP_WTOL_A$ and

$\delta Obj > ABS(ConvergedObj) * XSLP_WTOL_R$ 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

Description Number of SLP iterations over which to measure static objective (1) convergence

Type Integer

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	0
Affects routines	<code>XSLPmaxim</code> , <code>XSLPminim</code>
See also	<code>XSLP_XLIMIT</code> , <code>XSLP_XTOL_A</code> , <code>XSLP_XTOL_R</code>

XSLP_XLIMIT

Description	Number of SLP iterations up to which static objective (1) convergence testing starts
Type	Integer
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 <i>Iter</i> is the <code>XSLP_XCOUNT</code> most recent SLP iterations and <i>Obj</i> 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 <code>XSLP_XLIMIT</code> 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	0

Affects routines XSLPmaxim, XSLPminim

See also XSLP_XCOUNT, XSLP_XTOL_A, XSLP_XTOL_R

XSLP_ZEROCRITERION

Description	Bitmap determining the behavior of the placeholder deletion procedure	
Type	Integer	
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 <i>Management of zero placeholder entries</i> .	
Default value	0	
Affects routines	XSLPmaxim, XSLPminim	
See also	XSLP_ZEROCRITERIONCOUNT, XSLP_ZEROCRITERIONSTART, <i>Management of zero placeholder entries</i>	

XSLP_ZEROCRITERIONCOUNT

Description	Number of consecutive times a placeholder entry is zero before being considered for deletion
Type	Integer
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_ZEROCRITERIONSTART, <i>Management of zero placeholder entries</i>

XSLP_ZEROCRITERIONSTART

Description	SLP iteration at which criteria for deletion of placeholder entries are first activated.
Type	Integer
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>

9.3 Memory control parameters

Memory control parameters are integer controls which can be used to define a minimum number of items for which space should be provided. For example, to allow space for at least 5000 coefficients, set `XSLP_MEM_COEF` to 5000.

Normally, Xpress-SLP will expand the memory required for items as the number grows. However, this process can be inefficient in the use of available memory and can, in any case, take time. If the system runs out of memory, then an error message will be produced and normally a list of current memory requirements will be printed. Alternatively, the library function `XSLPuprintmemory` can be used to print the memory currently in use. The following is an example of the information produced:

Arrays and dimensions:					
Array	Item Size	Used Items	Max Items	Allocated Memory	Memory Control
MemList	28	103	129	4K	
String	1	206891	219888	215K	XSLP_MEM_STRING
Xv	16	1282	2000	32K	XSLP_MEM_XV
Xvitem	48	1382	1600	75K	XSLP_MEM_XVITEM
UserFunc	80	2	1000	79K	XSLP_MEM_UF
IntlFunc	80	45	48	4K	
Vars	136	1685	2000	266K	XSLP_MEM_VAR
Coef	40	4631	4633	181K	XSLP_MEM_COEF
Formula	48	1415	2000	94K	XSLP_MEM_FORMULA
ToknStak	16	10830	13107	205K	XSLP_MEM_STACK
Cols	48	8163	8192	384K	XSLP_MEM_COL
Rows	40	4596	5120	200K	XSLP_MEM_ROW
Xrows	48	1607	2000	94K	XSLP_MEM_XROW
FormValu	16	3182	3184	50K	XSLP_MEM_FORMULAVALUE
XPRSrow	4	12883	13155	52K	
XPRScol	4	12883	13155	52K	
XPRSccoef	8	12883	13155	103K	
XPRSsetyp	1	12883	13155	13K	
CalcStak	24	1	1000	24K	XSLP_MEM_CALCSTACK
XPRSrhwr	4	1492	1494	6K	
XPRSrhel	8	1492	1494	12K	

Used Items is the number of items actually in use; *Max Items* is the number currently allocated, which is reflected in the *Allocated Memory* figure. Where there is an option to change the size of the allocation, the name of the memory control parameter is given. So, for example, to set the initial size of the *Xrows* array to 1650, use the following:

```
XSLPsetintcontrol(Prob, XSLP_MEM_XROW, 1650);
```

This will have two effects: the array will be allocated from the start with 1650 items, so there will be no need to expand the array as items are loaded or created; the array will be large enough to hold the items required but will have less unused space, so there will be more memory available for other arrays if necessary.

The following is a list of the memory control parameters that can be set with an indication of the type of array for which they are used. The current value can be retrieved using `XSLPgetintcontrol`, or the full set can be listed using `XSLPuprintmemory`.

XSLP_MEM_CALCSTACK

Description	Memory allocation for formula calculations
--------------------	--

Type	Integer
------	---------

XSLP_MEM_COEF

Description	Memory allocation for nonlinear coefficients
-------------	--

Type	Integer
------	---------

XSLP_MEM_COL

Description	Memory allocation for additional information on matrix columns
-------------	--

Type	Integer
------	---------

XSLP_MEM_CVAR

Description	Memory allocation for character variables
-------------	---

Type	Integer
------	---------

XSLP_MEM_DERIVATIVES

Description	Memory allocation for analytic derivatives
-------------	--

Type	Integer
------	---------

XSLP_MEM_EXCELDOUBLE

Description	Memory allocation for return values from Excel user functions
-------------	---

Type	Integer
------	---------

XSLP_MEM_FORMULA

Description	Memory allocation for formulae
-------------	--------------------------------

Type	Integer
------	---------

XSLP_MEM_FORMULAHASH

Description Memory allocation for internal formula array

Type Integer

XSLP_MEM_FORMULAVALUE

Description Memory allocation for formula values and derivatives

Type Integer

XSLP_MEM_ITERLOG

Description Memory allocation for SLP iteration summary

Type Integer

XSLP_MEM_RETURNARRAY

Description Memory allocation for return values from multi-valued user function

Type Integer

XSLP_MEM_ROW

Description Memory allocation for additional information on matrix rows

Type Integer

XSLP_MEM_STACK

Description Memory allocation for parsed formulae, analytic derivatives

Type Integer

XSLP_MEM_STRING

Description Memory allocation for strings of all types
Type Integer

XSLP_MEM_TOL

Description Memory allocation for tolerance sets
Type Integer

XSLP_MEM_UF

Description Memory allocation for user functions
Type Integer

XSLP_MEM_VALSTACK

Description Memory allocation for intermediate values for analytic derivatives
Type Integer

XSLP_MEM_VAR

Description Memory allocation for SLP variables
Type Integer

XSLP_MEM_XF

Description Memory allocation for complicated functions
Type Integer

XSLP_MEM_XFNAMES

Description Memory allocation for complicated function input and return names
Type Integer

XSLP_MEM_XFVALUE

Description Memory allocation for complicated function values
Type Integer

XSLP_MEM_XROW

Description Memory allocation for extended row information
Type Integer

XSLP_MEM_XV

Description Memory allocation for XVs
Type Integer

XSLP_MEM_XVITEM

Description Memory allocation for individual XV entries
Type Integer

9.4 String control parameters

XSLP_CVNAME

Description	Name of the set of character variables to be used
Type	String
Notes	This variable may be required for input from a file using <code>XSLPreadprob</code> if there is more than one set of character variables in the file. If no name is set, then the first set of character variables will be used, and the name will be set accordingly. This variable may also be required for output using <code>XSLPwriteprob</code> where character variables 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> , <code>XSLP_TOLNAME</code>

XSLP_DELTAFORMAT

Description	Formatting string for creation of names for SLP delta vectors
Type	String
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 <code>XSLP_DELTAOFFSET</code>.
Default value	<code>pD_%s</code> where p is a unique prefix for names in the current problem
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_DELTAOFFSET</code> , <code>XSLP_UNIQUEPREFIX</code> <code>XSLPsetuniqueprefix</code>

XSLP_IVNAME

Description	Name of the set of initial values to be used
Type	String
Notes	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. 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.

Set by routines	<code>XSLPreadprob</code>
Default value	<code>none</code>
Affects routines	<code>XSLPreadprob</code> , <code>XSLPwriteprob</code>
See also	<code>XSLP_CVNAME</code> , <code>XSLP_SBNAME</code> , <code>XSLP_TOLNAME</code>

XSLP_MINUSDELTAFORMAT

Description	Formatting string for creation of names for SLP negative penalty delta vectors
Type	String
Note	This control can be used to create a specific naming structure for negative 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 <code>XSLP_DELTAOFFSET</code> .
Default value	<code>pD-%s</code> where p is a unique prefix for names in the current problem
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_DELTAOFFSET</code> , <code>XSLP_UNIQUEPREFIX</code> <code>XSLPsetuniqueprefix</code>

XSLP_MINUSERERRORFORMAT

Description	Formatting string for creation of names for SLP negative penalty error vectors
Type	String
Note	This control can be used to create a specific naming structure for negative 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 <code>XSLP_ERROROFFSET</code> .
Default value	<code>pE-%s</code> where p is a unique prefix for names in the current problem
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_ERROROFFSET</code> , <code>XSLP_UNIQUEPREFIX</code> <code>XSLPsetuniqueprefix</code>

XSLP_PENALTYCOLFORMAT

Description	Formatting string for creation of the names of the SLP penalty transfer vectors
Type	String

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 "DELT" 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
See also	XSLP_UNIQUEPREFIX XSLPsetuniqueprefix

XSLP_PENALTYROWFORMAT

Description	Formatting string for creation of the names of the SLP penalty rows
Type	String
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 "DELT" 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
See also	XSLP_UNIQUEPREFIX XSLPsetuniqueprefix

XSLP_PLUDELTAFORMAT

Description	Formatting string for creation of names for SLP positive penalty delta vectors
Type	String
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_UNIQUEPREFIX XSLPsetuniqueprefix

XSLP_PLUSERERRORFORMAT

Description	Formatting string for creation of names for SLP positive penalty error vectors
Type	String
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 <code>XSLP_ERROROFFSET</code> .
Default value	pE+%s where p is a unique prefix for names in the current problem
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_ERROROFFSET</code> , <code>XSLP_UNIQUEPREFIX</code> <code>XSLPsetuniqueprefix</code>

XSLP_SBLOROWFORMAT

Description	Formatting string for creation of names for SLP lower step bound rows
Type	String
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 <code>XSLP_SBROWOFFSET</code> .
Default value	pSB-%s where p is a unique prefix for names in the current problem
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_SBROWOFFSET</code> , <code>XSLP_UNIQUEPREFIX</code> <code>XSLPsetuniqueprefix</code>

XSLP_SBNAME

Description	Name of the set of initial step bounds to be used
Type	String
Notes	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. 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.
Set by routines	<code>XSLPreadprob</code>
Default value	none

Affects routines `XSLPreadprob`, `XSLPwriteprob`

See also `XSLP_CVNAME`, `XSLP_IVNAME`, `XSLP_TOLNAME`

XSLP_SBUPROWFORMAT

Description	Formatting string for creation of names for SLP upper step bound rows
Type	String
Note	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 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_SBROWOFFSET</code> .
Default value	pSB+%s where p is a unique prefix for names in the current problem
Affects routines	<code>XSLPconstruct</code>
See also	<code>XSLP_SBROWOFFSET</code> , <code>XSLP_UNIQUEPREFIX</code> <code>XSLPsetuniqueprefix</code>

XSLP_TOLNAME

Description	Name of the set of tolerance sets to be used
Type	String
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_CVNAME</code> , <code>XSLP_IVNAME</code> , <code>XSLP_SBNAME</code>

XSLP_UPDATEFORMAT

Description	Formatting string for creation of names for SLP update rows
Type	String

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> , <code>XSLP_UNIQUEPREFIX</code> <code>XSLPsetuniqueprefix</code>

Chapter 10

Library functions and the programming interface

10.1 Counting

All Xpress-SLP entities are numbered from 1. The 0th item is defined, and is an empty entity of the appropriate type. Therefore, whenever an Xpress-SLP function returns a zero value, it means that there is no data of that type.

In parsed and unparsed function arrays, the indices always count from 1. This includes types `XSLP_VAR` and `XSLP_CONSTRAINT`: the index is the matrix column or row index + `XPRS_CSTYLE`.

Note that for *input* of function arrays, types `XSLP_COL` and `XSLP_ROW` can be used; these do honor the setting of `XPRS_CSTYLE`, but will be converted into standard `XSLP_VAR` or `XSLP_CONSTRAINT` references. When a function array is returned from Xpress-SLP, the `XSLP_VAR` or `XSLP_CONSTRAINT` type will always be used.

Matrix row and column indices do honor the setting of `XPRS_CSTYLE`: that is, they count from zero if `XPRS_CSTYLE=1`, and from 1 if `XPRS_CSTYLE=0`.

Xpress-SLP functions affected by `XPRS_CSTYLE`

Function	Arguments affected
XSLPaddcoefs	RowIndex, ColIndex
XSLPadddc	RowIndex
XSLPaddvars	ColIndex, DetRow
XSLPchgcoef	RowIndex, ColIndex
XSLPchgdc	RowIndex
XSLPchgrow	RowIndex
XSLPchgvar	ColIndex, DetRow
XSLPevaluatecoef	RowIndex, ColIndex
XSLPgetcoef	RowIndex, ColIndex
XSLPgetdc	RowIndex
XSLPgetrow	RowIndex
XSLPgetvar	ColIndex, DetRow, Delta, PenaltyDelta, UpdateRow
XSLPloadcoefs	RowIndex, ColIndex
XSLPloaddc	RowIndex
XSLPloadvars	ColIndex, DetRow

also user callbacks defined by:

XSLPsetcbcascadevar	ColIndex
XSLPsetcbitervar	ColIndex

and all functions which have arguments which are indices into arrays.

10.2 The Xpress-SLP problem pointer

Xpress-SLP 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-SLP 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-SLP 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-SLP problems. See the documentation

below on the individual functions for more details.

`XSLPcopycontrols`(XSLPprob probl, 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-SLP problem

`XSLPdestroyprob`(XSLPprob probl)
Delete an Xpress-SLP problem from memory

`XSLPrestore`(XSLPprob probl)
Restore Xpress-SLP data structures from file

`XSLPsave`(XSLPprob probl)
Save Xpress-SLP data structures to file

10.3 The XSLPload... functions

The `XSLPload...` functions can be used to load an Xpress-SLP 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 linear part of the problem should be loaded first, using the normal Optimizer Library functions `XPRSloadlp` or `XPRSloadglobal`. Then the appropriate parts of the Xpress-SLP problem can be loaded. After all the `XSLPload...` 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-SLP 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 `XSLPadd...` function.

Xpress-SLP is compatible with the Xpress quadratic programming optimizer. `XPRSloadqp` and `XPRSloadqglobal` can be used to load quadratic problems. The quadratic objective will be optimized using the Xpress quadratic optimizer; the nonlinear constraints will be handled with the normal SLP procedures.

10.4 Library functions

A large number of routines are available for Library users of Xpress-SLP, 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.

<code>XSLPaddcoefs</code>	Add non-linear coefficients to the SLP problem	p. 139
<code>XSLPaddcvars</code>	Add character variables (CVs) to the SLP problem	p. 141
<code>XSLPadddcs</code>	Add delayed constraints (DCs) to the SLP problem	p. 142

XSLPadddfs	Add a set of distribution factors	p. 144
XSLPaddivfs	Add a set of initial value formulae	p. 145
XSLPaddnames	Set the names of a set of SLP entities in an SLP problem.	p. 147
XSLPaddtolsets	Add sets of standard tolerance values to an SLP problem	p. 148
XSLPadduserfuncs	Add user function definitions to an SLP problem.	p. 149
XSLPaddvars	Add SLP variables defined as matrix columns to an SLP problem	p. 151
XSLPaddxvs	Add a set of extended variable arrays (XVs) to an SLP problem	p. 153
XSLPcalluserfunc	Call a user function from a program or from within another user function	p. 155
XSLPcascade	Re-calculate consistent values for SLP variables. based on the current values of the remaining variables	p. 156
XSLPcascadeorder	Establish a re-calculation sequence for SLP variables with determining rows.	p. 157
XSLPchgccoef	Add or change a single matrix coefficient using a character string for the formula	p. 158
XSLPchgcoef	Add or change a single matrix coefficient using a parsed or unparsed formula	p. 159
XSLPchgcvvar	Add or change the value of the character string corresponding to an SLP character variable	p. 161
XSLPchgdc	Add or change the settings for a delayed constraint (DC)	p. 162
XSLPchgdf	Set or change a distribution factor	p. 164
XSLPchgfuncobject	Change the address of one of the objects which can be accessed by the user functions	p. 165
XSLPchgivf	Set or change the initial value formula for a variable	p. 166
XSLPchgrow	Change the status setting of a constraint	p. 167
XSLPchgrowwt	Set or change the initial penalty error weight for a row	p. 168
XSLPchgtolset	Add or change a set of convergence tolerances used for SLP variables	p. 169
XSLPchguserfunc	Add or change a user function in an SLP problem after the problem has been input	p. 170
XSLPchguserfuncaddress	Change the address of a user function	p. 172
XSLPchguserfuncobject	Change or define one of the objects which can be accessed by the user functions	p. 173
XSLPchgvar	Define a column as an SLP variable or change the characteristics and values of an existing SLP variable	p. 174
XSLPchg xv	Add or change an extended variable array (XV) in an SLP problem	p. 176
XSLPchg xvitem	Add or change an item of an existing XV in an SLP problem	p. 177

XSLPconstruct	Create the full augmented SLP matrix and data structures, ready for optimization	p. 179
XSLPcopycallbacks	Copy the user-defined callbacks from one SLP problem to another	p. 180
XSLPcopycontrols	Copy the values of the control variables from one SLP problem to another	p. 181
XSLPcopyprob	Copy an existing SLP problem to another	p. 182
XSLPcreateprob	Create a new SLP problem	p. 183
XSLPdecompose	Decompose nonlinear constraints into linear and nonlinear parts	p. 184
XSLPdestroyprob	Delete an SLP problem and release all the associated memory	p. 185
XSLPevaluatecoef	Evaluate a coefficient using the current values of the variables	p. 186
XSLPevaluateformula	Evaluate a formula using the current values of the variables	p. 187
XSLPformatvalue	Format a double-precision value in the style of Xpress-SLP	p. 188
XSLPfree	Free any memory allocated by Xpress-SLP and close any open Xpress-SLP files	p. 189
XSLPgetbanner	Retrieve the Xpress-SLP banner and copyright messages	p. 190
XSLPgetcccoef	Retrieve a single matrix coefficient as a formula in a character string	p. 191
XSLPgetcoef	Retrieve a single matrix coefficient as a formula split into tokens	p. 192
XSLPgetcvar	Retrieve the value of the character string corresponding to an SLP character variable	p. 193
XSLPgetdblattrib	Retrieve the value of a double precision problem attribute	p. 194
XSLPgetdblcontrol	Retrieve the value of a double precision problem control	p. 195
XSLPgetdc	Retrieve information about a delayed constraint in an SLP problem	p. 196
XSLPgetdf	Get a distribution factor	p. 197
XSLPgetdtime	Retrieve a double precision time stamp in seconds	p. 198
XSLPgetfuncinfo	Retrieve the argument information for a user function	p. 199
XSLPgetfuncinfoV	Retrieve the argument information for a user function	p. 200
XSLPgetfuncobject	Retrieve the address of one of the objects which can be accessed by the user functions	p. 201
XSLPgetfuncobjectV	Retrieve the address of one of the objects which can be accessed by the user functions	p. 202
XSLPgetindex	Retrieve the index of an Xpress-SLP entity with a given name	p. 203
XSLPgetintattrib	Retrieve the value of an integer problem attribute	p. 204
XSLPgetintcontrol	Retrieve the value of an integer problem control	p. 205
XSLPgetivf	Get the initial value formula for a variable	p. 206

<code>XSLPgetlasterror</code>	Retrieve the error message corresponding to the last Xpress-SLP error during an SLP run	p. 207
<code>XSLPgetmessagetype</code>	Retrieve the message type corresponding to a message number	p. 208
<code>XSLPgetnames</code>	Retrieve the names of a set of Xpress-SLP entities	p. 209
<code>XSLPgetparam</code>	Retrieve the value of a control parameter or attribute by name	p. 210
<code>XSLPgetptrattrib</code>	Retrieve the value of a problem pointer attribute	p. 211
<code>XSLPgetrow</code>	Retrieve the status setting of a constraint	p. 212
<code>XSLPgetrowwt</code>	Get the initial penalty error weight for a row	p. 213
<code>XSLPgetslpsol</code>	Obtain the solution values for the most recent SLP iteration	p. 214
<code>XSLPgetstrattrib</code>	Retrieve the value of a string problem attribute	p. 215
<code>XSLPgetstrcontrol</code>	Retrieve the value of a string problem control	p. 216
<code>XSLPgetstring</code>	Retrieve the value of a string in the Xpress-SLP string table	p. 217
<code>XSLPgettime</code>	Retrieve an integer time stamp in seconds and/or milliseconds	p. 218
<code>XSLPgettolset</code>	Retrieve the values of a set of convergence tolerances for an SLP problem	p. 219
<code>XSLPgetuserfunc</code>	Retrieve the type and parameters for a user function	p. 220
<code>XSLPgetuserfuncaddress</code>	Retrieve the address of a user function	p. 222
<code>XSLPgetuserfuncobject</code>	Retrieve the address of one of the objects which can be accessed by the user functions	p. 223
<code>XSLPgetvar</code>	Retrieve information about an SLP variable	p. 224
<code>XSLPgetversion</code>	Retrieve the Xpress-SLP major and minor version numbers	p. 226
<code>XSLPgetxv</code>	Retrieve information about an extended variable array	p. 227
<code>XSLPgetxvitem</code>	Retrieve information about an item in an extended variable array	p. 228
<code>XSLPglobal</code>	Initiate the Xpress-SLP mixed integer SLP (MISLP) algorithm	p. 230
<code>XSLPinit</code>	Initializes the Xpress-SLP system	p. 231
<code>XSLPitemname</code>	Retrieves the name of an Xpress-SLP entity or the value of a function token as a character string.	p. 232
<code>XSLPloadcoefs</code>	Load non-linear coefficients into the SLP problem	p. 233
<code>XSLPloadcvars</code>	Load character variables (CVs) into the SLP problem	p. 235
<code>XSLPloadcdcs</code>	Load delayed constraints (DCs) into the SLP problem	p. 236
<code>XSLPloaddfs</code>	Load a set of distribution factors	p. 238
<code>XSLPloaddivfs</code>	Load a set of initial value formulae	p. 239
<code>XSLPloadtolsets</code>	Load sets of standard tolerance values into an SLP problem	p. 241
<code>XSLPloaduserfuncs</code>	Load user function definitions into an SLP problem.	p. 242

<code>XSLPloadvars</code>	Load SLP variables defined as matrix columns into an SLP problem p. 244
<code>XSLPloadxvs</code>	Load a set of extended variable arrays (XVs) into an SLP problem p. 246
<code>XSLPmaxim</code>	Maximize an SLP problem p. 248
<code>XSLPminim</code>	Minimize an SLP problem p. 249
<code>XSLPopt</code>	Maximize or minimize an SLP problem p. 250
<code>XSLPparsecformula</code>	Parse a formula written as a character string into internal parsed (reverse Polish) format p. 251
<code>XSLPparseformula</code>	Parse a formula written as an unparsed array of tokens into internal parsed (reverse Polish) format p. 252
<code>XSLPpreparseformula</code>	Perform an initial scan of a formula written as a character string, identifying the operators but not attempting to identify the types of the individual tokens p. 253
<code>XSLPpresolve</code>	Perform a nonlinear presolve on the problem p. 254
<code>XSLPprintmsg</code>	Print a message string according to the current settings for Xpress-SLP output p. 255
<code>XSLPqparse</code>	Perform a quick parse on a free-format character string, identifying where each token starts p. 256
<code>XSLPreadprob</code>	Read an Xpress-SLP extended MPS format matrix from a file into an SLP problem p. 257
<code>XSLPpremaxim</code>	Continue the maximization of an SLP problem p. 258
<code>XSLPpreminim</code>	Continue the minimization of an SLP problem p. 259
<code>XSLPrestore</code>	Restore the Xpress-SLP problem from a file created by XSLPsave p. 260
<code>XSLPrevise</code>	Revise the unaugmented SLP matrix with data from a file p. 261
<code>XSLProwinfo</code>	Get or set row information p. 262
<code>XSLPsave</code>	Save the Xpress-SLP problem to file p. 263
<code>XSLPsaveas</code>	Save the Xpress-SLP problem to a named file p. 264
<code>XSLPscaling</code>	Analyze the current matrix for largest/smallest coefficients and ratios p. 265
<code>XSLPsetcbcascadeend</code>	Set a user callback to be called at the end of the cascading process, after the last variable has been cascaded p. 266
<code>XSLPsetcbcascadestart</code>	Set a user callback to be called at the start of the cascading process, before any variables have been cascaded p. 267
<code>XSLPsetcbcascadevar</code>	Set a user callback to be called after each column has been cascaded p. 268
<code>XSLPsetcbcascadevarF</code>	Set a user callback to be called after each column has been cascaded p. 270
<code>XSLPsetcbconstruct</code>	Set a user callback to be called during the Xpress-SLP augmentation process p. 272

<code>XSLPsetcbdestroy</code>	Set a user callback to be called when an SLP problem is about to be destroyed	p. 274
<code>XSLPsetcbformula</code>	Set a callback to be used in formula evaluation when an unknown token is found	p. 275
<code>XSLPsetcbintsol</code>	Set a user callback to be called during MISLP when an integer solution is obtained	p. 277
<code>XSLPsetcbiterend</code>	Set a user callback to be called at the end of each SLP iteration	p. 278
<code>XSLPsetcbiterstart</code>	Set a user callback to be called at the start of each SLP iteration	p. 279
<code>XSLPsetcbitervar</code>	Set a user callback to be called after each column has been tested for convergence	p. 280
<code>XSLPsetcbitervarF</code>	Set a user callback to be called after each column has been tested for convergence	p. 282
<code>XSLPsetcbmessage</code>	Set a user callback to be called whenever Xpress-SLP outputs a line of text	p. 284
<code>XSLPsetcbmessageF</code>	Set a user callback to be called whenever Xpress-SLP outputs a line of text	p. 286
<code>XSLPsetcboptnode</code>	Set a user callback to be called during MISLP when an optimal SLP solution is obtained at a node	p. 288
<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	p. 289
<code>XSLPsetcbslpend</code>	Set a user callback to be called at the end of the SLP optimization	p. 291
<code>XSLPsetcbslpnode</code>	Set a user callback to be called during MISLP after the SLP optimization at each node.	p. 292
<code>XSLPsetcbslpstart</code>	Set a user callback to be called at the start of the SLP optimization	p. 293
<code>XSLPsetdblcontrol</code>	Set the value of a double precision problem control	p. 294
<code>XSLPsetdefaultcontrol</code>	Set the values of one SLP control to its default value	p. 295
<code>XSLPsetdefaults</code>	Set the values of all SLP controls to their default values	p. 296
<code>XSLPsetfuncobject</code>	Change the address of one of the objects which can be accessed by the user functions	p. 297
<code>XSLPsetfunctionerror</code>	Set the function error flag for the problem	p. 298
<code>XSLPsetintcontrol</code>	Set the value of an integer problem control	p. 299
<code>XSLPsetlogfile</code>	Define an output file to be used to receive messages from Xpress-SLP	p. 300
<code>XSLPsetparam</code>	Set the value of a control parameter by name	p. 301
<code>XSLPsetstrcontrol</code>	Set the value of a string problem control	p. 302
<code>XSLPsetstring</code>	Set a value in the Xpress-SLP string table	p. 303
<code>XSLPsetuniqueprefix</code>	Find a prefix character string which is different from all the names currently in use within the SLP problem	p. 304

<code>XSLPsetuserfuncaddress</code>	Change the address of a user function	p. 305
<code>XSLPsetuserfuncinfo</code>	Set up the argument information array for a user function call	p. 306
<code>XSLPsetuserfuncobject</code>	Set or define one of the objects which can be accessed by the user functions	p. 307
<code>XSLPtime</code>	Print the current date and time	p. 308
<code>XSLPtokencount</code>	Count the number of tokens in a free-format character string	p. 309
<code>XSLPtoVBString</code>	Return a string to VB given its address in Xpress-SLP	p. 310
<code>XSLPuprintmemory</code>	Print the dimensions and memory allocations for a problem	p. 311
<code>XSLPuserfuncinfo</code>	Get or set user function declaration information	p. 312
<code>XSLPvalidate</code>	Validate the feasibility of constraints in a converged solution	p. 315
<code>XSLPvalidformula</code>	Check a formula in internal (parsed or unparsed) format for unknown tokens	p. 313
<code>XSLPwriteprob</code>	Write the current problem to a file in extended MPS or text format	p. 316

XSLPaddcoefs

Purpose

Add non-linear coefficients to the SLP problem

Synopsis

```
int XPRS_CC XSLPaddcoefs(XSLPprob Prob, int nSLPCoef, int *RowIndex,
    int *ColIndex, double *Factor, int *FormulaStart, int Parsed,
    int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPCoef	Number of non-linear coefficients to be added.
RowIndex	Integer array holding index of row for the coefficient. This respects the setting of XPRS_CSTYLE.
ColIndex	Integer array holding index of column for the coefficient This respects the setting of XPRS_CSTYLE.
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 nSLPCoef+1 holding the start position in the arrays Type and Value of the formula for the coefficients. FormulaStart[nSLPCoef] should be set to the next position after the end of the last formula.
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 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 RowIndex[3], ColIndex[3], FormulaStart[4], Type[8];
int n, nSLPCoef;
double Value[8];

RowIndex[0] = 1; ColIndex[0] = 2;
RowIndex[1] = 1; ColIndex[1] = 6;
RowIndex[2] = 3; ColIndex[2] = 2;

n = nSLPCoef = 0;
FormulaStart[nSLPCoef++] = n;
Type[n] = XSLP_COL; Value[n++] = 3;
Type[n++] = XSLP_EOF;

FormulaStart[nSLPCoef++] = n;
Type[n] = XSLP_COL; Value[n++] = 2;
Type[n] = XSLP_COL; Value[n++] = 2;
Type[n] = XSLP_OP; Value[n++] = XSLP_MULTIPLY;
Type[n++] = XSLP_EOF;

FormulaStart[nSLPCoef++] = n;
Type[n] = XSLP_COL; Value[n++] = 2;
```

```

Type[n++] = XSLP_EOF;

FormulaStart[nSLPCoef] = n;

XSLPaddcoefs(Prob, nSLPCoef,RowIndex, ColIndex,
             NULL, FormulaStart, 1, 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-SLP 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 "Formula Parsing".

The XSLPadd... functions load additional items into the SLP problem. The corresponding XSLPload... functions delete any existing items first.

Related topics

[XSLPchgcoef](#), [XSLPchgccoef](#), [XSLPgetcoef](#), [XSLPgetccoef](#), [XSLPloadcoefs](#)

XSLPaddcvars

Purpose

Add character variables (CVs) to the SLP problem

Synopsis

```
int XPRS_CC XSLPaddcvars(XSLPprob Prob, int nSLPCVar, char *cValue);
```

Arguments

Prob	The current SLP problem.
nSLPCVar	Number of character variables to be added.
cValue	Character buffer holding the values of the character variables; each one must be terminated by a null character.

Example

The following example adds three character variables to the problem, which contain "The first string", "String 2" and "A third set of characters" respectively

```
char *cValue="The first string\0"
            "String 2\0"
            "A third set of characters";
XSLPaddcvars(Prob, 3, cValue);
```

Further information

The XSLPadd... functions load additional items into the SLP problem. The corresponding XSLPload... functions delete any existing items first.

Related topics

[XSLPchgcvvar](#), [XSLPgetcvvar](#), [XSLPloadcvars](#)

XSLPadddc

Purpose

Add delayed constraints (DCs) to the SLP problem

Synopsis

```
int XPRS_CC XSLPadddc(XSLPprob Prob, int nSLPDC, int *RowIndex,
    int *Delay, int *DCStart, int Parsed, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPDC	Number of DCs to be added.
RowIndex	Integer array of the row indices of the DCs. This respects the setting of XPRS_CSTYLE.
Delay	Integer array of length nSLPDC holding the delay after initiation for each DC (see below).
DCStart	Integer array of length nSLPDC holding the start position in the arrays Type and Value of the formula for each DC. The DCStart entry should be negative for any DC which does not have a formula to determine the DC initiation.
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

The following example adds rows 3 and 5 to the list of delayed constraints. Row 3 is delayed until 2 SLP iterations after column 12 becomes nonzero; row 5 is delayed for 10 SLP iterations from the start (that is, until SLP iteration 11).

```
int RowIndex[2], Delay[2], DCStart[2], Type[2];
double Value[2];

RowIndex[0] = 3; Delay[0] = 2; DCStart[0] = 0;
Type[0] = XSLP_COL; Value[0] = 12;
Type[1] = XSLP_EOF;

RowIndex[1] = 5; Delay[1] = 10; DCStart[1] = -1;

XSLPadddc(Prob, 2, RowIndex, Delay, DCStart, 1, Type, Value);
```

Note that the entry for row 5 has a negative DCStart because there is no specific initiation formula (the countdown is started when the SLP optimization starts).

Further information

The XSLPadd... functions load additional items into the SLP problem. The corresponding XSLPload... functions delete any existing items first.

The token type and value arrays Type and Value follow the rules for parsed or unparsed formulae. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an XSLP_EOF token.

If a formula is provided, then the DC will be initiated when the formula first becomes nonzero. If no formula (or an empty formula) is given, the DC is initiated immediately.

The value of Delay is used to determine when a DC becomes active. If the value is zero then the value of XSLP_DCLIMIT is used instead. A value of 1 means that the DC becomes active immediately it is initiated; a value of 2 means that the DC will become active after 1 more

iteration and so on. DCs are normally checked at the end of each SLP iteration, so it is possible that a solution will be converged but activation of additional DCs will force optimization to continue. A negative value may be given for `Delay`, in which case the absolute value is used but the DC is not checked at the end of the optimization.

Related topics

[XSLPchgdc](#), [XSLPgetdc](#), [XSLPloadcs](#)

XSLPadddfs

Purpose

Add a set of distribution factors

Synopsis

```
int XSLP_CC XSLPadddfs(XSLPprob Prob, int nDF, const int *ColIndex,
    const int *RowIndex, const double *Value)
```

Arguments

Prob	The current SLP problem.
nDF	The number of distribution factors.
ColIndex	Array of indices of columns whose distribution factor is to be changed. This respects the setting of XPRS_CSTYLE.
RowIndex	Array of indices of the rows where each distribution factor applies. This respects the setting of XPRS_CSTYLE.
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 ColIndex[3], RowIndex[3];
double Value[3];
ColIndex[0] = 282;  RowIndex[0] = 134; Value[0] = 0.1;
ColIndex[1] = 282;  RowIndex[1] = 136; Value[1] = 0.15;
ColIndex[2] = 285;  RowIndex[2] = 133; Value[2] = 1.0;
XSLPadddfs(prob, 3, ColIndex, RowIndex, 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 XSLPadd... functions load additional items into the SLP problem. The corresponding XSLPload... functions delete any existing items first.

Related topics

[XSLPchgdf](#), [XSLPgetdf](#), [XSLPloaddfs](#)

XSLPaddivfs

Purpose

Add a set of initial value formulae

Synopsis

```
int XSLP_CC XSLPaddivfs(XSLPprob Prob, int nIVF, const int *ColIndex,
    const int *IVFStart, int Parsed, const int *Type,
    const double *Value)
```

Arguments

Prob	The current SLP problem.
nIVF	The number of initial value formulae.
ColIndex	Array of indices of columns whose initial value formula is to be added. This respects the setting of XPRS_CSTYLE.
IVStart	Array of start positions in the Type and Value arrays where the formula for a the corresponding column starts. This respects the setting of XPRS_CSTYLE.
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 for each formula.
Value	Array of values corresponding to the types in Type.

Example

The following example adds initial value formulae for the following:

column 282 = column 281 * 2
column 283 = column 281 * 2
column 285 = column 282 + 101

```
int ColIndex[3], IVStart[3];
int Type[20];
double Value[20];
int n;
n = 0
ColIndex[0] = 282; IVStart[0] = n;
Type[n] = XSLP_COL; Value[n++] = 281;
Type[n] = XSLP_CON; Value[n++] = 2;
Type[n] = XSLP_OP; Value[n++] = XSLP_MULTIPLY;
Type[n] = XSLP_EOF; Value[n++] = 0;

/* Use the same formula for column 283 */
ColIndex[1] = 283; IVStart[1] = IVStart[0];

ColIndex[2] = 285; IVStart[2] = n;
Type[n] = XSLP_COL; Value[n++] = 282;
Type[n] = XSLP_CON; Value[n++] = 101;
Type[n] = XSLP_OP; Value[n++] = XSLP_PLUS;
Type[n] = XSLP_EOF; Value[n++] = 0;

XSLPaddivfs(prob, 3, ColIndex, IVStart, 1, Type, Value);
```

Further information

For more details on initial value formulae see the "IV" part of the SLPDATA section in Extended MPS format.

A formula which starts with XSLP_EOF is empty and will not create an initial value formula.

The token type and value arrays `Type` and `Value` follow the rules for parsed or unparsed formulae. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an `XSLP_EOF` token.

The `XSLPadd...` functions load additional items into the SLP problem. The corresponding `XSLPload...` functions delete any existing items first.

Related topics

[XSLPchgivf](#), [XSLPgetivf](#), [XSLPloadivfs](#)

XSLPaddnames

Purpose

Set the names of a set of SLP entities in an SLP problem.

Synopsis

```
int XPRS_CC XSLPaddnames(XSLPprob Prob, int Type, char *cNames, int First,
    int Last);
```

Arguments

Prob	The current SLP problem.
Type	Type of entity. This can be one of the Xpress-SLP manifest constants XSLP_CVNAMES, XSLP_XVNAMES, XSLP_USERFUNCNAMES.
cNames	Character array holding the names, each one terminated by a null character.
First	Index of first item whose name is to be set. All entities count from 1.
Last	Index of last item whose name is to be set.

Example

The following example sets the name of user function 1 to `MyProfit` and of user function 2 to `ProfitCalcs`

```
char *cNames = "MyProfit\0ProfitCalcs";
XSLPaddnames(Prob, XSLP_USERFUNCNAMES, cNames, 1, 2);
```

Further information

It is not necessary to set names for Xpress-SLP entities because all entities can be referred to by their index. However, if a model is being output (for example by [XSLPwriteprob](#)) then any entities without names will have internally-generated names which may not be very meaningful.

Related topics

[XSLPgetnames](#)

XSLPaddtolsets

Purpose

Add sets of standard tolerance values to an SLP problem

Synopsis

```
int XPRS_CC XSLPaddtolsets(XSLPprob Prob, int nSLPTol, double *SLPTol);
```

Arguments

Prob	The current SLP problem.
nSLPTol	The number of tolerance sets to be added.
SLPTol	Double array of (nSLPTol * 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 SLPTol[18];
for (i=0;i<9;i++) SLPTol[i] = 0.005;
SLPTol[9] = 0;
for (i=10;i<18;i=i+2) SLPTol[i] = 0.01;
for (i=11;i<18;i=i+2) SLPTol[i] = 0.001;
XSLPaddtolsets(Prob, 2, SLPTol);
```

Further information

A tolerance set is an array of 9 values containing the following tolerances:

Entry	Tolerance
0	Closure tolerance (TC)
1	Absolute delta tolerance (TA)
2	Relative delta tolerance (RA)
3	Absolute coefficient tolerance (TM)
4	Relative coefficient tolerance (RM)
5	Absolute impact tolerance (TI)
6	Relative impact tolerance (RI)
7	Absolute slack tolerance (TS)
8	Relative slack tolerance (RS)

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 zero 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 `XSLPadd...` functions load additional items into the SLP problem. The corresponding `XSLPload...` functions delete any existing items first.

Related topics

[XSLPchgtolset](#), [XSLPgettolset](#), [XSLPloadtolsets](#)

XSLPadduserfuncs

Purpose

Add user function definitions to an SLP problem.

Synopsis

```
int XPRS_CC XSLPadduserfuncs(XSLPprob Prob, int nSLPUserFunc, int *Type,
    double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPUserFunc	Number of SLP user functions to be added.
Type	Integer array of token types.
Value	Double array of token values corresponding to the types in Type.

Example

Suppose we have the following user functions written in C in a library `lib01`:

`Func1` which takes two arguments and returns two values

`Func2` which takes one argument and returns the value and (optionally) the derivative of the function. Although the function is referred to as `Func2` in the problem, we are actually using the function `NewFunc2` from the library.

The following example adds the two functions to the SLP problem:

```
int nUserFuncs, ExtName, LibName, Type[10];
double Value[10];

XSLPsetstring(Prob, &LibName, "lib01");
Type[0] = XSLP_UFARGTYPE; Value[0] = (double) 023;
Type[1] = XSLP_UFEXETYPE; Value[1] = (double) 1;
Type[2] = XSLP_STRING; Value[2] = 0;
Type[3] = XSLP_STRING; Value[3] = LibName;
Type[4] = XSLP_EOF;

XSLPsetstring(Prob, &ExtName, "NewFunc2");
Type[5] = XSLP_UFARGTYPE; Value[5] = (double) 010023;
Type[6] = XSLP_UFEXETYPE; Value[6] = (double) 1;
Type[7] = XSLP_STRING; Value[7] = ExtName;
Type[8] = XSLP_STRING; Value[8] = LibName;
Type[9] = XSLP_EOF;

XSLPgetintattrib(Prob, XSLP_UFS, &nUserFuncs);
XSLPadduserfuncs(Prob, 2, Type, Value);
XSLPaddnames(Prob, XSLP_USERFUNCNAMES, "Func1\0Func2",
    nUserFuncs+1, nUserFuncs+2);
```

Note that the values for `XSLP_UFARGTYPE` are in octal

`XSLP_UFEXETYPE` describes the functions as taking a double array of values and an integer array of function information.

The remaining tokens hold the values for the external name and the three optional parameters (*file*, *item* and *template*). `Func01` has the same internal name (in the problem) and external name (in the library), so the library name is not required. A zero string index is used as a place holder, so that the next item is correctly recognized as the library name. `Func2` has a different external name, so this appears as the first string token, followed by the library name. As neither function needs the item or template names, these have been omitted.

The number of user functions already in the problem is in the integer problem attribute `XSLP_UFS`. The new internal names are added using `XSLPaddnames`.

Further information

The token type and value arrays `Type` and `Value` are formatted in a similar way to the unparsed internal format function stack. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an `XSLP_EOF` token.

The `XSLPadd...` functions load additional items into the SLP problem. The corresponding `XSLPload...` functions delete any existing items first.

Related topics

`XSLPchguserfunc`, `XSLPgetuserfunc`, `XSLPloaduserfuncs`

XSLPaddvars

Purpose

Add SLP variables defined as matrix columns to an SLP problem

Synopsis

```
int XPRS_CC XSLPaddvars(XSLPprob Prob, int nSLPVar, int *ColIndex,
    int *VarType, int *DetRow, int *SeqNum, int *TolIndex,
    double *InitValue, double *StepBound);
```

Arguments

Prob	The current SLP problem.
nSLPVar	The number of SLP variables to be added.
ColIndex	Integer array holding the index of the matrix column corresponding to each SLP variable. This respects the setting of <code>XPRS_CSTYLE</code> .
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 <code>NULL</code> 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 <code>NULL</code> 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 <code>NULL</code> if not required.
TolIndex	Integer array holding the index of the tolerance set for each SLP variable (a zero value means the default tolerances are used) May be <code>NULL</code> if not required.
InitValue	Double array holding the initial value for each SLP variable (use the <code>VarType</code> bit map to indicate if a value is being provided) May be <code>NULL</code> 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 <code>XPRS_PLUSINFINITY</code> is used for a value in <code>StepBound</code> , the delta will never have step bounds applied, and will almost always be regarded as converged. May be <code>NULL</code> 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 ColIndex[2], VarType[2];
double InitValue[2];

ColIndex[0] = 23; VarType[0] = 0;
ColIndex[1] = 25; VarType[1] = 2; InitValue[1] = 1.42;

XSLPaddvars(Prob, 2, ColIndex, VarType, NULL, NULL,
    NULL, InitValue, NULL);
```

`InitValue` 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 `XSLPadd...` functions load additional items into the SLP problem. The corresponding `XSLPload...` functions delete any existing items first.

Related topics

[XSLPchgvar](#), [XSLPgetvar](#), [XSLPloadvars](#)

XSLPaddxvs

Purpose

Add a set of extended variable arrays (XVs) to an SLP problem

Synopsis

```
int XPRS_CC XSLPaddxvs(XSLPprob Prob, int nSLPXV, int *XVStart, int Parsed,
    int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPXV	Number of XVs to be added.
XVStart	Integer array of length nSLPXV+1 holding the start position in the arrays Type and Value of the formula or value data for the XVs. XVStart[nSLPXV] should be set to one after the end of the last XV.
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 XV item.
Value	Array of values corresponding to the types in Type.

Example

The following example adds two XVs to the current problem. The first XV contains two items: columns 3 and 6, named "Temperature" and "Pressure" respectively. The second XV has four items: column 1, the constant 1.42, the square of column 3, and column 2.

```
int n, CType, TempIndex, PressIndex, XVStart[3], Type[10];
double Value[10];

XSLPgetintcontrol(Prob, XSLP_CTYPE, CType);
n = 0;
XSLPsetstring(Prob, &TempIndex, "Temperature");
XSLPsetstring(Prob, &PressIndex, "Pressure");

XVStart[0] = n;
Type[n] = XSLP_XVVARTYPE; Value[n++] = XSLP_VAR;
Type[n] = XSLP_XVVARINDEX; Value[n++] = 3 + CType;
Type[n] = XSLP_XVINTINDEX; Value[n++] = TempIndex;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_XVVARTYPE; Value[n++] = XSLP_VAR;
Type[n] = XSLP_XVVARINDEX; Value[n++] = 6 + CType;
Type[n] = XSLP_XVINTINDEX; Value[n++] = TempIndex;
Type[n++] = XSLP_EOF;

XVStart[1] = n;
Type[n] = XSLP_XVVARTYPE; Value[n++] = XSLP_VAR;
Type[n] = XSLP_XVVARINDEX; Value[n++] = 1 + CType;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_CON; Value[n++] = 1.42;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_VAR; Value[n++] = 3 + CType;
Type[n] = XSLP_CON; Value[n++] = 2;
Type[n] = XSLP_OP; Value[n++] = XSLP_EXPONENT;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_VAR; Value[n++] = 2 + CType;
Type[n++] = XSLP_EOF;
```

```
XVStart[2] = n;  
XSLPaddxvs(Prob, 2, XVStart, 1, Type, Value);
```

When a variable is used directly as an item in an XV, it is described by two tokens: `XSLP_XVVARTYPE` and `XSLP_VARINDEX`. When used in a formula, it appears as `XSLP_VAR` or `XSLP_COL`.

Note that `XSLP_COL` cannot be used in an `XSLP_XVVARINDEX`; instead, use the setting of `XPRS_CTYPE` to convert it to a value which counts from 1, and use `XSLP_VAR`.

Because `Parsed` is set to 1, the formulae are written in internal parsed (reverse Polish) form.

Further information

The token type and value arrays `Type` and `Value` are formatted in a similar way to the unparsed internal format function stack. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an `XSLP_EOF` token.

The `XSLPadd...` functions load additional items into the SLP problem. The corresponding `XSLPload...` functions delete any existing items first.

Related topics

[XSLPchg xv](#), [XSLPget xv](#), [XSLPload xvs](#)

XSLPcalluserfunc

Purpose

Call a user function from a program or from within another user function

Synopsis

```
double XPRS_CC XSLPcalluserfunc(XSLPprob Prob, int FuncNumber, void *Arg1,
                                void *Arg2, void *Arg3, void *Arg4, void *Arg5, void *Arg6)
```

Arguments

Prob	The current SLP problem.
FuncNumber	The internal number of the function to be called.
Arg1	address of an array of double precision values holding the input values for the function. May be <code>NULL</code> if not required.
Arg2	address of an array of integer values. This must be dimensioned at least <code>XSLP_FUNCINFOSIZE</code> and is normally populated by using <code>XSLPsetuserfuncinfo</code> . This array must always be provided, even if the user function does not use it.
Arg3	address of a string buffer, normally used to hold the names of the input variables. May be <code>NULL</code> if not required.
Arg4	address of a string buffer, normally used to hold the names of the return variables. May be <code>NULL</code> if not required.
Arg5	address of an array of double precision values, normally used to hold the array of perturbations or flags for calculating first derivatives. May be <code>NULL</code> if not required.
Arg6	address of an array of double precision values, used to hold the array of return values from the function. This argument can always be provided and, if not null, will be used to hold the return value(s) from the function. May be <code>NULL</code> if not required.

Return value

If the called function returns a single value, the return value of `XSLPcalluserfunc` is the called function value; if the called function returns the address of array of values, the return value of `XSLPcalluserfunc` is the value of the first item in the array.

Example

The following example sets up the data to call user function number 2 with three input values, and prints the first return value from the function.

```
double InputArray[3], ReturnArray[4];
double FuncInfo[XSLP_FUNCINFOSIZE];

InputArray[0] = 1.42; InputArray[1] = 5;
InputArray[2] = -99;

XSLPsetuserfuncinfo(Prob, FuncInfo, 0, 3, 1, 0, 0, 0);
XSLPcalluserfunc(Prob, 2, InputArray, FuncInfo,
                 NULL, NULL, NULL, ReturnArray);
printf("Result = %lg\n", ReturnArray[0]);
```

Further information

Apart from `Arg2` (which is always required) and `Arg6` (which will always be used if it is provided), any argument required by the function must not be `NULL`. So, for example, if the function expects an array of input names then `Arg3` must be provided.

It is the user's responsibility to ensure that any arrays used are large enough to hold the data.

Related topics

[XSLPsetuserfuncinfo](#)

XSLPcascade

Purpose

Re-calculate consistent values for SLP variables. based on the current values of the remaining variables

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

Purpose

Establish a re-calculation sequence for SLP variables with determining rows.

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

XSLPchgcoef

Purpose

Add or change a single matrix coefficient using a character string for the formula

Synopsis

```
int XPRS_CC XSLPchgcoef(XSLPprob Prob, int RowIndex, int ColIndex,
    double *Factor, char *Formula);
```

Arguments

Prob	The current SLP problem.
RowIndex	The index of the matrix row for the coefficient. This respects the setting of XPRS_CSTYLE.
ColIndex	The index of the matrix column for the coefficient. This respects the setting of XPRS_CSTYLE.
Factor	Address of a double precision variable holding the constant multiplier for the formula. If <code>Factor</code> is <code>NULL</code> , 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-SLP 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, XVs, character variables or functions. Therefore, if a formula refers to a new column or XV, that new item must be added to the Xpress-SLP problem first.

Related topics

[XSLPaddcoefs](#), [XSLPchgcoef](#), [XSLPgetcoef](#), [XSLPloadcoefs](#)

XSLPchgcoef

Purpose

Add or change a single matrix coefficient using a parsed or unparsed formula

Synopsis

```
int XPRS_CC XSLPchgcoef(XSLPprob Prob, int RowIndex, int ColIndex,
    double *Factor, int Parsed, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
RowIndex	The index of the matrix row for the coefficient. This respects the setting of XPRS_CSTYLE.
ColIndex	The index of the matrix column for the coefficient. This respects the setting of XPRS_CSTYLE.
Factor	Address of a double precision variable holding the constant multiplier for the formula. If <code>Factor</code> is <code>NULL</code> , a value of 1.0 will be used.
Parsed	Integer indicating the whether the token arrays are formatted as internal unparsed (<code>Parsed=0</code>) or internal parsed reverse Polish (<code>Parsed=1</code>).
Type	Array of token types providing the description and formula for each item.
Value	Array of values corresponding to the types in <code>Type</code> .

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.

```
int n, iSin, Type[4];
double Value[4];
double Factor;

XSLPgetindex(Prob, XSLP_INTERNALFUNCNAME_SNO_CASE,
    "sin", &iSin);

n = 0;
Type[n] = XSLP_IFUN; Value[n++] = iSin;
Type[n] = XSLP_VAR; Value[n++] = 1;
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_VAR` always counts from 1, so `Col1` is always 1, regardless of the setting of `XPRS_CSTYLE`.

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-SLP 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

[XSLPaddcoefs](#), [XSLPchgcoef](#), [XSLPgetcoef](#), [XSLPloadcoefs](#)

XSLPchgcvvar

Purpose

Add or change the value of the character string corresponding to an SLP character variable

Synopsis

```
int XPRS_CC XSLPchgcvvar(XSLPprob Prob, int nSLPCVar, char *cValue);
```

Arguments

Prob	The current SLP problem.
nSLPCVar	The index of the character variable being changed. An index of zero will create a new variable.
cValue	Character buffer holding the <i>value</i> of the character variable (not its <i>name</i> , which is created by XSLPaddnames if required).

Example

Assuming that character variable 7 has already been created, the following example changes its value to "new value" and creates a new character variable called `BoxName` with the value "Jewel box"

```
XSLPchgcvvar(Prob, 7, "new value");  
  
XSLPchgcvvar(Prob, 0, "Jewel box");  
  
XSLPgetintattrib(Prob, XSLP_CVS, &n);  
XSLPaddnames(Prob, XSLP_CVNAMES, "BoxName", n, n);
```

Integer attribute `XSLP_CVS` holds the number of character variables in the problem.

Further information

Character variables can be used in formulae instead of strings, and are required in certain cases where the strings contain embedded spaces.

Related topics

[XSLPaddcvars](#), [XSLPgetcvar](#), [XSLPloadcvars](#)

XSLPchgdc

Purpose

Add or change the settings for a delayed constraint (DC)

Synopsis

```
int XPRS_CC XSLPchgdc(XSLPprob Prob, int RowIndex, char *RowType,
    int *Delay, int *IterCount, int Parsed, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
RowIndex	Index of row whose DC status is to be changed. This respects the setting of XPRS_CSTYLE.
RowType	Character buffer holding the type of the row when it is constraining. May be NULL if not required.
Delay	Address of an integer holding the delay after the DC is initiated (see below). May be NULL if not required.
IterCount	Address of an integer holding the number of SLP iterations since the DC was initiated. May be NULL if not required.
Parsed	integer indicating whether the formula is in internal unparsed (Parsed=0) or internal parsed reverse Polish (Parsed=1) format.
Type	Integer array of token types (see the section on Formula Parsing for a full list). May be NULL if not required.
Value	Array of values corresponding to the types in Type. May be NULL if not required.

Example

The following example delays row 3 until 2 SLP iterations after column 12 becomes nonzero

```
int Delay, Type[2];
double Value[2];

Delay = 2;
Type[0] = XSLP_COL; Value[0] = 12;
Type[1] = XSLP_EOF;

XSLPchgdc(Prob, 3, NULL, 2, &Delay, NULL, 0 Type, Value);
```

Further information

The formula is used to determine when the DC is initiated. If a formula is given, the DC is initiated when the formula first becomes nonzero. An empty formula and Delay = 1 means that the DC is initiated after the first SLP iteration.

If any of the addresses is NULL then the current information for the DC will be left unaltered. For a new DC, the defaults will be left unchanged.

The array of formula tokens must be terminated by an XSLP_EOF token.

If RowType is not given, the type of the row in the current matrix will be used.

If Delay is not given or is zero, the default delay from XSLP_DCLIMIT will be used. The DC is initiated when the formula (if given) first becomes nonzero. To activate a DC immediately, set Delay to 1 and provide an empty formula.

If IterCount is less than Delay, then the DC is inactive. A nonzero value for IterCount implies that the DC is initiated, and IterCount will be incremented at each subsequent SLP iteration.

If Type and/or Value is NULL the existing formula will not be changed.

If an empty formula (Type[0] = XSLP_EOF) is given, then the DC will be initiated after the delay; Delay = 1 means after the first SLP iteration.

Related topics

[XSLPaddcs](#), [XSLPgetdc](#), [XSLPladdcs](#),

XSLPchgdf

Purpose

Set or change a distribution factor

Synopsis

```
int XSLP_CC XSLPchgdf(XSLPprob Prob, int ColIndex, int RowIndex,  
    const double *Value)
```

Arguments

Prob	The current SLP problem.
ColIndex	The index of the column whose distribution factor is to be set or changed. This respects the setting of XPRS_CSTYLE.
RowIndex	The index of the row where the distribution applies. This respects the setting of XPRS_CSTYLE.
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-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](#), [XSLPgetdf](#), [XSLPloaddfs](#)

XSLPchgfuncobject

Purpose

Change the address of one of the objects which can be accessed by the user functions

Synopsis

```
int XPRS_CC XSLPchgfuncobject(int *ArgInfo, int ObjType, void **Address)
```

Arguments

ArgInfo	The array of argument information for the user function.
ObjType	An integer indicating which object is to be changed XSLP_GLOBALFUNCOBJECT The <i>Global Function Object</i> ; XSLP_USERFUNCOBJECT The <i>User Function Object</i> for the function; XSLP_INSTANCEFUNCOBJECT The <i>Instance Function Object</i> for the instance of the function.
Address	Pointer holding the address of the object.

Example

The following example from within a user function checks if there is a function instance. If so, it gets the *Instance Function Object*. If it is `NULL` an array is allocated and its address is saved as the new *Instance Function Object*.

```
int Instance;
XSLPgetfuncinfo(ArgInfo, NULL, NULL, NULL, NULL,
                 NULL, NULL, NULL, &Instance);
if (Instance) {
    XSLPgetfuncobject(ArgInfo, XSLP_INSTANCEFUNCOBJECT,
                      &Object);
    if (Object == NULL) {
        Object = calloc(4*3, sizeof(double));
        XSLPchgfuncobject(ArgInfo, XSLP_INSTANCEFUNCOBJECT,
                          &Object);
    }
}
```

Further information

This function changes the address of one of the objects which can be accessed by any user function. It requires the `ArgInfo` array of argument information. This is normally provided as one of the arguments to a user function, or it can be created by using the function

[`XSLPsetuserfuncinfo`](#)

The identity of the function and the instance are obtained from the `ArgInfo` array. Within a user function, therefore, using the `ArgInfo` array passed to the user function will change the objects accessible to that function.

If, instead, `XSLPchgfuncobject` is used with an array which has been populated by [`XSLPsetuserfuncinfo`](#), the *Global Function Object* can be set as usual. The *User Function Object* cannot be set (use [`XSLPchguserfuncobject`](#) for this purpose). There is no *Instance Function Object* as such; however, a value can be set by `XSLPchgfuncobject` which can be used by the function subsequently called by [`XSLPcalluserfunc`](#). It is the user's responsibility to manage the object and save and restore the address as necessary, because Xpress-SLP will not retain the information itself.

If `Address` is `NULL`, then the corresponding information will be unchanged.

Related topics

[`XSLPchguserfuncobject`](#), [`XSLPgetfuncobject`](#), [`XSLPgetuserfuncobject`](#),
[`XSLPsetfuncobject`](#), [`XSLPsetuserfuncobject`](#)

XSLPchgivf

Purpose

Set or change the initial value formula for a variable

Synopsis

```
int XSLP_CC XSLPchgivf(XSLPprob Prob, int ColIndex, int Parsed,
    const int *Type, const double *Value)
```

Arguments

Prob	The current SLP problem.
ColIndex	The index of the column whose initial value formula is to be set or changed. This respects the setting of XPRS_CSTYLE.
Parsed	Integer indicating the whether the token array is formatted as internal unparsed (Parsed=0) or internal parsed reverse Polish (Parsed=1).
Type	Array of token types for the formula.
Value	Array of values corresponding to the types in Type.

Example

The following example sets the initial value formula for column 282 to be column 281 * 2

```
int Type[20];
double Value[20];
int n;
n = 0
Type[n] = XSLP_COL;  Value[n++] = 281;
Type[n] = XSLP_CON;  Value[n++] = 2;
Type[n] = XSLP_OP;   Value[n++] = XSLP_MULTIPLY;
Type[n] = XSLP_EOF;  Value[n++] = 0;

XSLPchgivf(prob, 282, 1, Type, Value);
```

Further information

For more details on initial value formulae see the "IV" part of the SLPDATA section in Extended MPS format.

If the first token in Type is XSLP_EOF, any existing initial value formula will be deleted.

The token type and value arrays Type and Value follow the rules for parsed or unparsed formulae. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an XSLP_EOF token.

Related topics

[XSLPaddivfs](#), [XSLPgetivf](#), [XSLPloadivfs](#)

XSLPchgrow

Purpose

Change the status setting of a constraint

Synopsis

```
int XPRS_CC XSLPchgrow(XSLPprob Prob, int RowIndex, int *Status);
```

Arguments

Prob	The current SLP problem.
RowIndex	The index of the matrix row to be changed. This respects the setting of XPRS_CSTYLE.
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 XSLPgetrow) 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 RowIndex, Status;  
RowIndex = 9;  
XSLPgetrow(Prob, RowIndex, &Status);  
Status = Status | (1<<11);  
XSLPchgrow(Prob, RowIndex, &Status);
```

Further information

If Status is NULL the current status will remain unchanged.

Related topics

[XSLPgetrow](#)

XSLPchgrowwt

Purpose

Set or change the initial penalty error weight for a row

Synopsis

```
int XSLP_CC XSLPchgrowwt(XSLPprob Prob, int RowIndex, const double *Value)
```

Arguments

Prob	The current SLP problem.
RowIndex	The index of the row whose weight is to be set or changed. This respects the setting of <code>XPRS_CSTYLE</code> .
Value	Address of a double precision variable holding the new value of the weight. May be <code>NULL</code> 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 Value;  
Value = -3.6;  
XSLPchgrowwt(Prob, 2, &Value);  
Value = 2.0;  
XSLPchgrowwt(Prob, 4, &Value);
```

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. After that, the current weighting can be accessed and changed using [XSLProwinfo](#).

Related topics

[XSLPgetrowwt](#), [XSLProwinfo](#)

XSLPchgtolset

Purpose

Add or change a set of convergence tolerances used for SLP variables

Synopsis

```
int XPRS_CC XSLPchgtolset(XSLPprob Prob, int nSLPTol, int *Status,
    double *Tols);
```

Arguments

Prob	The current SLP problem.
nSLPTol	Tolerance set for which values are to be changed. A zero value for nSLPTol 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];

Tols[2] = 0.005;
Status = 1<<2;
XSLPchgtolset(Prob, 0, Status, Tols);
Tols[1] = Tols[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	Tolerance
0	Closure tolerance (TC)
1	Absolute delta tolerance (TA)
2	Relative delta tolerance (RA)
3	Absolute coefficient tolerance (TM)
4	Relative coefficient tolerance (RM)
5	Absolute impact tolerance (TI)
6	Relative impact tolerance (RI)
7	Absolute slack tolerance (TS)
8	Relative slack tolerance (RS)

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](#), [XSLPgettolset](#), [XSLPloadtolsets](#)

XSLPchguserfunc

Purpose

Add or change a user function in an SLP problem after the problem has been input

Synopsis

```
int XPRS_CC XSLPchguserfunc(XSLPprob Prob, int nSLPUF, char *xName,
    int *ArgType, int *ExeType, char *Param1, char *Param2,
    char *Param3);
```

Arguments

Prob	The current SLP problem.
nSLPUF	The number of the user function. This always counts from 1 and is not affected by the setting of XPRS_CSTYLE. A value of zero will create a new function.
xName	Character string containing the null-terminated external name of the user function. Note that this is not the name used in written formulae, which is created by the XSLPaddnames function if required.
ArgType	bitmap specifying existence and type of arguments: Bits 0-2 Type of DVALUE. 0=omitted, 1=NULL, 3=DOUBLE, 4=VARIANT; Bits 3-5 Type of ARGINFO. 0=omitted, 1=NULL, 2=INTEGER, 4=VARIANT; Bits 6-8 Type of ARGNAME. 0=omitted, 4=VARIANT, 6=CHAR; Bits 9-11 Type of RETNAME. 0=omitted, 4=VARIANT, 6=CHAR; Bits 12-14 Type of DELTA. 0=omitted, 1=NULL, 3=DOUBLE, 4=VARIANT; Bits 15-17 Type of RESULTS. 0=omitted, 1=NULL, 3=DOUBLE.
ExeType	type of function: Bits 0-2 determine the type of linkage: 1 = User library or DLL; 2 = Excel spreadsheet XLS; 3 = Excel macro XLF; 5 = MOSEL; 6 = VB; 7 = COM Bits 3-7 re-evaluation and derivatives flags: Bit 3-4 re-evaluation setting: 0: default; Bit 3 = 1: re-evaluation at each SLP iteration; Bit 4 = 1: re-evaluation when independent variables are outside tolerance; Bit 5 RESERVED Bit 6-7 derivatives setting: 0: default; Bit 6 = 1: tangential derivatives; Bit 7 = 1: forward derivatives Bit 8 calling mechanism: 0= standard, 1=CDECL (Windows only) Bit 9 instance setting: 0=standard, 1=function calls are grouped by instance Bit 24 multi-valued function Bit 28 non-differentiable function
Param1	null-terminated character string for first parameter (FILE).
Param2	null-terminated character string for second parameter (ITEM).
Param3	null-terminated character string for third parameter (HEADER).

Example

Suppose we have the following user functions written in C in a library `lib01`:

`Func1` which takes two arguments and returns two values

`Func2` which takes one argument and returns the value and (optionally) the derivative of the function. Although the function is referred to as `Func2` in the problem, we are actually using the function `NewFunc2` from the library.

The following example adds the two functions to the SLP problem:

```
int nUF;
```

```

XSLPgetintattrib(Prob,XSLP_UFS,&nUF);
XSLPchguserfunc(Prob, 0, NULL, 023, 1,
                 "lib01", NULL, NULL);
XSLPchguserfunc(Prob, 0, "NewFunc2", 010023, 1,
                 "lib01", NULL, NULL);

XSLPaddnames(Prob,XSLP_USERFUNCNAMES,"Func1\0Func2",
             nUF+1,nUF+2);

```

Note the use of zero as the number of the user function in order to create a new user function. A value of `NULL` for `xName` means that the internal and external function names are the same.

Further information

A `NULL` value for any of the arguments leaves the existing value (if any) unchanged. If the call is defining a new user function, a `NULL` value will leave the default value unchanged.

The following manifest constants are provided for setting evaluation and derivative bits in

`ExeType`:

Setting bit 3: `XSLP_RECALC`

Setting bit 4: `XSLP_TOLCALC`

Setting bit 6: `XSLP_2DERIVATIVE`

Setting bit 7: `XSLP_1DERIVATIVE`

Setting bit 9: `XSLP_INSTANCEFUNCTION`

Setting bit 24: `XSLP_MULTIVALUED`

Setting bit 28: `XSLP_NODERIVATIVES`

If bit 9 (`XSLP_INSTANCEFUNCTION`) is set, then calls to the function will be grouped according to the argument list, so that the function is called only once for each unique set of arguments. This happens automatically if the function is "complicated" (see the section on "User function interface" for more details).

Bit 24 (`XSLP_MULTIVALUED`) does not have to be set if the function is multi-valued and it requires `RETNAME`, `DELTA` or `RESULTS`. It must be set if the function is multi-valued, does not use any of those arrays, and may be called directly by the user application using `XSLPcalluserfunc`.

If bit 28 (`XSLP_NODERIVATIVES`) is set, then formulae involving the function will always be evaluated using numerical derivatives.

Related topics

[XSLPadduserfuncs](#), [XSLPgetuserfunc](#), [XSLPloaduserfuncs](#)

XSLPchguserfuncaddress

Purpose

Change the address of a user function

Synopsis

```
int XPRS_CC XSLPchguserfuncaddress(XSLPprob Prob, int nSLPUF,
    void **Address);
```

Arguments

Prob	The current SLP problem.
nSLPUF	The index of the user function.
Address	Pointer holding the address of the user function.

Example

The following example defines a user function via [XSLPchguserfunc](#) and then re-defines the address.

```
double InternalFunc(double *, int *);
int nUF;

XSLPchguserfunc(Prob, 0, NULL, 023, 1,
    NULL, NULL, NULL);

XSLPgetintattrib(Prob, XSLP_UFS, &nUF);
XSLPaddnames(Prob, XSLP_USERFUNCNAMES, "Func1",
    nUF, nUF);

XSLPchguserfuncaddress(Prob, nUF, &InternalFunc);
```

Note that `InternalFunc` is defined as taking two arguments (`double*` and `int*`). This matches the `ArgType` setting in [XSLPchguserfunc](#). The external function name is `NULL` because it is not required when the address is given.

Further information

`nSLPUF` is an Xpress-SLP index and always counts from 1. It is not affected by the setting of `XPRS_CSTYLE`.

If `Address` is `NULL`, then the corresponding information will be left unaltered.

The address of the function is changed to the one provided. `XSLPchguserfuncaddress` should only be used for functions declared as of type DLL or VB. Its main use is where a user function is actually internal to the system rather than being provided in an external library. In such a case, the function is initially defined as an external function using [XSLPloaduserfuncs](#), [XSLPadduserfuncs](#) or [XSLPchguserfunc](#) and the address of the function is then provided using `XSLPchguserfuncaddress`.

Related topics

[XSLPadduserfuncs](#) [XSLPchguserfunc](#), [XSLPgetuserfunc](#), [XSLPloaduserfuncs](#),
[XSLPsetuserfuncaddress](#)

XSLPchguserfuncobject

Purpose

Change or define one of the objects which can be accessed by the user functions

Synopsis

```
int XPRS_CC XSLPchguserfuncobject(XSLPprob Prob, int Entity,
    void **Address);
```

Arguments

Prob	The current SLP problem.
Entity	An integer indicating which object is to be defined. The value is interpreted as follows: 0 The <i>Global Function Object</i> ; n > 0 The <i>User Function Object</i> for user function number n; n < 0 The <i>Instance Function Object</i> for user function instance number -n.
Address	The address of a pointer to the object. If Address is NULL, then any setting of the user function object is left unaltered.

Example

The following example sets the *Global Function Object*. It then sets the *User Function Object* for the function ProfitCalcs.

```
double *GlobObj;
void *ProfitObj;
int iUF;

XSLPchguserfuncobject(Prob, 0, &GlobObj);

if (!XSLPgetindex(Prob, XSLP_USERFUNC_NAMESNOCASE,
    "ProfitCalcs", &iUF)) {
    XSLPchguserfuncobject(Prob, iUF, &ProfitObj);
}
```

The function objects can be of any type. The index of the user function is obtained using the case-insensitive search for names. If the name is not found, [XSLPgetindex](#) returns a nonzero value.

Further information

As instance numbers are not normally meaningful, this function should only be used with a negative value of n to reset all *Instance Function Objects* to NULL when a model is being re-optimized within the same program execution.

Related topics

[XSLPchgfuncobject](#), [XSLPsetfuncobject](#), [XSLPsetuserfuncobject](#)

XSLPchgvar

Purpose

Define a column as an SLP variable or change the characteristics and values of an existing SLP variable

Synopsis

```
int XPRS_CC XSLPchgvar(XSLPprob Prob, int ColIndex, int *DetRow,
    double *InitStepBound, double *StepBound, double *Penalty,
    double *Damp, double *InitValue, double *Value, int *TolSet,
    int *History, int *Converged, int *VarType);
```

Arguments

Prob	The current SLP problem.
ColIndex	The index of the matrix column. This respects the setting of XPRS_CSTYLE. May be NULL if not required.
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.
InitValue	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 InitValue;
int VarType, TolSet;

InitValue = 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 `Value`, `History` or `Converged` is only effective during SLP iterations.

Changing `InitValue` 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`, `XSLPgetvar`, `XSLPloadvars`

XSLPchg xv

Purpose

Add or change an extended variable array (XV) in an SLP problem

Synopsis

```
int XPRS_CC XSLPchg xv(XSLPprob Prob, int nSLPXV, int *nXVitems);
```

Arguments

Prob	The current SLP problem.
nSLPXV	integer holding the index of the XV. A zero index will create a new XV.
nXVitems	Address of an integer holding the number of items in the XV.

Example

The following example creates a new XV, and deletes the last item from XV number 4.

```
int nXVitem;

XSLPchg xv(Prob, 0, NULL);
XSLPget xv(Prob, 4, &nXVitem);
nXVitem--;
XSLPchg xv(Prob, 4, &nXVitem);
```

Note the use of [XSLPget xv](#) to find the current number of items in the XV.

Further information

If `nXVitems` is `NULL` then the existing value is retained. For a new XV, `nXVitems` should always be zero or `NULL`. For an existing XV, `nXVitems` can be less than or equal to the current number of items in the XV. If it is less, then items will be deleted from the end of the XV.

[XSLPchg xvitem](#) is used to add items to an existing or newly-created XV.

Related topics

[XSLPaddxvs](#), [XSLPchg xvitem](#), [XSLPget xv](#), [XSLPget xvitem](#), [XSLPloadxvs](#)

XSLPchgXvitem

Purpose

Add or change an item of an existing XV in an SLP problem

Synopsis

```
int XPRS_CC XSLPchgXvitem(XSLPprob Prob, int nSLPXV, int nXVitem,
    int Parsed, int *VarType, int *VarIndex, int *IntIndex,
    double *Reserved1, double *Reserved2, int *Reserved3, int *Type,
    double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPXV	index of the XV.
nXVitem	index of the item in the XV. If this is zero then a new item will be added to the end of the XV.
Parsed	integer indicating whether the formula of the item is in internal unparsed format (Parsed=0) or internal parsed (reverse Polish) format (Parsed=1).
VarType	Address of an integer holding the token type of the XV variable. This can be zero (there is no variable), XSLP_VAR, XSLP_CVAR or XSLP_XV.
VarIndex	Address of an integer holding the index within the VarType of the XV variable.
IntIndex	Address of an integer holding the index within the Xpress-SLP string table of the internal variable name. Zero means there is no internal name.
Reserved1	Reserved for future use.
Reserved2	Reserved for future use.
Reserved3	Reserved for future use.
Type	Integer array of token types to describe the value or formula for the XVitem.
Value	Double array of values corresponding to Type, describing the value or formula for the XVitem.

Example

The following example adds two items to XV number 4. The first is column number 25, the second is named "SQ" and is the square root of column 19.

```
int n, CType, VarType, VarIndex, IntIndex, Type[4];
double Value[4];

XSLPgetintcontrol(Prob, XPRS_CSTYLE, &CStyle);

VarType = XSLP_VAR;
VarIndex = 25 + CType;
XSLPchgXvitem(Prob, 4, 0, 1, &VarType, &VarIndex,
    NULL, NULL, NULL, NULL, NULL, NULL);

n = 0;
Type[n] = XSLP_COL;   Var[n++] = 19;
Type[n] = XSLP_CON;   Var[n++] = 0.5;
Type[n] = XSLP_OP;    Var[n++] = XSLP_EXPONENT;
Type[n++] = XSLP_EOF;

VarType = 0;
XSLPsetstring(Prob, "SQ", &IntIndex);
XSLPchgXvitem(Prob, 4, 0, 1, &VarType, NULL,
    &IntIndex, NULL, NULL, NULL,
```

`Type, Value);`

Note that columns used as XVitems are specified as `XSLP_VAR` which always counts from 1. `XSLP_COL` can be used within formulae. The formula is provided in parsed (reverse Polish) format (`Parsed=1`) which is more efficient than the unparsed form.

Further information

The XVitems for an XV will always be used in the order in which they are added.

A `NULL` value for any of the addresses will leave the existing value unchanged. If the XVitem is new, the default value will be used.

If `VarType` is zero (meaning that the XVitem is not a variable), then `VarIndex` is not used. If the variable is a column, do not use a `VarType` of `XSLP_COL` — use `XSLP_VAR` instead, and adjust the index if necessary.

The formula in `Type` and `Value` must be terminated by an `XSLP_EOF` token.

Related topics

[XSLPaddxvs](#), [XSLPgetxvitem](#), [XSLPloadxvs](#)

XSLPconstruct

Purpose

Create the full augmented SLP matrix and data structures, ready for optimization

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", "l");
```

The "l" 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

Synopsis

```
int XPRS_CC XSLPcopycallbacks(XSLPprob NewProb, XSLPprob OldProb);
```

Arguments

NewProb	The SLP problem to receive the callbacks.
OldProb	The SLP problem from which the callbacks are to be copied.

Example

The following example creates a new problem and copies only the Xpress-SLP 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-SLP callbacks and the Optimizer Library callbacks for the underlying problem. If only the Xpress-SLP 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

Synopsis

```
int XPRS_CC XSLPcopycontrols(XSLPprob NewProb, XSLPprob OldProb);
```

Arguments

NewProb	The SLP problem to receive the controls.
OldProb	The SLP problem from which the controls are to be copied.

Example

The following example creates a new problem and copies only the Xpress-SLP 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-SLP controls and the Optimizer Library controls for the underlying problem. If only the Xpress-SLP controls are required, set the integer control variable `XSLP_CONTROL` appropriately.

Related topics

`XSLP_CONTROL`

XSLPcopyprob

Purpose

Copy an existing SLP problem to another

Synopsis

```
int XPRS_CC XSLPcopyprob(XSLPprob NewProb, XSLPprob OldProb,  
    char *ProbName);
```

Arguments

NewProb	The SLP problem to receive the copy.
OldProb	The SLP problem from which to copy.
ProbName	The name to be given to the problem.

Example

The following example creates a new Xpress-SLP 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-SLP problem and the underlying Optimizer Library problem. If only the Xpress-SLP problem is required, set the integer control variable `XSLP_CONTROL` appropriately.

This function does not copy controls or callbacks. These must be copied separately using `XSLPcopycontrols` and `XSLPcopycallbacks` if required.

Related topics

`XSLP_CONTROL`

XSLPcreateprob

Purpose

Create a new SLP problem

Synopsis

```
int XPRS_CC XSLPcreateprob(XSLPprob *Prob, XPRSPprob *xProb);
```

Arguments

Prob	The address of the SLP problem variable.
xProb	The address of the underlying Optimizer Library problem variable.

Example

The following example creates an optimizer problem, and then a new Xpress-SLP problem.

```
XSLPprob nProb;  
XPRSPprob xProb;  
  
XPRScreateprob(&xProb);  
XSLPcreateprob(&nProb, &xProb);
```

Further information

An Xpress-SLP 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](#)

XSLPdecompose

Purpose

Decompose nonlinear constraints into linear and nonlinear parts

Synopsis

```
int XSLP_CC XSLPdecompose(XSLPprob Prob, int nItems, const int *Index)
```

Arguments

Prob	The current SLP problem.
nItems	The number of entries in the array Index
Index	Integer array holding the indices of the constraints to be processed. Index respects the setting of XPRS_CSTYLE. This array may be NULL, in which case all eligible constraints in the problem will be processed

Example

The following example reads a nonlinear problem and then decomposes the first ten constraints into linear and nonlinear parts:

```
int Index[10];
int Count;
XSLPreadprob(Prob, "MyProb", "");
for (Count=0; Count<10; Count++) Index[Count] = Count;
XSLPdecompose(Prob, Count, Index);
```

Further information

If nItems is zero or Index is NULL, then all constraints will be processed.

The selected constraints which contain coefficient formulae (rather than constants) are processed according to the settings of XSLP_DECOMPOSE. If a variable appears only linearly in a formula (that is, its coefficient is constant), then the variable will be deleted from the formula, and a constant coefficient will be added in the constraint for the variable. For example, the constraint $(x + 2) * y \leq 10$

would be replaced by the two terms:

$(x * y)$ (nonlinear) + $y * 2$ (linear)

If XSLP_DECOMPOSEPASSLIMIT is set to a positive number, then the SLP presolve procedure will be called after decomposition is complete. If presolve changes the problem further, then the decomposition procedure will be repeated. This sequence will be continued up to XSLP_DECOMPOSEPASSLIMIT times. To make best use of XSLPpresolve in this context, set XSLP_PRESOLVE to 15.

Related topics

XSLP_DECOMPOSE, XSLP_DECOMPOSEPASSLIMIT

XSLPdestroyprob

Purpose

Delete an SLP problem and release all the associated memory

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

Purpose

Evaluate a coefficient using the current values of the variables

Synopsis

```
int XPRS_CC XSLPevaluatecoef(XSLPprob Prob, int RowIndex, int ColIndex,  
    double *dValue);
```

Arguments

Prob	The current SLP problem.
RowIndex	Integer index of the row. This respects the setting of XPRS_CSTYLE.
ColIndex	Integer index of the column. This respects the setting of XPRS_CSTYLE.
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 `Value` setting of an SLP variable (see [XSLPchgvar](#) and [XSLPgetvar](#)).

Related topics

[XSLPchgvar](#), [XSLPevaluateformula](#) [XSLPgetvar](#)

XSLPevaluateformula

Purpose

Evaluate a formula using the current values of the variables

Synopsis

```
int XPRS_CC XSLPevaluateformula(XSLPprob Prob, int Parsed, int *Type,
    double *Value, double *dValue);
```

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.
Value	Double array of values corresponding to Type.
dValue	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 dValue, Value[10];

n = 0;
Type[n] = XSLP_COL; Value[n++] = 3;
Type[n] = XSLP_COL; Value[n++] = 6;
Type[n] = XSLP_OP; Value[n++] = XSLP_DIVIDE;
Type[n++] = XSLP_EOF;

XSLPevaluateformula(Prob, 1, Type, Value, &dValue);
```

Further information

The formula in Type and Value 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](#)

XSLPformatvalue

Purpose

Format a double-precision value in the style of Xpress-SLP

Synopsis

```
int XPRS_CC XSLPformatvalue(double dValue, char *Buffer);
```

Arguments

dValue	Double precision value to be formatted.
Buffer	Character buffer to hold the formatted result. The result will never be more than 15 characters in length including the terminating null character.

Example

The following example formats the powers of 16 from -6 to +6 and prints the results:

```
int i;
double Value;
char Buffer[16];

Value = 1;
for (i=0;i<=6;i++) {
    XSLPformatvalue(Value,Buffer);
    printf("\n16^%d = %s",i,Buffer);
    Value = Value * 16;
}
Value = 1.0/16.0;
for (i=1;i<=6;i++) {
    XSLPformatvalue(Value,Buffer);
    printf("\n16^-%d = %s",i,Buffer);
    Value = Value / 16;
}
```

The results are as follows:

```
16^0 = 1
16^1 = 16
16^2 = 256
16^3 = 4096
16^4 = 65536
16^5 = 1.048576e+006
16^6 = 1.677722e+007
16^-1 = 0.0625
16^-2 = 0.00390625
16^-3 = 0.00024414063
16^-4 = 1.525879e-005
16^-5 = 9.536743e-007
16^-6 = 5.960464e-008
```

Further information

Trailing zeroes are removed. The decimal point is removed for integers. Numbers with absolute value less than 1.0e-04 or greater than 1.0e+06 are printed in scientific format.

XSLPfree

Purpose

Free any memory allocated by Xpress-SLP and close any open Xpress-SLP files

Synopsis

```
int XPRS_CC XSLPfree(void);
```

Example

The following code frees the Xpress-SLP memory and then frees the optimizer memory:

```
XSLPfree();  
XPRSfree();
```

Further information

A call to `XSLPfree` only frees the items specific to Xpress-SLP. `XPRSfree` must be called after `XSLPfree` to free the optimizer structures.

Related topics

[XSLPinit](#)

XSLPgetbanner

Purpose

Retrieve the Xpress-SLP banner and copyright messages

Synopsis

```
int XPRS_CC XSLPgetbanner(char *Banner);
```

Argument

Banner	Character buffer to hold the banner. This will be at most 256 characters including the null terminator.
---------------	---

Example

The following example retrieves the Xpress-SLP banner and prints it

```
char Buffer[260];  
XSLPgetbanner(Buffer);  
printf("%s\n", Buffer);
```

Further information

Note that XSLPgetbanner does not take the normal Prob argument.

If XSLPgetbanner is called before XPRSinit, then it will return only the Xpress-SLP information; otherwise it will include the XPRSgetbanner information as well.

XSLPgetccoef

Purpose

Retrieve a single matrix coefficient as a formula in a character string

Synopsis

```
int XPRS_CC XSLPgetccoef(XSLPprob Prob, int RowIndex, int ColIndex,
    double *Factor, char *Formula, int fLen);
```

Arguments

Prob	The current SLP problem.
RowIndex	Integer holding the row index for the coefficient. This respects the setting of XPRS_CSTYLE.
ColIndex	Integer holding the column index for the coefficient. This respects the setting of XPRS_CSTYLE.
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.
fLen	Maximum length of returned formula.

Return value

0	Normal return.
1	Formula is too long for the buffer and has been truncated.
other	Error.

Example

The following example displays the formula for the coefficient in row 2, column 3:

```
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 1: printf("\nFormula is too long for the buffer");
        break;
default: printf("\nError accessing coefficient");
        break;
}
```

Further information

If the requested coefficient is constant, then `Factor` will be set to 1.0 and the value will be formatted in `Formula`.

If the length of the formula would exceed `fLen-1`, the formula is truncated to the last token that will fit, and the (partial) formula is terminated with a null character.

Related topics

[XSLPchgccoef](#), [XSLPchgcoef](#), [XSLPgetcoef](#)

XSLPgetcoef

Purpose

Retrieve a single matrix coefficient as a formula split into tokens

Synopsis

```
int XPRS_CC XSLPgetcoef(XSLPprob Prob, int RowIndex, int ColIndex,
    double *Factor, int Parsed, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
RowIndex	Integer holding the row index for the coefficient. This respects the setting of XPRS_CSTYLE.
ColIndex	Integer holding the column index for the coefficient. This respects the setting of XPRS_CSTYLE.
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).
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 in row 2, column 3 in unparsed form:

```
int n, Type[10];
double Value[10];

XSLPgetcoef(Prob, 2, 3, &Factor, 0, Type, Value);

for (n=0; Type[n] != XSLP_EOF; n++)
    printf("\nType=%-3d  Value=%lg", Type[n], Value[n]);
```

Further information

The `Type` and `Value` arrays are terminated by an `XSLP_EOF` token.

If the requested coefficient is constant, then `Factor` will be set to 1.0 and the value will be returned with token type `XSLP_CON`.

Related topics

[XSLPchgcoef](#), [XSLPchgcoef](#), [XSLPgetcoef](#)

XSLPgetcvar

Purpose

Retrieve the value of the character string corresponding to an SLP character variable

Synopsis

```
int XPRS_CC XSLPgetcvar(XSLPprob Prob, int nSLPCV, char *cValue);
```

Arguments

Prob	The current SLP problem.
nSLPCV	Integer holding the index of the requested character variable.
cValue	Character buffer to receive the value of the variable. The buffer must be large enough to hold the character string, which will be terminated by a null character.

Example

The following example retrieves the string stored in the character variable named `BoxType`:

```
int iCVar;  
char Buffer[200];  
XSLPgetindex(Prob, XSLP_CVNames, "BoxType", &iCVar);  
XSLPgetcvar(Prob, iCVar, Buffer);
```

Further information

Related topics

[XSLPaddcvars](#), [XSLPchgcvvar](#), [XSLPloadcvars](#)

XSLPgetdblattrib

Purpose

Retrieve the value of a double precision problem attribute

Synopsis

```
int XPRS_CC XSLPgetdblattrib(XSLPprob Prob, int Param, double *dValue);
```

Arguments

Prob	The current SLP problem.
Param	attribute (SLP or optimizer) whose value is to be returned.
dValue	Address of a double precision variable to receive the value.

Example

The following example retrieves the value of the Xpress-SLP 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`.

Related topics

`XSLPgetintattrib`, `XSLPgetstrattrib`

XSLPgetdblcontrol

Purpose

Retrieve the value of a double precision problem control

Synopsis

```
int XPRS_CC XSLPgetdblcontrol(XSLPprob Prob, int Param, double *dValue);
```

Arguments

Prob	The current SLP problem.
Param	control (SLP or optimizer) whose value is to be returned.
dValue	Address of a double precision variable to receive the value.

Example

The following example retrieves the value of the Xpress-SLP 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`.

Related topics

[XSLPgetintcontrol](#), [XSLPgetstrcontrol](#), [XSLPsetdblcontrol](#)

XSLPgetdc

Purpose

Retrieve information about a delayed constraint in an SLP problem

Synopsis

```
int XPRS_CC XSLPgetdc(XSLPprob Prob, int RowIndex, char *RowType,
    int *Delay, int *IterCount, int Parsed, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
RowIndex	The index of the matrix row . This respects the setting of XPRS_CSTYLE.
RowType	Address of character buffer to receive the type of the row when it is constraining. May be <code>NULL</code> if not required. May be <code>NULL</code> if not required.
Delay	Address of an integer to receive the delay after the DC is initiated. May be <code>NULL</code> if not required.
IterCount	Address of an integer to receive the number of SLP iterations since the DC was initiated. May be <code>NULL</code> if not required.
Parsed	Integer indicating whether the formula is to be in internal unparsed (<code>Parsed=0</code>) or parsed reverse Polish (<code>Parsed=1</code>) format.
Type	Integer array to receive the token types. May be <code>NULL</code> if not required.
Value	Array of values corresponding to the types in <code>Type</code> . May be <code>NULL</code> if not required.

Example

The following example gets the formula for the delayed constraint row 3:

```
int Type[10];
double Value[10];
XSLPgetdc(Prob, 3, NULL, NULL, 0, Type, Value);
```

The formula is returned as tokens in unparsed form.

Further information

If `RowType` is returned as zero, then the row is not currently a delayed constraint.

The formula is used to determine when the DC is initiated. An empty formula means that the DC is initiated after the first SLP iteration.

If any of the addresses is `NULL` then the corresponding information for the DC will not be provided.

The array of formula tokens will be terminated by an `XSLP_EOF` token.

Related topics

[XSLPaddcs](#), [XSLPchgdc](#), [XSLPloadcs](#)

XSLPgetdf

Purpose

Get a distribution factor

Synopsis

```
int XSLP_CC XSLPgetdf(XSLPprob Prob, int ColIndex, int RowIndex,  
    double *Value)
```

Arguments

Prob	The current SLP problem.
ColIndex	The index of the column whose distribution factor is to be retrieved. This respects the setting of XPRS_CSTYLE.
RowIndex	The index of the row from which the distribution factor is to be taken. This respects the setting of XPRS_CSTYLE.
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](#), [XSLPloaddfs](#)

XSLPgetdtime

Purpose

Retrieve a double precision time stamp in seconds

Synopsis

```
int XPRS_CC XSLPgetdtime(XSLPprob Prob, double *Seconds);
```

Arguments

Prob	The current SLP problem.
Seconds	Address of double precision variable of the time in seconds.

Example

The following example measures the elapsed time to read a problem:

```
double Start, Finish;  
XSLPgetdtime(Prob, &Start);  
XSLPreadprob(Prob, "NewMat", "");  
XSLPgetdtime(Prob, &Finish);  
printf("\nElapsed time to read = %lg secs", Finish-Start);
```

Further information

If Seconds is `NULL`, then the information will not be returned.

The timing information returned is provided by the operating system and is typically accurate to no more than 1 millisecond.

The clock is not initialized when Xpress-SLP starts, so it is necessary to save an initial time and then measure all times by difference.

Related topics

[XSLPgettime](#)

XSLPgetfuncinfo

Purpose

Retrieve the argument information for a user function

Synopsis

```
int XPRS_CC XSLPgetfuncinfo(int *ArgInfo, int *CallFlag, int *nInput,
    int *nOutput, int *nDelta, int *nInStr, int *nOutStr, int *nSLPUF,
    int *nInst)
```

Arguments

ArgInfo	The array of argument information for the user function.
CallFlag	The address of an integer to receive the caller flag value. May be <code>NULL</code> if not required.
nInput	The address of an integer to receive the number of input values. May be <code>NULL</code> if not required.
nOutput	The address of an integer to receive the number of return values. May be <code>NULL</code> if not required.
nDelta	The address of an integer to receive the number of deltas (first derivatives) required. May be <code>NULL</code> if not required.
nInStr	The address of an integer to receive the number of strings in the <code>ARGNAME</code> array. May be <code>NULL</code> if not required.
nOutStr	The address of an integer to receive the number of strings in the <code>RETNAME</code> array. May be <code>NULL</code> if not required.
nSLPUF	The address of an integer to receive the number of the function. May be <code>NULL</code> if not required.
nInst	The address of an integer to receive the instance number for the call. May be <code>NULL</code> if not required.

Example

The following example retrieves the number of the function and the problem pointer. It then retrieves the internal name by which the function is known.

```
char fName[60];
int fNum;
XSLPprob Prob;
void *Object;

XSLPgetfuncinfo(ArgInfo, NULL, NULL,
    NULL, NULL, NULL, NULL,
    &fNum, NULL);
XSLPgetfuncobject(ArgInfo, XSLP_XSLPPROBLEM, &Object);
Prob = (XSLPprob) Object;
XSLPgetnames(Prob, XSLP_USERFUNCNAMES, fName, fNum, fNum);
```

Further information

If any of the addresses is `NULL` the corresponding information will not be returned.

Related topics

[XSLPgetfuncinfoV](#), [XSLPsetuserfuncinfo](#)

XSLPgetfuncinfoV

Purpose

Retrieve the argument information for a user function

Synopsis

```
int XPRS_CC XSLPgetfuncinfoV(VARIANT *ArgInfo, int *CallFlag, int *nInput,
                             int *nOutput, int *nDelta, int *nInStr, int *nOutStr, int *nSLPUF,
                             int *nInst)
```

Arguments

ArgInfo	The array of argument information for the user function.
CallFlag	The address of an integer to receive the caller flag value. May be <code>NULL</code> if not required.
nInput	The address of an integer to receive the number of input values. May be <code>NULL</code> if not required.
nOutput	The address of an integer to receive the number of return values. May be <code>NULL</code> if not required.
nDelta	The address of an integer to receive the number of deltas (first derivatives) required. May be <code>NULL</code> if not required.
nInStr	The address of an integer to receive the number of strings in the <code>ARGNAME</code> array. May be <code>NULL</code> if not required.
nOutStr	The address of an integer to receive the number of strings in the <code>RETNAME</code> array. May be <code>NULL</code> if not required.
nSLPUF	The address of an integer to receive the number of the function. May be <code>NULL</code> if not required.
nInst	The address of an integer to receive the instance number for the call. May be <code>NULL</code> if not required.

Example

The following example retrieves the number of the function and the problem pointer. It then retrieves the internal name by which the function is known.

```
char fName[60];
int fNum;
XSLPprob Prob;
void *Object;

XSLPgetfuncinfo(ArgInfo, NULL, NULL,
                NULL, NULL, NULL, NULL,
                &fNum, NULL);
XSLPgetfuncobjectV(ArgInfo, XSLP_XSLPPROBLEM, &Object);
Prob = (XSLPprob) Object;
XSLPgetnames(Prob, XSLP_USERFUNCNames, fName, fNum, fNum);
```

Further information

This function is identical to [XSLPgetfuncinfo](#) except that `ArgInfo` is of type `VARIANT` rather than `int`. It is used in VB or COM functions when the argument information array is passed as one of the arguments. To use this version of the function, pass the first member of array as the first argument to the function — e.g.

```
XSLPgetfuncinfoV(ArgInfo(0), ...)
```

If any of the addresses is `NULL` the corresponding information will not be returned.

Related topics

[XSLPgetfuncinfo](#), [XSLPsetuserfuncinfo](#)

XSLPgetfuncobject

Purpose

Retrieve the address of one of the objects which can be accessed by the user functions

Synopsis

```
int XPRS_CC XSLPgetfuncobject(int *ArgInfo, int ObjType, void **Address)
```

Arguments

ArgInfo	The array of argument information for the user function.
ObjType	An integer indicating which object is to be returned. The following values are defined: XSLP_XSLPPROBLEM The Xpress-SLP problem pointer; XSLP_XPRSPROBLEM The Xpress Optimizer problem pointer; XSLP_GLOBALFUNCOBJECT The <i>Global Function Object</i> ; XSLP_USERFUNCOBJECT The <i>User Function Object</i> for the current function; XSLP_INSTANCEFUNCOBJECT The <i>Instance Function Object</i> for the current instance;
Address	Pointer to hold the address of the object.

Example

The following example retrieves the number of the function and the problem pointer. It then retrieves the internal name by which the function is known.

```
char fName[60];
int fNum;
XSLPprob Prob;
void *Object;

XSLPgetfuncinfo(ArgInfo, NULL, NULL,
                NULL, NULL, NULL, NULL,
                &fNum, NULL);
XSLPgetfuncobject(ArgInfo, XSLP_XSLPPROBLEM, &Object);
Prob = (XSLPprob) Object;
XSLPgetnames(Prob, XSLP_USERFUNCNames, fName, fNum, fNum);
```

Further information

For functions which have no current instance because the function does not have instances, the *Instance Function Object* will be NULL.

For functions which have no current instance because the function was called directly from another user function, the *Instance Function Object* will be that set by the calling function.

Related topics

[XSLPchgfuncobject](#), [XSLPchguserfuncobject](#), [XSLPgetfuncobjectV](#),
[XSLPgetuserfuncobject](#), [XSLPsetfuncobject](#), [XSLPsetuserfuncobject](#)

XSLPgetfuncobjectV

Purpose

Retrieve the address of one of the objects which can be accessed by the user functions

Synopsis

```
int XPRS_CC XSLPgetfuncobjectV(VARIANT *ArgInfo, int ObjType,
    void **Address)
```

Arguments

ArgInfo	The array of argument information for the user function.
ObjType	An integer indicating which object is to be returned. The following values are defined: XSLP_XSLPPROBLEM The Xpress-SLP problem pointer; XSLP_XPRSPROBLEM The Xpress Optimizer problem pointer; XSLP_GLOBALFUNCOBJECT The <i>Global Function Object</i> ; XSLP_USERFUNCOBJECT The <i>User Function Object</i> for the current function; XSLP_INSTANCEFUNCOBJECT The <i>Instance Function Object</i> for the current instance;
Address	Pointer to hold the address of the object.

Example

The following example retrieves the number of the function and the problem pointer. It then retrieves the internal name by which the function is known.

```
char fName[60];
int fNum;
XSLPprob Prob;
void *Object;

XSLPgetfuncinfoV(ArgInfo, NULL, NULL,
    NULL, NULL, NULL, NULL,
    &fNum, NULL);
XSLPgetfuncobjectV(ArgInfo, XSLP_XSLPPROBLEM, &Object);
Prob = (XSLPprob) Object;
XSLPgetnames(Prob, XSLP_USERFUNCNAMES, fName, fNum, fNum);
```

Further information

This function is identical to [XSLPgetfuncobject](#) except that `ArgInfo` is of type `VARIANT` rather than `int`. It is used in VB or COM functions when the argument information array is passed as one of the arguments. To use this version of the function, pass the first member of array as the first argument to the function — e.g.

```
XSLPgetfuncobjectV(ArgInfo(0),.....)
```

For functions which have no current instance because the function does not have instances, the *Instance Function Object* will be `NULL`.

For functions which have no current instance because the function was called directly from another user function, the *Instance Function Object* will be that set by the calling function.

Related topics

[XSLPchgfuncobject](#), [XSLPchguserfuncobject](#), [XSLPgetfuncobject](#),
[XSLPgetuserfuncobject](#), [XSLPsetfuncobject](#), [XSLPsetuserfuncobject](#)

XSLPgetindex

Purpose

Retrieve the index of an Xpress-SLP entity with a given name

Synopsis

```
int XPRS_CC XSLPgetindex(XSLPprob Prob, int Type, char *cName, int *Index);
```

Arguments

Prob	The current SLP problem.
Type	Type of entity. The following are defined: XSLP_CV NAMES (=3) Character variables; XSLP_XV NAMES (=4) Extended variable arrays; XSLP_USERFUNC NAMES (=5) User functions; XSLP_INTERNALFUNC NAMES (=6) Internal functions; XSLP_USERFUNC NAMES NOCASE (=7) User functions, case insensitive; XSLP_INTERNALFUNC NAMES NOCASE (=8) Internal functions, case insensitive; The constants 1 (for row names) and 2 (for column names) may also be used.
cName	Character string containing the name, terminated by a null character.
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_INTERNALFUNC NAMES NOCASE,  
             "SIN", &UpperIndex);  
XSLPgetindex(Prob, XSLP_INTERNALFUNC NAMES NOCASE,  
             "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` which ignore the setting of `XPRS_CSTYLE`. A value of zero returned in `Index` means there is no matching item. The case-insensitive types will find the first match regardless of the case of `cName` or of the defined function.

Related topics

[XSLPgetnames](#)

XSLPgetintattrib

Purpose

Retrieve the value of an integer problem attribute

Synopsis

```
int XPRS_CC XSLPgetintattrib(XSLPprob Prob, int Param, int *iValue);
```

Arguments

Prob	The current SLP problem.
Param	attribute (SLP or optimizer) whose value is to be returned.
iValue	Address of an integer variable to receive the value.

Example

The following example retrieves the value of the Xpress-SLP attribute `XSLP_CVS` and of the optimizer attribute `XPRS_COLS`:

```
int nCV, nCol;  
XSLPgetintattrib(Prob, XSLP_CVS, &nCV);  
XSLPgetintattrib(Prob, XPRS_COLS, &nCol);
```

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`.

Related topics

[XSLPgetdblattrib](#), [XSLPgetstrattrib](#)

XSLPgetintcontrol

Purpose

Retrieve the value of an integer problem control

Synopsis

```
int XPRS_CC XSLPgetintcontrol(XSLPprob Prob, int Param, int *iValue);
```

Arguments

Prob	The current SLP problem.
Param	control (SLP or optimizer) whose value is to be returned.
iValue	Address of an integer variable to receive the value.

Example

The following example retrieves the value of the Xpress-SLP 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`.

Related topics

`XSLPgetdblcontrol`, `XSLPgetstrcontrol`, `XSLPsetintcontrol`

XSLPgetivf

Purpose

Get the initial value formula for a variable

Synopsis

```
int XSLP_CC XSLPgetivf(XSLPprob Prob, int ColIndex, int Parsed, int *Type,
    double *Value)
```

Arguments

Prob	The current SLP problem.
ColIndex	The index of the column whose initial value formula is to be retrieved. This respects the setting of XPRS_CSTYLE.
Parsed	Integer indicating the whether the token array is formatted as internal unparsed (Parsed=0) or internal parsed reverse Polish (Parsed=1).
Type	Array to receive token types for the formula.
Value	Array to receive values corresponding to the types in Type.

Example

The following example gets the initial value formula for column 282 in unparsed form and then prints it:

```
int Type[100];
double Value[100];
char Buffer[256];
int i;

XSLPgetivf(prob, 282, 0, Type, Value);

for (i=0; Type[i]; i++) {
    XSLPitemname(prob, &Type[i], &Value[i], Buffer);
    printf("%s ", Buffer);
}
printf("\n");
```

Further information

For more details on initial value formulae see the "IV" part of the SLPDATA section in Extended MPS format.

If there is no formula for the initial value but there is a constant initial value, then a formula containing the constant value will be returned. That is:

```
XSLP_CON value
XSLP_EOF 0
```

If there is no initial value formula and no constant initial value, an empty formula will be returned. That is:

```
XSLP_EOF 0
```

The token type and value arrays `Type` and `Value` follow the rules for parsed or unparsed formulae. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an `XSLP_EOF` token.

Related topics

[XSLPaddivfs](#), [XSLPchgivf](#), [XSLPloadivfs](#)

XSLPgetlasterror

Purpose

Retrieve the error message corresponding to the last Xpress-SLP error during an SLP run

Synopsis

```
int XPRS_CC XSLPgetlasterror(XSLPprob Prob, int *Code, char *Buffer);
```

Arguments

Prob	The current SLP problem.
Code	Address of an integer to receive the message number of the last error. May be NULL if not required.
Buffer	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-SLP functions return a value of 32 to indicate a non-recoverable error. XSLPgetlasterror can retrieve the actual error number and message.

Related topics

[XSLPgetmessagetype](#)

XSLPgetmessagetype

Purpose

Retrieve the message type corresponding to a message number

Synopsis

```
int XPRS_CC XSLPgetmessagetype(int Code, int *Type);
```

Arguments

Code	Integer holding the message number.
Type	Integer to receive the message type.

Example

The following example retrieves the last error message and finds its type.

```
int Code, Type;
XSLPgetlasterror(Prob, &Code, NULL);
XSLPgetlasterror(Code, &Type);
printf("\nError %d is of type %d", Code, Type);
```

Further information

The possible values returned in `Type` are:

- 0 no such message number
- 1 information
- 3 warning
- 4 error

Related topics

[XSLPgetlasterror](#)

XSLPgetnames

Purpose

Retrieve the names of a set of Xpress-SLP entities

Synopsis

```
int XPRS_CC XSLPgetnames(XSLPprob Prob, int Type, char *cNames, int First,
    int Last);
```

Arguments

Prob	The current SLP problem.
Type	Type of entity. The following are defined: XSLP_CVNAMES (=3) Character variables XSLP_XVNAMES (=4) Extended variable arrays XSLP_USERFUNCNAMES (=5) User functions XSLP_INTERNALFUNCNAMES (=6) Internal functions For compatibility with XSLPgetindex , values for Type of 1 (rows) and 2 (columns) are also possible.
cNames	Character buffer to receive the names. Each name will be terminated by a null character.
First	Index of first item to be returned.
Last	Index of last item to be returned.

Example

The following example retrieves the names of internal function numbers 3 and 4.

```
char ch, Buffer[60];
XSLPgetnames(Prob, XSLP_INTERNALNAMES, Buffer, 3, 4);
ch = Buffer;
printf("\nFunction #3 is %s",ch);
for (;;)ch++ if (*ch == '\0') break;
ch++;
printf("\nFunction #4 is %s",ch);
```

Names are returned in `Buffer` separated by null characters. `ch` finds the null character and hence the start of the next name.

Further information

`First` and `Last` always count from 1. The setting of `XPRS_CSTYLE` has no effect in any case.

Related topics

[XSLPgetindex](#)

XSLPgetparam

Purpose

Retrieve the value of a control parameter or attribute by name

Synopsis

```
int XPRS_CC XSLPgetparam(XSLPprob Prob, const char *Param, char *cValue);
```

Arguments

Prob	The current SLP problem.
Param	Name of the control or attribute whose value is to be returned.
cValue	Character buffer to receive the value.

Example

The following example retrieves the value of the Xpress-SLP pointer attribute `XSLP_XPRSPROBLEM` which is the underlying optimizer problem pointer:

```
XSLPprob Prob;  
XPRSprob xprob;  
char Buffer[32];  
XSLPgetparam(Prob, "XSLP_XPRSPROBLEM", Buffer);  
xprob = (XPRSprob) strtol(Buffer, NULL, 16);
```

Further information

This function can be used to retrieve any Xpress-SLP or Optimizer attribute or control. The value is always returned as a character string and the receiving buffer must be large enough to hold it. It is the user's responsibility to convert the character string into an appropriate value.

Related topics

`XSLPgetdblattrib`, `XSLPgetdblcontrol`, `XSLPgetintattrib`, `XSLPgetintcontrol`
`XSLPgetstrattrib`, `XSLPgetstrcontrol`, `XSLPsetparam`

XSLPgetptrattrib

Purpose

Retrieve the value of a problem pointer attribute

Synopsis

```
int XPRS_CC XSLPgetptrattrib(XSLPprob Prob, int Param, void **Value);
```

Arguments

Prob	The current SLP problem.
Param	attribute whose value is to be returned.
Value	Address of a pointer to receive the value.

Example

The following example retrieves the value of the Xpress-SLP 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`

XSLPgetrow

Purpose

Retrieve the status setting of a constraint

Synopsis

```
int XPRS_CC XSLPgetrow(XSLPprob Prob, int RowIndex, int *Status);
```

Arguments

Prob	The current SLP problem.
RowIndex	The index of the matrix row whose data is to be obtained. This respects the setting of XPRS_CSTYLE.
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, CStyle, Status;
XSLPgetintattrib(Prob, XPRS_ROWS, &nRow);
XSLPgetintattrib(Prob, XPRS_CSTYLE, &CStyle);
CStyle = 1-CStyle;
for (iRow=0; iRow<nRow; iRow++) {
    XSLPgetrow(Prob, iRow+CStyle, &Status);
    if (Status & 0x800) printf("\nRow %d is enforced");
}
```

Further information

See the section on bitmap settings for details on the possible information in `Status`.

Related topics

[XSLPchgrow](#)

XSLPgetrowwt

Purpose

Get the initial penalty error weight for a row

Synopsis

```
int XSLP_CC XSLPgetrowwt(XSLPprob Prob, int RowIndex, double *Value)
```

Arguments

Prob	The current SLP problem.
RowIndex	The index of the row whose weight is to be retrieved. This respects the setting of XPRS_CSTYLE.
Value	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 Value;  
XSLPgetrowwt(Prob, 2, &Value)
```

Further information

The initial row weight is used only when the augmented structure is created. After that, the current weighting can be accessed and changed using [XSLProwinfo](#).

Related topics

[XSLPchgrowwt](#), [XSLProwinfo](#)

XSLPgetslpsol

Purpose

Obtain the solution values for the most recent SLP iteration

Synopsis

```
int XPRS_CC XSLPgetslpsol(XSLPprob Prob, double *x, double *slack,  
    double *dual, double *dj);
```

Arguments

Prob	The current SLP problem.
x	Double array of length XPRS_COLS to hold the values of the primal variables. May be NULL if not required.
slack	Double array of length XPRS_ROWS to hold the values of the slack variables. May be NULL if not required.
dual	Double array of length XPRS_ROWS to hold the values of the dual variables. May be NULL if not required.
dj	Double array of length XPRS_COLS 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 from the most recent SLP iteration:

```
XSLPprob prob;  
int nCol;  
double *val, *dj;  
XSLPgetintattrib(prob, XPRS_COLS, &nCol);  
val = malloc(nCol*sizeof(double));  
dj = malloc(nCol*sizeof(double));  
XSLPgetslpsol(prob, val, NULL, NULL, dj);
```

Further information

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.

XSLPgetstrattrib

Purpose

Retrieve the value of a string problem attribute

Synopsis

```
int XPRS_CC XSLPgetstrattrib(XSLPprob Prob, int Param, char *cValue);
```

Arguments

Prob	The current SLP problem.
Param	attribute (SLP or optimizer) whose value is to be returned.
cValue	Character buffer to receive the value.

Example

The following example retrieves the value of the Xpress-SLP 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`.

Related topics

`XSLPgetdblattrib`, `XSLPgetintattrib`

XSLPgetstrcontrol

Purpose

Retrieve the value of a string problem control

Synopsis

```
int XPRS_CC XSLPgetstrcontrol(XSLPprob Prob, int Param, char *cValue);
```

Arguments

Prob	The current SLP problem.
Param	control (SLP or optimizer) whose value is to be returned.
cValue	Character buffer to receive the value.

Example

The following example retrieves the value of the Xpress-SLP control `XSLP_CVNAME` and of the optimizer control `XPRS_MP SOBJNAME`:

```
char CVName[200], ObjName[200];
XSLPgetstrcontrol(Prob, XSLP_CVNAME, CVName);
XSLPgetstrcontrol(Prob, XPRS_MP SOBJNAME, ObjName);
```

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`.

Related topics

[XSLPgetdblcontrol](#), [XSLPgetintcontrol](#), [XSLPsetstrcontrol](#)

XSLPgetstring

Purpose

Retrieve the value of a string in the Xpress-SLP string table

Synopsis

```
int XPRS_CC XSLPgetstring(XSLPprob Prob, int Param, char *cValue);
```

Arguments

Prob	The current SLP problem.
Param	Index of the string whose value is to be returned.
cValue	Character buffer to receive the value.

Example

The following example retrieves string number 3

```
char Buffer[60];  
XSLPgetstring(Prob, 3, Buffer);
```

Further information

The value will be terminated by a null character. The buffer must be long enough to hold the string including the null terminator.

Strings are placed in the Xpress-SLP string table by [XSLPsetstring](#) and also by the formula parsing routines for the `XSLP_UNKNOWN` token type.

Related topics

[XSLPsetstring](#)

XSLPgettime

Purpose

Retrieve an integer time stamp in seconds and/or milliseconds

Synopsis

```
int XPRS_CC XSLPgettime(XSLPprob Prob, int *Seconds, int *MSeconds);
```

Arguments

Prob	The current SLP problem.
Seconds	Address of integer to receive the number of seconds.
MSeconds	Address of integer to receive the number of milliseconds. May be <code>NULL</code> if not required.

Example

The following example prints the time elapsed in milliseconds for reading a matrix.

```
int Secs, MSecs, StartSecs, StartMSecs, Elapsed;
XSLPgettime(Prob, &StartSecs, &StartMSecs);
XSLPreadprob(Prob, "Matrix", "");
XSLPgettime(Prob, &Secs, &MSecs);

Elapsed = (Secs-StartSecs)*1000
          + (MSecs - StartMSecs);
printf("\nElapsed time = %d", Elapsed);
```

Further information

If `Seconds` or `Milliseconds` is `NULL`, then the corresponding information will not be returned.

This routine relies on the accuracy of the system clock.

The clock is not initialized when Xpress-SLP starts, so it is necessary to save an initial time and then measure all times by difference.

Related topics

[XSLPgetdtime](#)

XSLPgettolset

Purpose

Retrieve the values of a set of convergence tolerances for an SLP problem

Synopsis

```
int XPRS_CC XSLPgettolset(XSLPprob Prob, int nSLPTol, int *Status,
    double *Tols);
```

Arguments

Prob	The current SLP problem.
nSLPTol	The index of the tolerance set.
Status	Address of integer to receive the bit-map of status settings. May be <code>NULL</code> if not required.
Tols	Array of 9 double-precision values to hold the tolerances. May be <code>NULL</code> 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 `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. `Status` indicates which of these values is active as follows. Bit `n` of `Status` is set if `Tols[n]` is active, where `n` is:

- 0 Closure tolerance (TC)
- 1 Absolute delta tolerance (TA)
- 2 Relative delta tolerance (RA)
- 3 Absolute coefficient tolerance (TM)
- 4 Relative coefficient tolerance (RM)
- 5 Absolute impact tolerance (TI)
- 6 Relative impact tolerance (RI)
- 7 Absolute slack tolerance (TS)
- 8 Relative slack tolerance (RS)

Related topics

Related topics

[XSLPaddtolsets](#), [XSLPchgtolset](#), [XSLPloadtolsets](#)

XSLPgetuserfunc

Purpose

Retrieve the type and parameters for a user function

Synopsis

```
int XPRS_CC XSLPgetuserfunc(XSLPprob Prob, int nSLPUF, char *xName,
    int *ArgType, int *ExeType, char *Param1, char *Param2,
    char *Param3);
```

Arguments

Prob	The current SLP problem.
nSLPUF	The number of the user function. This always counts from 1 and is not affected by the setting of XPRS_CSTYLE.
xName	Character string to receive the null-terminated external name of the user function. May be NULL if not required. Note that the external name is not the name used in written formulae, which is created by the XSLPaddnames function if required.
ArgType	Address of an integer to receive the bitmap specifying existence and type of arguments: Bits 0-2 Type of DVALUE. 0=omitted, 1=NULL, 3=DOUBLE, 4=VARIANT; Bits 3-5 Type of ARGINFO. 0=omitted, 1=NULL, 2=INTEGER, 4=VARIANT; Bits 6-8 Type of ARGNAME. 0=omitted, 4=VARIANT, 6=CHAR; Bits 9-11 Type of RETNAME. 0=omitted, 4=VARIANT, 6=CHAR; Bits 12-14 Type of DELTA. 0=omitted, 1=NULL, 3=DOUBLE, 4=VARIANT; Bits 15-17 Type of RESULTS. 0=omitted, 1=NULL, 3=DOUBLE. May be NULL if not required.
ExeType	Address of an integer to receive the bitmap holding the type of function: Bits 0-2 determine the type of linkage: 1 = User library or DLL; 2 = Excel spreadsheet XLS; 3 = Excel macro XLF; 5 = MOSEL; 6 = VB; 7 = COM Bits 3-7 re-evaluation and derivatives flags: Bit 3-4 re-evaluation setting: 0: default; Bit 3 = 1: re-evaluation at each SLP iteration; Bit 4 = 1: re-evaluation when independent variables are outside tolerance; Bit 5 RESERVED Bit 6-7 derivatives setting: 0: default; Bit 6 = 1: tangential derivatives; Bit 7 = 1: forward derivatives Bit 8 calling mechanism: 0= standard, 1=CDECL (Windows only) Bit 9 instance setting: 0=standard, 1=function calls are grouped by instance Bit 24 multi-valued function Bit 28 non-differentiable function May be NULL if not required.
Param1	Character buffer to hold the first parameter (FILE). May be NULL if not required.
Param2	Character buffer to hold the second parameter (ITEM). May be NULL if not required.
Param3	Character buffer to hold the third parameter (HEADER). May be NULL if not required.

Example

The following example retrieves the argument type and external name for user function number 3 and prints a simplified description of the function prototype.

```
int ArgType;
char ExtName[60];

XSLPgetuserfunc(Prob, 1, ExtName, &ArgType, NULL,
                NULL, NULL, NULL);
printf("\nFunction is %s(", ExtName);
for (i=0; i<6; i++) {
    if (i) printf(",");
    if (ArgType & (07 << i*3))
        printf("Arg%d", i+1);
}
printf(")");;
```

Related topics

[XSLPadduserfuncs](#), [XSLPchguserfunc](#), [XSLPloaduserfuncs](#)

XSLPgetuserfuncaddress

Purpose

Retrieve the address of a user function

Synopsis

```
int XPRS_CC XSLPgetuserfuncaddress(XSLPprob Prob, int nSLPUF,  
    void **Address);
```

Arguments

Prob	The current SLP problem.
nSLPUF	The number of the user function. This always counts from 1 and is not affected by the setting of XPRS_CSTYLE.
Address	Pointer to hold the address of the user function.

Example

The following example retrieves the addresses of user functions 3 and 5 and checks if they are the same.

```
void *Func3, *Func5;  
XSLPgetuserfuncaddress(Prob, 3, &Func3);  
XSLPgetuserfuncaddress(Prob, 5, &Func5);  
if (Func3 && (Func3 == Func5))  
    printf("\nFunctions are the same");
```

Further information

The address returned is the address in memory of the function for functions of type `DLL`. It will be `NULL` for functions of other types.

Related topics

[XSLPadduserfuncs](#), [XSLPchguserfunc](#), [XSLPloaduserfuncs](#)

XSLPgetuserfuncobject

Purpose

Retrieve the address of one of the objects which can be accessed by the user functions

Synopsis

```
int XPRS_CC XSLPgetuserfuncobject(XSLPprob prob, int Entity,
    void **Address);
```

Arguments

Prob	The current SLP problem.
Entity	An integer indicating which object is to be defined. The value is interpreted as follows: 0 The <i>Global Function Object</i> ; n > 0 The <i>User Function Object</i> for user function number n; n < 0 The <i>Instance Function Object</i> for user function instance number -n.
Address	Pointer to hold the address of the object.

Example

The following example retrieves the *Function Object* for user function number 3.

```
void *Obj;
XSLPgetuserfuncobject(Prob, 3, &Obj);
```

Further information

This function returns the address of one of the objects previously defined by [XSLPsetuserfuncobject](#) or [XSLPchguserfuncobject](#) . As instance numbers are not normally meaningful, this function should only be used to get the values of all *Instance Function Objects* in order, for example, to free any allocated memory.

Related topics

[XSLPgetfuncobject](#), [XSLPsetfuncobject](#), [XSLPsetuserfuncobject](#)

XSLPgetvar

Purpose

Retrieve information about an SLP variable

Synopsis

```
int XPRS_CC XSLPgetvar(XSLPprob prob, int ColIndex, int *DetRow,
    double *InitStepBound, double *StepBound, double *Penalty,
    double *Damp, double *InitValue, double *Value, int *TolSet,
    int *History, int *Converged, int *VarType, int *Delta,
    int *PenaltyDelta, int *UpdateRow, double *OldValue);
```

Arguments

Prob	The current SLP problem.
ColIndex	The index of the column. This respects the setting of XPRS_CSTYLE.
DetRow	Address of an integer to receive the index of the determining row. This respects the setting of XPRS_CSTYLE. May be NULL if not required.
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.
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.
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.
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.
InitValue	Address of a double precision variable to receive the value of the initial value of the variable. May be NULL if not required.
Value	Address of a double precision variable to receive the current activity of the variable. May be NULL if not required.
TolSet	Address of an integer to receive the index of the tolerance set of the variable. May be NULL if not required.
History	Address of an integer to receive the SLP history of the variable. May be NULL if not required.
Converged	Address of an integer to receive the convergence status of the variable (see the section on "Convergence Criteria" for more information). May be NULL if not required.
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.
Delta	Address of an integer to receive the index of the delta vector for the variable. This respects the setting of XPRS_CSTYLE. May be NULL if not required.
PenaltyDelta	Address of an integer to receive the index of the first penalty delta vector for the variable. This respects the setting of XPRS_CSTYLE. The second penalty delta immediately follows the first. May be NULL if not required.
UpdateRow	Address of an integer to receive the index of the update row for the variable. This respects the setting of XPRS_CSTYLE. May be NULL if not required.
OldValue	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 Status, History;  
double Value;  
  
XSLPgetvar(Prob, 3, NULL, NULL, NULL,  
           NULL, NULL, NULL, &Value,  
           NULL, &History, &Converged,  
           NULL, NULL, NULL, NULL, NULL);
```

Further information

If `ColIndex` refers to a column which is not an SLP variable, then all the return values will indicate that there is no corresponding data.

`DetRow` will be set to -1 if there is no determining row.

`Delta`, `PenaltyDelta` and `UpdateRow` will be set to 0-`XPRS_CSTYLE` if there is no corresponding item.

Related topics

[XSLPaddvars](#), [XSLPchgvar](#), [XSLPloadvars](#)

XSLPgetversion

Purpose

Retrieve the Xpress-SLP major and minor version numbers

Synopsis

```
int XPRS_CC XSLPgetversion(int *Major, int *Minor);
```

Arguments

Major	Address of integer to receive the major version number. May be <code>NULL</code> if not required.
Minor	Address of integer to receive the minor version number. May be <code>NULL</code> if not required.

Example

The following example retrieves the major version number of Xpress-SLP

```
int Num;  
XSLPgetversion(&Num, NULL);
```

Further information

`XSLPgetversion` can be called before `XSLPinit`.

XSLPgetxv

Purpose

Retrieve information about an extended variable array

Synopsis

```
int XPRS_CC XSLPgetxv(XSLPprob Prob, int nSLPXV, int *nXVitems);
```

Arguments

Prob	The current SLP problem.
nSLPXV	The index of the XV.
nXVitems	Address of integer to receive the number of items in the XV.

Example

The following example retrieves the number of items in extended variable array number 3.

```
int nItems;  
XSLPgetxv(Prob, 3, &nItems);
```

Further information

To obtain information on the individual items in an XV, use [XSLPgetxvitem](#).

Related topics

[XSLPgetxvitem](#)

XSLPgetxvitem

Purpose

Retrieve information about an item in an extended variable array

Synopsis

```
int XPRS_CC XSLPgetxvitem(XSLPprob Prob, int nSLPXV, int nXVitem,
    int Parsed, int *VarType, int *VarIndex, int *IntIndex,
    double *Reserved1, double *Reserved2, int *Reserved3, int *Type,
    double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPXV	index of the XV.
nXVitem	index of the item in the XV. This always counts from 1 and is independent of the setting of XPRS_CSTYLE.
Parsed	integer indicating whether the formula of the item is to be retrieved in internal unparsed format (Parsed=0) or internal parsed (reverse Polish) format (Parsed=1).
VarType	Address of an integer holding the token type of the XV variable. This can be zero (there is no variable), XSLP_VAR, XSLP_CVAR or XSLP_XV. May be NULL if not required.
VarIndex	Address of an integer holding the index within the VarType of the XV variable. May be NULL if not required.
IntIndex	Address of an integer holding the index within the Xpress-SLP string table of the internal variable name. Zero means there is no internal name. May be NULL if not required.
Reserved1	Reserved for future use.
Reserved2	Reserved for future use.
Reserved3	Reserved for future use.
Type	Integer array of token types to describe the value or formula for the XVitem. May be NULL if not required.
Value	Double array of values corresponding to Type, describing the value or formula for the XVitem. May be NULL if not required.

Example

The following example retrieves the information for the second item in XV number 3.

```
int VarType, VarIndex, IntIndex, Type[10];
double Value[10];
char Buffer[60];
XSLPgetxvitem(Prob, 3, 2, 0,
               &VarType, &VarIndex, &IntIndex,
               NULL, NULL, NULL, Type, Value);
if (VarType)
    printf("\nVariable type %d index %d", VarType, VarIndex);
if (IntIndex) {
    XSLPgetstring(Prob, IntIndex, Buffer);
    printf("\nName %s", Buffer);
}

if (!VarType)
    for (i=0; Type[i] != XSLP_EOF; i++) {
        printf("\nType=%d Value=%lg", Type[i], Value[i]);
    }
```

The formula is retrieved in unparsed format. It is assumed that there will never be more than 10 tokens in the formula, including the terminator.

Further information

If `VarType` is zero (meaning that the `XVitem` is not a variable), then `VarIndex` is not used.

The formula in `Type` and `Value` will be terminated by an `XSLP_EOF` token. `Type` and `Value` must be large enough to hold the formula.

Related topics

[XSLPaddxvs](#), [XSLPgetxvitem](#), [XSLPloadxvs](#)

XSLPglobal

Purpose

Initiate the Xpress-SLP mixed integer SLP (MISLP) algorithm

Synopsis

```
int XPRS_CC XSLPglobal(XSLPprob Prob);
```

Argument

Prob The current SLP problem.

Example

The following example optimizes the problem and then finds the integer solution.

```
XSLPmaxim(Prob, "");  
XSLPglobal(Prob);
```

Further information

The current Xpress-SLP mixed integer problem will be maximized or minimized using the algorithm defined by the control variable [XSLP_MIPALGORITHM](#).

It is recommended that [XSLPminim](#) or [XSLPmaxim](#) is used first to obtain a converged solution to the relaxed problem. If this is not done, ensure that [XSLP_OBJSENSE](#) is set appropriately.

See the chapter on "Mixed Integer Successive Linear Programming" for more information about the Xpress-SLP MISLP algorithms.

Related topics

[XSLPmaxim](#), [XSLPminim](#), [XSLP_MIPALGORITHM](#), [XSLP_OBJSENSE](#)

XSLPinit

Purpose

Initializes the Xpress-SLP system

Synopsis

```
int XPRS_CC XSLPinit();
```

Argument

none

Example

The following example initiates the Xpress-SLP 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-SLP 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](#)

XSLPitemname

Purpose

Retrieves the name of an Xpress-SLP entity or the value of a function token as a character string.

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-SLP entity. This can be any one of the token types described in the section on Formula Parsing.
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 Formula Parsing.
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];

XSLPgetcoef(Prob, 2, 3, &Factor, 0, 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-SLP entity, then an internally-generated name will be used.

Numerical values will be formatted as fixed-point or floating-point depending on their size.

Related topics

[XSLPformatvalue](#)

XSLPloadcoefs

Purpose

Load non-linear coefficients into the SLP problem

Synopsis

```
int XPRS_CC XSLPloadcoefs(XSLPprob Prob, int nSLPCoef, int *RowIndex,
    int *ColIndex, double *Factor, int *FormulaStart, int Parsed,
    int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPCoef	Number of non-linear coefficients to be loaded.
RowIndex	Integer array holding index of row for the coefficient. This respects the setting of XPRS_CSTYLE.
ColIndex	Integer array holding index of column for the coefficient This respects the setting of XPRS_CSTYLE.
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 nSLPCoef+1 holding the start position in the arrays Type and Value of the formula for the coefficients. FormulaStart[nSLPCoef] should be set to the next position after the end of the last formula.
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 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 loads coefficients representing:

Col2 * Col3 + Col6 * Col2^2 into Row1 and
Col2 ^ 2 into Row3.

```
int RowIndex[3], ColIndex[3], FormulaStart[4], Type[8];
int n, nSLPCoef;
double Value[8];

RowIndex[0] = 1; ColIndex[0] = 2;
RowIndex[1] = 1; ColIndex[1] = 6;
RowIndex[2] = 3; ColIndex[2] = 2;

n = nSLPCoef = 0;
FormulaStart[nSLPCoef++] = n;
Type[n] = XSLP_COL; Value[n++] = 3;
Type[n++] = XSLP_EOF;

FormulaStart[nSLPCoef++] = n;
Type[n] = XSLP_COL; Value[n++] = 2;
Type[n] = XSLP_COL; Value[n++] = 2;
Type[n] = XSLP_OP; Value[n++] = XSLP_MULTIPLY;
Type[n++] = XSLP_EOF;

FormulaStart[nSLPCoef++] = n;
Type[n] = XSLP_COL; Value[n++] = 2;
```

```

Type[n++] = XSLP_EOF;

FormulaStart[nSLPCoef] = n;

XSLPloadcoefs(Prob, nSLPCoef,RowIndex, ColIndex,
              NULL, FormulaStart, 1, 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-SLP 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 "Formula Parsing".

The XSLPload... functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding XSLPadd... functions add or replace items leaving other items of the same type unchanged.

Related topics

[XSLPaddcoefs](#), [XSLPchgcoef](#), [XSLPchgcccoef](#), [XSLPgetcoef](#), [XSLPgetcccoef](#)

XSLPloadcvars

Purpose

Load character variables (CVs) into the SLP problem

Synopsis

```
int XPRS_CC XSLPloadcvars(XSLPprob Prob, int nSLPCVar, char *cValue);
```

Arguments

Prob	The current SLP problem.
nSLPCVar	Number of character variables to be loaded.
cValue	Character buffer holding the values of the character variables; each one must be terminated by a null character.

Example

The following example loads three character variables into the problem, which contain "The first string", "String 2" and "A third set of characters" respectively

```
char *cValue="The first string\0"
           "String 2\0"
           "A third set of characters";
XSLPloadcvars(Prob, 3, cValue);
```

Further information

The `XSLPload...` functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding `XSLPadd...` functions add or replace items leaving other items of the same type unchanged.

Related topics

[XSLPaddcvars](#), [XSLPchgcvr](#), [XSLPgetcvar](#)

XSLPloaddc

Purpose

Load delayed constraints (DCs) into the SLP problem

Synopsis

```
int XPRS_CC XSLPloaddc(XSLPprob Prob, int nSLPDC, int *RowIndex,
    int *Delay, int *DCStart, int Parsed, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPDC	Number of DCs to be loaded.
RowIndex	Integer array of the row indices of the DCs. This respects the setting of XPRS_CSTYLE.
Delay	Integer array of length nSLPDC holding the delay after initiation for each DC (see below).
DCStart	Integer array of length nSLPDC holding the start position in the arrays Type and Value of the formula for each DC. The DCStart entry should be negative for any DC which does not have a formula to determine the DC initiation.
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

The following example loads rows 3 and 5 as the list of delayed constraints. Row 3 is delayed until 2 SLP iterations after column 12 becomes nonzero; row 5 is delayed for 10 SLP iterations from the start (that is, until SLP iteration 11).

```
int RowIndex[2], Delay[2], DCStart[2], Type[2];
double Value[2];

RowIndex[0] = 3; Delay[0] = 2; DCStart[0] = 0;
Type[0] = XSLP_COL; Value[0] = 12;
Type[1] = XSLP_EOF;

RowIndex[1] = 5; Delay[1] = 10; DCStart[1] = -1;

XSLPloaddc(Prob, 2, RowIndex, Delay, DCStart, 1, Type, Value);
```

Note that the entry for row 5 has a negative DCStart because there is no specific initiation formula (the countdown is started when the SLP optimization starts).

Further information

The XSLPload. . . functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding XSLPadd. . . functions add or replace items leaving other items of the same type unchanged.

The token type and value arrays Type and Value follow the rules for parsed or unparsed formulae. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an XSLP_EOF token.

If a formula is provided, then the DC will be initiated when the formula first becomes nonzero. If no formula (or an empty formula) is given, the DC is initiated immediately.

The value of Delay is used to determine when a DC becomes active. If the value is zero then the value of XSLP_DCLIMIT is used instead. A value of 1 means that the DC becomes active

immediately it is initiated; a value of 2 means that the DC will become active after 1 more iteration and so on. DCs are normally checked at the end of each SLP iteration, so it is possible that a solution will be converged but activation of additional DCs will force optimization to continue. A negative value may be given for `Delay`, in which case the absolute value is used but the DC is not checked at the end of the optimization.

Related topics

[XSLPaddcs](#), [XSLPchgdc](#), [XSLPgetdc](#)

XSLPloaddfs

Purpose

Load a set of distribution factors

Synopsis

```
int XSLP_CC XSLPloaddfs(XSLPprob Prob, int nDF, const int *ColIndex,
    const int *RowIndex, const double *Value)
```

Arguments

Prob	The current SLP problem.
nDF	The number of distribution factors.
ColIndex	Array of indices of columns whose distribution factor is to be changed. This respects the setting of XPRS_CSTYLE.
RowIndex	Array of indices of the rows where each distribution factor applies. This respects the setting of XPRS_CSTYLE.
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 ColIndex[3], RowIndex[3];
double Value[3];
ColIndex[0] = 282;  RowIndex[0] = 134; Value[0] = 0.1;
ColIndex[1] = 282;  RowIndex[1] = 136; Value[1] = 0.15;
ColIndex[2] = 285;  RowIndex[2] = 133; Value[2] = 1.0;
XSLPloaddfs(prob, 3, ColIndex, RowIndex, 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 `XSLPadd...` functions load additional items into the SLP problem. The corresponding `XSLPload...` functions delete any existing items first.

Related topics

[XSLPadddfs](#), [XSLPchgdf](#), [XSLPgetdf](#)

XSLPloadivfs

Purpose

Load a set of initial value formulae

Synopsis

```
int XSLP_CC XSLPloadivfs(XSLPprob Prob, int nIVF, const int *ColIndex,
    const int *IVFStart, int Parsed, const int *Type,
    const double *Value)
```

Arguments

Prob	The current SLP problem.
nIVF	The number of initial value formulae.
ColIndex	Array of indices of columns whose initial value formulae are to be loaded. This respects the setting of XPRS_CSTYLE.
IVStart	Array of start positions in the Type and Value arrays where the formula for a the corresponding column starts. This respects the setting of XPRS_CSTYLE.
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 for each formula.
Value	Array of values corresponding to the types in Type.

Example

The following example loads initial value formulae for the following:

column 282 = column 281 * 2

column 283 = column 281 * 2

column 285 = column 282 + 101

Any existing initial value formulae (for any variables) will be deleted.

```
int ColIndex[3], IVStart[3];
int Type[20];
double Value[20];
int n;
n = 0
ColIndex[0] = 282; IVStart[0] = n;
Type[n] = XSLP_COL; Value[n++] = 281;
Type[n] = XSLP_CON; Value[n++] = 2;
Type[n] = XSLP_OP; Value[n++] = XSLP_MULTIPLY;
Type[n] = XSLP_EOF; Value[n++] = 0;

/* Use the same formula for column 283 */
ColIndex[1] = 283; IVStart[1] = IVStart[0];

ColIndex[2] = 285; IVStart[2] = n;
Type[n] = XSLP_COL; Value[n++] = 282;
Type[n] = XSLP_CON; Value[n++] = 101;
Type[n] = XSLP_OP; Value[n++] = XSLP_PLUS;
Type[n] = XSLP_EOF; Value[n++] = 0;

XSLPloadivfs(prob, 3, ColIndex, IVStart, 1, Type, Value);
```

Further information

For more details on initial value formulae see the "IV" part of the SLPDATA section in Extended MPS format.

A formula which starts with XSLP_EOF is empty and will not create an initial value formula.

The token type and value arrays `Type` and `Value` follow the rules for parsed or unparsed formulae. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an `XSLP_EOF` token.

The `XSLPadd...` functions load additional items into the SLP problem. The corresponding `XSLPload...` functions delete any existing items first.

Related topics

[XSLPaddivfs](#), [XSLPchgivf](#), [XSLPgetivf](#)

XSLPloadtolsets

Purpose

Load sets of standard tolerance values into an SLP problem

Synopsis

```
int XPRS_CC XSLPloadtolsets(XSLPprob Prob, int nSLPTol, double *SLPTol);
```

Arguments

Prob	The current SLP problem.
nSLPTol	The number of tolerance sets to be loaded.
SLPTol	Double array of (nSLPTol * 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 SLPTol[18];
for (i=0;i<9;i++) SLPTol[i] = 0.005;
SLPTol[9] = 0;
for (i=10;i<18;i=i+2) SLPTol[i] = 0.01;
for (i=11;i<18;i=i+2) SLPTol[i] = 0.001;
XSLPloadtolsets(Prob, 2, SLPTol);
```

Further information

A tolerance set is an array of 9 values containing the following tolerances:

Entry	Tolerance
0	Closure tolerance (TC)
1	Absolute delta tolerance (TA)
2	Relative delta tolerance (RA)
3	Absolute coefficient tolerance (TM)
4	Relative coefficient tolerance (RM)
5	Absolute impact tolerance (TI)
6	Relative impact tolerance (RI)
7	Absolute slack tolerance (TS)
8	Relative slack tolerance (RS)

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 zero 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 `XSLPload...` functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding `XSLPadd...` functions add or replace items leaving other items of the same type unchanged.

Related topics

[XSLPaddtolsets](#), [XSLPchgtolset](#), [XSLPgettolset](#)

XSLPloaduserfuncs

Purpose

Load user function definitions into an SLP problem.

Synopsis

```
int XPRS_CC XSLPloaduserfuncs(XSLPprob Prob, int nSLPUserFunc, int *Type,
    double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPUserFunc	Number of SLP user functions to be loaded.
Type	Integer array of token types.
Value	Double array of token values corresponding to the types in Type.

Example

Suppose we have the following user functions written in C in a library `lib01`:

`Func1` which takes two arguments and returns two values

`Func2` which takes one argument and returns the value and (optionally) the derivative of the function. Although the function is referred to as `Func2` in the problem, we are actually using the function `NewFunc2` from the library.

The following example loads the two functions into the SLP problem:

```
int ExtName, LibName, Type[10];
double Value[10];

XSLPsetstring(Prob, &LibName, "lib01");
Type[0] = XSLP_UFARGTYPE; Value[0] = (double) 023;
Type[1] = XSLP_UFEXETYPE; Value[1] = (double) 1;
Type[2] = XSLP_STRING; Value[2] = 0;
Type[3] = XSLP_STRING; Value[3] = LibName;
Type[4] = XSLP_EOF;

XSLPsetstring(Prob, &ExtName, "NewFunc2");
Type[5] = XSLP_UFARGTYPE; Value[5] = (double) 010023;
Type[6] = XSLP_UFEXETYPE; Value[6] = (double) 1;
Type[7] = XSLP_STRING; Value[7] = ExtName;
Type[8] = XSLP_STRING; Value[8] = LibName;
Type[9] = XSLP_EOF;

XSLPloaduserfuncs(Prob, 2, Type, Value);
XSLPaddnames(Prob, XSLP_USERFUNCNAMES, "Func1\0Func2",
    1, 2);
```

Note that the values for `XSLP_UFARGTYPE` are in octal

`XSLP_UFEXETYPE` describes the functions as taking a double array of values and an integer array of function information.

The remaining tokens hold the values for the external name and the three optional parameters (*file*, *item* and *template*). `Func01` has the same internal name (in the problem) and external name (in the library), so the library name is not required. A zero string index is used as a place holder, so that the next item is correctly recognized as the library name. `Func2` has a different external name, so this appears as the first string token, followed by the library name. As neither function needs the item or template names, these have been omitted.

The number of user functions already in the problem is in the integer problem attribute `XSLP_UFS`. The new internal names are added using `XSLPaddnames`.

Further information

The token type and value arrays `Type` and `Value` are formatted in a similar way to the unparsed internal format function stack. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an `XSLP_EOF` token.

The `XSLPload...` functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding `XSLPadd...` functions add or replace items leaving other items of the same type unchanged.

Related topics

[XSLPadduserfuncs](#), [XSLPchguserfunc](#), [XSLPgetuserfunc](#)

XSLPloadvars

Purpose

Load SLP variables defined as matrix columns into an SLP problem

Synopsis

```
int XPRS_CC XSLPloadvars(XSLPprob Prob, int nSLPVar, int *ColIndex,
    int *VarType, int *DetRow, int *SeqNum, int *TolIndex,
    double *InitValue, double *StepBound);
```

Arguments

Prob	The current SLP problem.
nSLPVar	The number of SLP variables to be loaded.
ColIndex	Integer array holding the index of the matrix column corresponding to each SLP variable. This respects the setting of <code>XPRS_CSTYLE</code> .
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 <code>NULL</code> 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 <code>NULL</code> 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 <code>NULL</code> if not required.
TolIndex	Integer array holding the index of the tolerance set for each SLP variable (a zero value means the default tolerances are used) May be <code>NULL</code> if not required.
InitValue	Double array holding the initial value for each SLP variable (use the <code>VarType</code> bit map to indicate if a value is being provided) May be <code>NULL</code> 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 <code>XPRS_PLUSINFINITY</code> is used for a value in <code>StepBound</code> , the delta will never have step bounds applied, and will almost always be regarded as converged. May be <code>NULL</code> 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 ColIndex[2], VarType[2];
double InitValue[2];

ColIndex[0] = 23; VarType[0] = 0;
ColIndex[1] = 25; VarType[1] = 2; InitValue[1] = 1.42;

XSLPloadvars(Prob, 2, ColIndex, VarType, NULL, NULL,
    NULL, InitValue, NULL);
```

`InitValue` 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 `XSLPload...` functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding `XSLPadd...` functions add or replace items leaving other items of the same type unchanged.

Related topics

[XSLPaddvars](#), [XSLPchgvar](#), [XSLPgetvar](#)

XSLPloadxvs

Purpose

Load a set of extended variable arrays (XVs) into an SLP problem

Synopsis

```
int XPRS_CC XSLPloadxvs(XSLPprob Prob, int nSLPXV, int *XVStart,
    int Parsed, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
nSLPXV	Number of XVs to be loaded.
XVStart	Integer array of length nSLPXV+1 holding the start position in the arrays Type and Value of the formula or value data for the XVs. XVStart[nSLPXV] should be set to one after the end of the last XV.
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 XV item.
Value	Array of values corresponding to the types in Type.

Example

The following example loads two XVs into the current problem. The first XV contains two items: columns 3 and 6, named "Temperature" and "Pressure" respectively. The second XV has four items: column 1, the constant 1.42, the square of column 3, and column 2.

```
int n, CType, TempIndex, PressIndex, XVStart[3], Type[10];
double Value[10];

XSLPgetintcontrol(Prob, XSLP_CTYPE, CType);
n = 0;
XSLPsetstring(Prob, &TempIndex, "Temperature");
XSLPsetstring(Prob, &PressIndex, "Pressure");

XVStart[0] = n;
Type[n] = XSLP_XVVARTYPE; Value[n++] = XSLP_VAR;
Type[n] = XSLP_XVVARINDEX; Value[n++] = 3 + CType;
Type[n] = XSLP_XVINTINDEX; Value[n++] = TempIndex;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_XVVARTYPE; Value[n++] = XSLP_VAR;
Type[n] = XSLP_XVVARINDEX; Value[n++] = 6 + CType;
Type[n] = XSLP_XVINTINDEX; Value[n++] = TempIndex;
Type[n++] = XSLP_EOF;

XVStart[1] = n;
Type[n] = XSLP_XVVARTYPE; Value[n++] = XSLP_VAR;
Type[n] = XSLP_XVVARINDEX; Value[n++] = 1 + CType;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_CON; Value[n++] = 1.42;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_VAR; Value[n++] = 3 + CType;
Type[n] = XSLP_CON; Value[n++] = 2;
Type[n] = XSLP_OP; Value[n++] = XSLP_EXPONENT;
Type[n++] = XSLP_EOF;
Type[n] = XSLP_VAR; Value[n++] = 2 + CType;
Type[n++] = XSLP_EOF;
```

```
XVStart[2] = n;  
XSLPloadxvs(Prob, 2, XVStart, 1, Type, Value);
```

When a variable is used directly as an item in an XV, it is described by two tokens: `XSLP_XVVARTYPE` and `XSLP_VARINDEX`. When used in a formula, it appears as `XSLP_VAR` or `XSLP_COL`.

Note that `XSLP_COL` cannot be used in an `XSLP_XVVARINDEX`; instead, use the setting of `XPRS_CTYPE` to convert it to a value which counts from 1, and use `XSLP_VAR`.

Because `Parsed` is set to 1, the formulae are written in internal parsed (reverse Polish) form.

Further information

The token type and value arrays `Type` and `Value` are formatted in a similar way to the unparsed internal format function stack. For possible token types and values see the chapter on "Formula Parsing". Each formula must be terminated by an `XSLP_EOF` token.

The `XSLPload...` functions load items into the SLP problem. Any existing items of the same type are deleted first. The corresponding `XSLPadd...` functions add or replace items leaving other items of the same type unchanged.

Related topics

[XSLPaddxvs](#), [XSLPchg xv](#), [XSLPget xv](#)

XSLPmaxim

Purpose

Maximize an SLP problem

Synopsis

```
int XPRS_CC XSLPmaxim(XSLPprob Prob, char *Flags);
```

Arguments

Prob The current SLP problem.

Flags These have the same meaning as for XPRsmaxim.

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, **XSLPglocal**, **XSLPminim**, **XSLPopt**, **XSLPpresolve**

XSLPminim

Purpose

Minimize an SLP problem

Synopsis

```
int XPRS_CC XSLPminim(XSLPprob Prob, char *Flags);
```

Arguments

Prob The current SLP problem.
Flags These have the same meaning as for XPRSminim.

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, **XSLPglocal**, **XSLPmaxim**, **XSLPopt**, **XSLPpresolve**

XSLPopt

Purpose

Maximize or minimize an SLP problem

Synopsis

```
int XPRS_CC XSLPopt(XSLPprob Prob, char *Flags);
```

Arguments

Prob The current SLP problem.

Flags These have the same meaning as for XPRSmxim and XPRSminim.

Example

The following example reads an SLP problem from file and then maximizes it using the primal simplex optimizer.

```
XSLPreadprob("Matrix", "");  
XSLPsetdblcontrol(Prob, XSLP_OBJSENSE, -1);  
XSLPopt(Prob, "p");
```

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

XSLPopt is equivalent to XSLPmaxim (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, XSLPglobal, XSLPmaxim, XSLPminim, XSLPpresolve

XSLPparsecformula

Purpose

Parse a formula written as a character string into internal parsed (reverse Polish) format

Synopsis

```
int XPRS_CC XSLPparsecformula(XSLPprob Prob, char *Formula, int *nToken,
    int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
Formula	Character string containing the formula, written in the same free-format style as used in formulae in Extended MPS format, with spaces separating tokens.
nToken	Address of an integer to receive the number of tokens in the parsed formula (not counting the terminating XSLP_EOF token). May be NULL if not required.
Type	Array of token types providing the parsed formula.
Value	Array of values corresponding to the types in Type.

Example

Assuming that `x` and `y` are already defined as columns, the following example converts the formula "sin(x+y)" into internal parsed format, and then writes it out as a sequence of tokens.

```
int n, Type[20];
double Value[20];
XSLPparsecformula(Prob, "sin ( x + y )", NULL, Type, Value);
printf("\n");
for (n=0; Type[n] != XSLP_EOF; n++) {
    XSLPitemname(Prob, Type[n], Value[n], Buffer);
    printf(" %s", Buffer);
}
```

Further information

Tokens are identified by name, so any columns or user functions which appear in the formula must already have been defined. Unidentified tokens will appear as type XSLP_UNKNOWN.

Related topics

[XSLPparseformula](#), [XSLPpreparseformula](#)

XSLPparseformula

Purpose

Parse a formula written as an unparsed array of tokens into internal parsed (reverse Polish) format

Synopsis

```
int XPRS_CC XSLPparseformula(XSLPprob Prob, int *inType, double *inValue,
                             int *nToken, int *Type, double *Value);
```

Arguments

Prob	The current SLP problem.
inType	Array of token types providing the unparsed formula.
inValue	Array of values corresponding to the types in inType.
nToken	Address of an integer to receive the number of tokens in the parsed formula (not counting the terminating XSLP_EOF token). May be NULL if not required.
Type	Array of token types providing the parsed formula.
Value	Array of values corresponding to the types in Type.

Example

Assuming that `x` and `y` are already defined as columns with index `iX` and `iY` respectively, the following example converts the formula "sin(x+y)" into internal parsed format, and then writes it out as a sequence of tokens.

```
int n, iSin, iX, iY;
int inType[7], Type[20];
double inValue[7], Value[20];
n = 0;
XSLPgetIndex(Prob, XSLP_INTERNALFUNCNAME_SNO_CASE,
             "SIN", &iSin);
Type[n] = XSLP_IFUN; Value[n++] = iSin;
Type[n++] = XSLP_LB;
Type[n] = XSLP_COL; Value[n++] = iX;
Type[n] = XSLP_OP; Value[n++] = XSLP_PLUS;
Type[n] = XSLP_COL; Value[n++] = iY;
Type[n++] = XSLP_RB;
Type[n++] = XSLP_EOF;
XSLPparseformula(Prob, inType, inValue,
                 NULL, Type, Value);

printf("\n");
for (n=0; Type[n] != XSLP_EOF; n++) {
    XSLPitemname(Prob, Type[n], Value[n], Buffer);
    printf(" %s", Buffer);
}
```

Further information

For possible token types and values see the chapter on "Formula Parsing".

Related topics

[XSLPparseformula](#), [XSLPpreparseformula](#)

XSLPpreparseformula

Purpose

Perform an initial scan of a formula written as a character string, identifying the operators but not attempting to identify the types of the individual tokens

Synopsis

```
int XPRS_CC XSLPpreparseformula(XSLPprob Prob, char *Formula, int *nToken,
                                int *Type, double *Value, char *StringTable, int *SizeTable);
```

Arguments

Prob	The current SLP problem.
Formula	Character string containing the formula, written in the same free-format style as formulae in Extended MPS format, with spaces separating tokens.
nToken	Address of an integer to receive the number of tokens in the parsed formula (not counting the terminating XSLP_EOF token). May be NULL if not required.
Type	Array of token types providing the parsed formula.
Value	Array of values corresponding to the types in Type.
StringTable	Character buffer to receive the names of the unidentified tokens.
SizeTable	Address of an integer variable to hold the size of StringTable actually used. May be NULL if not required.

Example

The following example converts the formula "sin(x+y)" into internal parsed format without trying to identify the tokens apart from operands and numbers, and then writes it out as a sequence of tokens.

```
int n, Type[20];
double Value[20];
char Strings[200];
XSLPpreparseformula(Prob, "sin ( x + y )", NULL,
                    Type, Value, Strings, NULL);
printf("\n");
for (n=0; Type[n] != XSLP_EOF; n++) {
    if (Type[n] == XSLP_UNKNOWN)
        printf("\n? %s", Strings[(int)Value[n]]);
    else {
        XSLPitemname(Prob, Type[n], Value[n], Buffer);
        printf(" %s", Buffer);
    }
}
```

Further information

Only operands and numbers are identified by XSLPpreparseformula. All other operands, including names of variables, functions and XVs, are left as strings of type XSLP_UNKNOWN. The Value of such a type is the index in StringTable of the start of the token name.

The parsed formula can be converted into a calculable formula by replacing the XSLP_UNKNOWN tokens by the correct types and values.

Related topics

[XSLPparseformula](#), [XSLPparseformula](#)

XSLPpresolve

Purpose

Perform a nonlinear presolve on the problem

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 maximizes it.

```
XSLPreadprob(Prob, "Matrix", "");  
XSLPsetintcontrol(Prob, XSLP_PRESOLVE, 1);  
XSLPpresolve(Prob);  
XSLPmaximize(Prob, "");
```

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

XSLPprintmsg

Purpose

Print a message string according to the current settings for Xpress-SLP output

Synopsis

```
int XPRS_CC XSLPprintmsg(XSLPprob Prob, int MsgType, char *Msg);
```

Arguments

Prob	The current SLP problem.
MsgType	Integer containing the message type. The following types are system-defined: 1 Information message 3 Warning message 4 Error message Other message types can be used and passed to a user-supplied message handler.
Msg	Character string containing the message.

Example

The following example checks the SLP optimization status and prints an informative message for some of the possible values.

```
int Status;  
XSLPgetintattrib(Prob, XSLP_STATUS, &Status);  
if (!Status)  
    XSLPprintmsg(Prob, 1, "Fully converged solution");  
if (Status & XSLP_MAXTIME)  
    XSLPprintmsg(Prob, 3, "Max time exceeded");  
if (Status & XSLP_CONVERGEDOBJUCC)  
    XSLPprintmsg(Prob, 1, "Solution with unimportant "  
                  "unconverged values");
```

Further information

If `MsgType` is outside the range 1 to 4, any message handler written to handle the standard message types may not print the message correctly.

XSLPqparse

Purpose

Perform a quick parse on a free-format character string, identifying where each token starts

Synopsis

```
int XPRS_CC XSLPqparse(char *Record, char *Token[], int NumFields);
```

Arguments

Record	Character string to be parsed. Each token must be separated by one or more spaces from the next one.
Token	Array of character pointers to receive the start address of each token.
NumFields	Maximum number of fields to be parsed.

Return value

The number of fields processed.

Example

The following example does a quick parse of the formula "sin(x+y)" to identify where the tokens start, and then prints the first character of each token.

```
char *Token[20];
int i, n;

n = XSLPqparse("sin ( x + y )",Token,20);
for (i=0;i<n;i++)
    printf("\nToken[%d] starts with %c",i,Token[i][0]);
```

Further information

XSLPqparse does not change Record in any way. Although Token[i] will contain the address of the start of the ith token, the end of the token is still indicated by a space or the end of the record.

The return value of XSLPqparse is the number of fields processed. This will be less than NumFields if there are fewer fields in the record.

XSLPreadprob

Purpose

Read an Xpress-SLP extended MPS format matrix from a file into an SLP problem

Synopsis

```
int XPRS_CC XSLPreadprob(XSLPprob Prob, char *Probname, char *Flags);
```

Arguments

Prob	The current SLP problem.
Probname	Character string containing the name of the file from which the matrix is to be read.
Flags	Character string containing any flags needed for the input routine. No flag settings are currently recognized.

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.

For details of the format of the file, see the section on "Extended MPS Format".

Related topics

Extended MPS Format (Chapter 2), [XSLPwriteprob](#)

XSLPmaxim

Purpose

Continue the maximization of an SLP problem

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-SLP 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](#), [XSLPpreminim](#)

XSLPminim

Purpose

Continue the minimization of an SLP problem

Synopsis

```
int XPRS_CC XSLPminim(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. If it has not converged, it saves the file and continues for another 10.

```
int Status;

XSLPsetintcontrol(Prob, XSLP_ITERLIMIT, 10);
XSLPminim(Prob, "");
XSLPgetintattrib(Prob, XSLP_STATUS, &Status);
if (Status & XSLP_MAXSLPITERATIONS) {
    XSLPsave(Prob);
    XSLPsetintcontrol(Prob, XSLP_ITERLIMIT, 20);
    XSLPminim(Prob, "");
}
```

Further information

This allows Xpress-SLP 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-SLP problem from a file created by [XSLPsave](#)

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-SLP problem and the underlying optimizer problem. If only the Xpress-SLP problem is required, set the integer control variable [XSLP_CONTROL](#) appropriately.

The problem is saved into two files *save.svf* which is the optimizer save file, and *save.svx* which is the SLP save file. Both files are required for a full restore; only the *svx* file is required when the underlying optimizer problem is not being restored.

Related topics

[XSLP_CONTROL](#), [XSLPsave](#)

XSLPrevis

Purpose

Revise the unaugmented SLP matrix with data from a file

Synopsis

```
int XPRS_CC XSLPrevis(XSLPprob Prob, char *Filename);
```

Arguments

Prob	The current SLP problem.
Filename	Character string containing the name of the file with the revise data.

Example

The following example reads a matrix from file and then revises it according to the data in file "ReviseData.dat".

```
XSLPreadprob(Prob, "Matrix", "");  
XSLPrevis(Prob, "ReviseData.dat");
```

Further information

`XSLPrevis` does not implement a full revise facility. In particular, there is no provision for adding or deleting rows or columns. However, coefficients can be deleted with an explicit zero entry.

The data in the revise file is written in Extended MPS format and can change ROWS, COLUMNS, RHS, BOUNDS and RANGES data. The MODIFY, BEFORE and AFTER keywords are recognized but ignored.

`XSLPrevis` must be called before the matrix is augmented by `XSLPconstruct`.

XSLProwinfo

Purpose

Get or set row information

Synopsis

```
int XSLP_CC XSLProwinfo(XSLPprob Prob, int RowIndex, int InfoType,
    void *Info);
```

Arguments

Prob	The current SLP problem.
RowIndex	Index of the row whose information is to be handled. This respects the setting of XPRS_CSTYLE
InfoType	Type of information (see below)
Info	Address of information to be set or retrieved

Example

The following example retrieves the number of times that the penalty error vector has been active, and the total of the error activities, for row number 4:

```
int NumError;
double TotalError;
XSLProwinfo(Prob, 4, XSLP_GETROWNUMPENALTYERRORS, &NumError);
XSLProwinfo(Prob, 4, XSLP_GETROWTOTALPENALTYERROR, &TotalError);
```

Further information

The following constants are provided for row information handling:

XSLP_GETROWNUMPENALTYERRORS	Get the number of times the penalty error vector has been active
XSLP_GETROWMAXPENALTYERROR	Get the maximum size of the penalty error vector activity
XSLP_GETROWTOTALPENALTYERROR	Get the total of the penalty error vector activities
XSLP_GETROWAVERAGEPENALTYERROR	Get the average size of the penalty error vector activity
XSLP_GETROWCURRENTPENALTYERROR	Get the size of the penalty error vector activity in the current iteration. The value is negative for constraints of type L and for equalities where the left hand side is greater than the right hand side.
XSLP_GETROWCURRENTPENALTYFACTOR	Get the size of the penalty error factor for the current iteration
XSLP_SETROWPENALTYFACTOR	Set the size of the penalty error factor for the next iteration
XSLP_GETROWPENALTYCOLUMN1	Get the index of the penalty column for the row (the error column with a positive entry for an equality row)
XSLP_GETROWPENALTYCOLUMN2	Get the index of the second penalty column for an equality row (the error column with a negative entry)

Related topics

[XSLP_PENALTYINFOSTART](#)

XSLPsave

Purpose

Save the Xpress-SLP problem to file

Synopsis

```
int XPRS_CC XSLPsave(XSLPprob Prob);
```

Argument

Prob The current SLP problem.

Example

The following example saves the current problem to files named *prob1.svf* and *prob1.svx*.

```
XPRSProb xprob;  
XSLPgetptrattrib(Prob, XSLP_XPRSPROBLEM, &xprob);  
XPRSsetprobname(xprob, "prob1");  
XSLPsave(Prob);
```

Further information

The problem is saved into two files *prob.svf* which is the optimizer save file, and *prob.svx* which is the SLP save file, where *prob* is the name of the problem. Both files are used in a full save; only the *svx* file is required when the underlying optimizer problem is not being saved.

Normally `XSLPsave` saves both the Xpress-SLP problem and the underlying optimizer problem. If only the Xpress-SLP problem is required, set the integer control variable `XSLP_CONTROL` appropriately.

Related topics

`XSLP_CONTROL`, `XSLPrestore` `XSLPsaveas`

XSLPsaveas

Purpose

Save the Xpress-SLP problem to a named file

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.

Example

The following example saves the current problem to files named *MyProb.svf* and *MyProb.svx*.

```
XSLPsaveas(Prob, "MyProb");
```

Further information

The problem is saved into two files *filename.svf* which is the optimizer save file, and *filename.svx* which is the SLP save file, where *filename* is the second argument to the function. Both files are used in a full save; only the svx file is required when the underlying optimizer problem is not being saved.

Normally `XSLPsaveas` saves both the Xpress-SLP problem and the underlying optimizer problem. If only the Xpress-SLP 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

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

Purpose

Set a user callback to be called at the end of the cascading process, after the last variable has been cascaded

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called at the end of the cascading process. <code>UserFunc</code> returns an integer value. The return value is noted by Xpress-SLP but it has no effect on the optimization.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcbcascadeend</code> .
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</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 `Object` 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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPcascade](#), [XSLPsetcbcascadestart](#), [XSLPsetcbcascadevar](#), [XSLPsetcbcascadevarF](#)

XSLPsetcbcascadestart

Purpose

Set a user callback to be called at the start of the cascading process, before any variables have been cascaded

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called at the start of the cascading process. <code>UserFunc</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.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcbcascadestart</code> .
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</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 `Object` 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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPcascade](#), [XSLPsetcbcascadeend](#), [XSLPsetcbcascadevar](#), [XSLPsetcbcascadevarF](#)

XSLPsetcbcascadevar

Purpose

Set a user callback to be called after each column has been cascaded

Synopsis

```
int XPRS_CC XSLPsetcbcascadevar(XSLPprob Prob, int (XPRS_CC *UserFunc)
    (XSLPprob myProb, void *myObject, int ColIndex), void *Object);
```

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called after each column has been cascaded. <i>UserFunc</i> 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.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <i>Object</i> to <i>XSLPsetcbcascadevar</i> .
ColIndex	The number of the column which has been cascaded. <i>ColIndex</i> respects the setting of <i>XPRS_CSTYLE</i> .
Object	Address of a user-defined object, which can be used for any purpose by the function. <i>Object</i> is passed to <i>UserFunc</i> as <i>myObject</i> .

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 *Object* 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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long,
    ByVal varindex As Long) As Long
```

Related topics

[XSLPcascade](#), [XSLPsetcbcascadeend](#), [XSLPsetcbcascadestart](#), [XSLPsetcbcascadevarF](#)

XSLPsetcbcascadevarF

Purpose

Set a user callback to be called after each column has been cascaded

Synopsis

```
int XPRS_CC XSLPsetcbcascadevarF(XSLPprob Prob, int (XPRS_CC *UserFunc)
    (XSLPprob myProb, void *myObject, int *ColIndex), void *Object);
```

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called after each column has been cascaded. <i>UserFunc</i> 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.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <i>Object</i> to <i>XSLPsetcbcascadevarF</i> .
ColIndex	Address of an integer containing the number of the column which has been cascaded. The column number respects the setting of <i>XPRS_CSTYLE</i> .
Object	Address of a user-defined object, which can be used for any purpose by the function. <i>Object</i> is passed to <i>UserFunc</i> as <i>myObject</i> .

Example

The following example sets up a callback to be executed after each variable has been cascaded:

```
double *cSol;
XSLPsetcbcascadevarF(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 *pCol) {
    int iCol;
    double *cSol, Value;
    cSol = * (double **) Obj;
    iCol = *pCol;
    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 *Object* 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.

XSLPsetcbcascadevarF is identical to *XSLPsetcbcascadevar* except that the column number is passed by reference rather than by value.

Related topics

[XSLPcascade](#), [XSLPsetcbcascadeend](#), [XSLPsetcbcascadestart](#), [XSLPsetcbcascadevar](#)

XSLPsetcbconstruct

Purpose

Set a user callback to be called during the Xpress-SLP augmentation process

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called during problem augmentation. UserFunc returns an integer value. See below for an explanation of the values.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbconstruct.
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

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 Object 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, n, NULL, NULL, NULL, NULL,
                       NULL, NULL, &cValue[n], NULL, NULL, NULL,
                       NULL);
        /* set Object 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

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called when the SLP problem is about to be destroyed. UserFunc returns an integer value. At present the return value is ignored.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbdestroy.
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

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 Object 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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPdestroyprob](#)

XSLPsetcbformula

Purpose

Set a callback to be used in formula evaluation when an unknown token is found

Synopsis

```
int XPRS_CC XSLPsetcbformula(XSLPprob Prob, int (XPRS_CC *UserFunc)
    (XSLPprob myProb, void *myObject, double Value, double *Result),
    void *Object);
```

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called during formula evaluation. UserFunc returns an integer value. At present the value is ignored.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbformula.
Value	The Value of the unknown token.
Result	Address of a double precision value to hold the result of the calculation.
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

Example

The following example sets a callback to process unknown tokens in formulae. It then creates a formula with an unknown token, and evaluates it.

```
int XPRS_CC MyCB(XSLPprob MyProb, void *MyObject, double MyValue, double *Result)
{
    union { char *p; double d; } z;
    z.d = MyValue;
    if (z.p != NULL) *Result = atof(z.p);
    else *Result = 0;
    return(0);
}
...
int Type[10];
double Value[10];
int nToken;
double Answer;
union { char *p; double d; } z;

XSLPsetcbformula(prob,MyCB,NULL);

nToken = 0;
Type[nToken] = XSLP_CON; Value[nToken++] = 10;
Type[nToken] = XSLP_UNKNOWN; z.p = "25.2"; Value[nToken++] = z.d;
Type[nToken] = XSLP_OP; Value[nToken++] = XSLP_PLUS;
Type[nToken] = XSLP_EOF; Value[nToken++] = 0;

XSLPevaluateformula(prob,1,Type,Value,&Answer);

printf("Answer = %lg",Answer);
```

This demonstrates how the Value of an unknown token can be set in any way, as long as the routine that sets the token up and the callback agree on how it is to be interpreted.

In this case, the value actually contains the address of a character string, which is converted by the callback into a real number.

Related topics

[XSLPevaluateformula](#)

XSLPsetcbintsol

Purpose

Set a user callback to be called during MISLP when an integer solution is obtained

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called when an integer solution is obtained. <code>UserFunc</code> returns an integer value. At present, the return value is ignored.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcbintsol</code> .
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</code> .

Example

The following example sets up a callback to be executed whenever an integer solution is found during MISLP:

```
double *cSol;
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;
    double *cSol;
    cSol = * (double **) Obj;
    XSLPgetptrattrib(MyProb, XSLP_XPRSPROBLEM, &xprob);
    XPRSgetsol(xprob, cSol, NULL, NULL, NULL);
    return 0;
}
```

The `Object` 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 must be used during MISLP instead of the `XPRSsetcbintsol` callback which is used for MIP problems.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPsetcboptnode](#), [XSLPsetcbprenode](#)

XSLPsetcbiterend

Purpose

Set a user callback to be called at the end of each SLP iteration

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called at the end of each SLP iteration. <code>UserFunc</code> returns an integer value. If the return value is nonzero, the SLP iterations will stop.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcbiterend</code> .
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</code> .

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;
    XPRSPprob xprob;
    XSLPgetptrattrib(MyProb, XSLP_XPRSPROBLEM, &xprob);
    XSLPgetintattrib(xprob, XPRS_SIMPLEXITER, &nIter);
    if (nIter < 10) return 1;
    return 0;
}
```

The `Object` 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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPsetcbiterstart](#), [XSLPsetcbitervar](#), [XSLPsetcbitervarF](#)

XSLPsetcbiterstart

Purpose

Set a user callback to be called at the start of each SLP iteration

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called at the start of each SLP iteration. UserFunc returns an integer value. If the return value is nonzero, the SLP iterations will stop.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbiterstart.
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

Example

The following example sets up a callback to be executed at the start of the optimization to save 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);
    XPRSgetsol(xprob, cSol, NULL, NULL, NULL);
    return 0;
}
```

The Object 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. When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPsetcbiterend](#), [XSLPsetcbitervar](#), [XSLPsetcbitervarF](#)

XSLPsetcbitervar

Purpose

Set a user callback to be called after each column has been tested for convergence

Synopsis

```
int XPRS_CC XSLPsetcbitervar(XSLPprob Prob, int (XPRS_CC *UserFunc)
                             (XSLPprob myProb, void *myObject, int ColIndex), void *Object);
```

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called after each column has been tested for convergence. UserFunc returns an integer value. The return value is interpreted as a convergence status. The possible values are: < 0 The variable has not converged; 0 The convergence status of the variable is unchanged; 1 to 9 The column has converged on a system-defined convergence criterion (these values should not normally be returned); > 9 The variable has converged on user criteria.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbitervar.
ColIndex	The number of the column which has been tested for convergence. The column number respects the setting of XPRS_CSTYLE.
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

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 `Object` 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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long,  
                        ByVal varindex As Long) As Long
```

Related topics

[XSLPsetcbiterend](#), [XSLPsetcbiterstart](#), [XSLPsetcbitervarF](#)

XSLPsetcbitervarF

Purpose

Set a user callback to be called after each column has been tested for convergence

Synopsis

```
int XPRS_CC XSLPsetcbitervarF(XSLPprob Prob, int (XPRS_CC *UserFunc)
                              (XSLPprob myProb, void *myObject, int *ColIndex), void *Object);
```

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called after each column has been tested for convergence. UserFunc returns an integer value. The return value is interpreted as a convergence status. The possible values are: < 0 The variable has not converged; 0 The convergence status of the variable is unchanged; 1 to 9 The column has converged on a system-defined convergence criterion (these values should not normally be returned); > 9 The variable has converged on user criteria.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbitervarF.
ColIndex	Address of an integer holding the number of the column which has been tested for convergence. The column number respects the setting of XPRS_CSTYLE.
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

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;
XSLPsetcbitervarF(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 *pCol) {
    int *Important, Converged, iCol;
    Important = *(int **) Obj;
    iCol = *pCol;
    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 `Object` 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.

`XSLPsetcbitervarF` is identical to `XSLPsetcbitervar` except that the column number is passed by reference rather than by value.

Related topics

`XSLPsetcbiterend`, `XSLPsetcbiterstart`, `XSLPsetcbitervarF`

XSLPsetcbmessage

Purpose

Set a user callback to be called whenever Xpress-SLP outputs a line of text

Synopsis

```
int XPRS_CC XSLPsetcbmessage(XSLPprob Prob, void (XPRS_CC *UserFunc)
    (XSLPprob myProb, void *myObject, char *msg, int len, int msgtype),
    void *Object);
```

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called whenever Xpress-SLP outputs a line of text. UserFunc does not return a value.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbmessage.
msg	Character buffer holding the string to be output.
len	Length in characters of msg excluding the null terminator.
msgtype	Type of message. The following are system-defined: 1 Information message 3 Warning message 4 Error message A negative value indicates that the Optimizer is about to finish and any buffers should be flushed at this time. User-defined values are also possible for msgtype which can be passed using <code>XSLPprintmsg</code>
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

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 len, 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](#).

Visual Basic users must use [XSLPtoVBString](#) or an equivalent to convert `msg` into a VB-type string.

When used with VB, the callback function has the prototype:

```

Public Sub mycbfunc (ByVal prob As Long, ByVal obj As Long,
                    ByVal msg As Long, ByVal length As Long,
                    ByVal msgtype As Long)

```

Related topics

[XSLPsetcbmessageF](#), [XSLPsetlogfile](#), [XSLPtoVBString](#)

XSLPsetcbmessageF

Purpose

Set a user callback to be called whenever Xpress-SLP outputs a line of text

Synopsis

```
int XPRS_CC XSLPsetcbmessageF(XSLPprob Prob, void (XPRS_CC *UserFunc)
    (XSLPprob myProb, void *myObject, char *msg, int *len, int *msgtype),
    void *Object);
```

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called whenever Xpress-SLP outputs a line of text. UserFunc does not return a value.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbmessage.
msg	Character buffer holding the string to be output.
len	Address of an integer holding the length in characters of msg excluding the null terminator.
msgtype	Address of an integer holding the type of message. The following are system-defined: 1 Information message 3 Warning message 4 Error message A negative value indicates that the Optimizer is about to finish and any buffers should be flushed at this time. User-defined values are also possible for msgtype which can be passed using XSLPprintmsg
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

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 *plen, int *pmsgtype) {
    FILE *logfile;
    int len, msgtype;
    logfile = (FILE *) Obj;
    len = *plen;
    msgtype = *pmsgtype;
    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](#).

Visual Basic users must use [XSLPtoVBString](#) or an equivalent to convert `msg` into a VB-type string.

`XSLPsetcbmessageF` is identical to [XSLPsetcbmessage](#) except that the callback function receives the message length and type by reference rather than by value.

Related topics

[XSLPsetcbmessage](#), [XSLPsetlogfile](#), [XSLPtoVBString](#)

XSLPsetcboptnode

Purpose

Set a user callback to be called during MISLP when an optimal SLP solution is obtained at a node

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called when an optimal SLP solution is obtained at a node. <code>UserFunc</code> 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).
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcboptnode</code> .
feas	Address of an integer containing the feasibility flag. If <code>UserFunc</code> sets the flag nonzero, the node is declared infeasible.
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</code> .

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 myProb, void *Obj, int *feas) {
    double Total, ObjVal;
    XSLPgetdblattrib(myProb, XSLP_ERRORCOSTS, &Total);
    XSLPgetdblattrib(myProb, XSLP_OBJVAL, &ObjVal);
    if (fabs(Total) > fabs(ObjVal) * 0.001 &&
        fabs(Total) > 1) *feas = 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. When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

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

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called after the set-up of the SLP problem to be solved at a node. <code>UserFunc</code> 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).
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcbprenode</code> .
feas	Address of an integer containing the feasibility flag. If <code>UserFunc</code> sets the flag nonzero, the node is declared infeasible.
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</code> .

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 myProb, void *Obj, int *feas) {
    XPRSProb xprob;
    int i, *IntList;
    double LO, UP;
    IntList = (int *) Obj;
    XSLPgetptrattrib(myProb, 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. When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPsetchoptnode](#), [XSLPsetcbslnode](#)

XSLPsetcbslpend

Purpose

Set a user callback to be called at the end of the SLP optimization

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called at the end of the SLP optimization. <code>UserFunc</code> 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.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcbslpend</code> .
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</code> .

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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPsetcbslpstart](#)

XSLPsetcbslpnode

Purpose

Set a user callback to be called during MISLP after the SLP optimization at each node.

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called after the set-up of the SLP problem to be solved at a node. <code>UserFunc</code> 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).
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as <code>Object</code> to <code>XSLPsetcbslpnode</code> .
feas	Address of an integer containing the feasibility flag. If <code>UserFunc</code> sets the flag nonzero, the node is declared infeasible.
Object	Address of a user-defined object, which can be used for any purpose by the function. <code>Object</code> is passed to <code>UserFunc</code> as <code>myObject</code> .

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 myProb, void *Obj, int *feas) {
    double TargetValue, LPValue;
    XSLPgetdblattrib(prob, XPRS_LPOBJVAL, &LPValue);
    TargetValue = * (double *) Obj;
    if (LPValue < TargetValue) *feas = 1;
    return 0;
}
```

Further information

If a node can be cut off by the callback function, then further exploration of the node is avoided. When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPsetcboptnode](#), [XSLPsetcbprenode](#)

XSLPsetcbstart

Purpose

Set a user callback to be called at the start of the SLP optimization

Synopsis

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

Arguments

Prob	The current SLP problem.
UserFunc	The function to be called at the start of the SLP optimization. UserFunc returns an integer value. If the return value is nonzero, the optimization will not be carried out.
myProb	The problem passed to the callback function.
myObject	The user-defined object passed as Object to XSLPsetcbstart.
Object	Address of a user-defined object, which can be used for any purpose by the function. Object is passed to UserFunc as myObject.

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;
XSLPsetcbstart(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.

When used with VB, the callback function has the prototype:

```
Public Function mycbfunc (ByVal prob As Long, ByVal object As Long) As Long
```

Related topics

[XSLPsetcbpend](#)

XSLPsetdblcontrol

Purpose

Set the value of a double precision problem control

Synopsis

```
int XPRS_CC XSLPsetdblcontrol(XSLPprob Prob, int Param, double dValue);
```

Arguments

Prob	The current SLP problem.
Param	control (SLP or optimizer) whose value is to be returned.
dValue	Double precision value to be set.

Example

The following example sets the value of the Xpress-SLP 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`.

Related topics

`XSLPgetdblcontrol`, `XSLPsetintcontrol`, `XSLPsetstrcontrol`

XSLPsetdefaultcontrol

Purpose

Set the values of one SLP control to its default value

Synopsis

```
int XPRS_CC XSLPsetdefaultcontrol(XSLPprob Prob, int Param);
```

Arguments

Prob	The current SLP problem.
Param	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 and then reads and optimizes another problem using the default setting.

```
XSLPreadprob(Prob, "Matrix1", "");  
XSLPsetintcontrol(Prob, XSLP_LOG, 4);  
XSLPmaxim(Prob, "");  
XSLPsetdefaultcontrol(Prob, XSLP_LOG);  
XSLPreadprob(Prob, "Matrix2", "");  
XSLPmaxim(Prob, "");
```

Further information

This function cannot reset the optimizer controls. Use `XPRSsetdefaults` or `XPRSsetdefaultcontrol` as well to reset optimizer controls to their default values.

Related topics

[XSLPsetdblcontrol](#), [XSLPsetdefaults](#), [XSLPsetintcontrol](#), [XSLPsetstrcontrol](#)

XSLPsetdefaults

Purpose

Set the values of all SLP controls to their default values

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 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);
XSLPmaxim(Prob, "");
XSLPsetdefaults(Prob);
XSLPreadprob(Prob, "Matrix2", "");
XSLPmaxim(Prob, "");
```

Further information

This function does not reset the optimizer controls. Use `XPRSsetdefaults` as well to reset all the controls to their default values.

Related topics

[XSLPsetdblcontrol](#), [XSLPsetintcontrol](#), [XSLPsetstrcontrol](#)

XSLPsetfuncobject

Purpose

Change the address of one of the objects which can be accessed by the user functions

Synopsis

```
int XPRS_CC XSLPsetfuncobject(int *ArgInfo, int ObjType, void *Address)
```

Arguments

ArgInfo	The array of argument information for the user function.
ObjType	An integer indicating which object is to be changed XSLP_GLOBALFUNCOBJECT The <i>Global Function Object</i> ; XSLP_USERFUNCOBJECT The <i>User Function Object</i> for the function; XSLP_INSTANCEFUNCOBJECT The <i>Instance Function Object</i> for the instance of the function.
Address	The address of the object.

Example

The following example from within a user function checks if there is a function instance. If so, it gets the *Instance Function Object*. If it is `NULL` an array is allocated and its address is saved as the new *Instance Function Object*.

```
int Instance;  
XSLPgetfuncinfo(ArgInfo, NULL, NULL, NULL, NULL,  
                NULL, NULL, NULL, &Instance);  
if (Instance) {  
    XSLPgetfuncobject(ArgInfo, XSLP_INSTANCEFUNCOBJECT,  
                     &Object);  
    if (Object == NULL) {  
        Object = calloc(4*3, sizeof(double));  
        XSLPsetfuncobject(ArgInfo, XSLP_INSTANCEFUNCOBJECT,  
                          Object);  
    }  
}
```

Further information

This function changes the address of one of the objects which can be accessed by any user function. It requires the `ArgInfo` array of argument information. This is normally provided as one of the arguments to a user function, or it can be created by using the function

[XSLPsetuserfuncinfo](#)

The identity of the function and the instance are obtained from the `ArgInfo` array. Within a user function, therefore, using the `ArgInfo` array passed to the user function will change the objects accessible to that function.

If, instead, `XSLPsetfuncobject` is used with an array which has been populated by [XSLPsetuserfuncinfo](#), the *Global Function Object* can be set as usual. The *User Function Object* cannot be set (use [XSLPchguserfuncobject](#) for this purpose). There is no *Instance Function Object* as such; however, a value can be set by `XSLPsetfuncobject` which can be used by the function subsequently called by [XSLPcalluserfunc](#). It is the user's responsibility to manage the object and save and restore the address as necessary, because Xpress-SLP will not retain the information itself.

If `Address` is `NULL`, then the corresponding information will be unchanged.

Related topics

[XSLPchgfuncobject](#), [XSLPchguserfuncobject](#), [XSLPgetfuncobject](#),
[XSLPgetuserfuncobject](#), [XSLPsetuserfuncobject](#)

XSLPsetfunctionerror

Purpose

Set the function error flag for the problem

Synopsis

```
int XPRS_CC XSLPsetfunctionerror(XSLPprob Prob);
```

Argument

Prob The current SLP problem.

Example

The following example from within a user function sets the function error flag if there is an error during the function evaluation:

```
double XPRS_CC ProfitCalc(double *Value, int *ArgInfo) {
    XSLPprob Prob;
    double Factor, Size;
    Factor = Value[0];
    Size = Value[1];
    if (Factor < 0) {
        XSLPgetfuncobject(ArgInfo, XSLP_XSLPPROBLEM, &Prob);
        XSLPsetfunctionerror(Prob);
        return 0.0;
    }
    return pow(Factor, Size);
}
```

Note the use of `XSLPgetfuncobject` to retrieve the Xpress-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.

XSLPsetintcontrol

Purpose

Set the value of an integer problem control

Synopsis

```
int XPRS_CC XSLPsetintcontrol(XSLPprob Prob, int Param, int iValue);
```

Arguments

Prob	The current SLP problem.
Param	control (SLP or optimizer) whose value is to be returned.
iValue	The value to be set.

Example

The following example sets the value of the Xpress-SLP 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 requested, the return value will be the same as that from `XPRSsetintcontrol`.

Related topics

`XSLPgetintcontrol`, `XSLPsetdblcontrol`, `XSLPsetintcontrol`, `XSLPsetstrcontrol`

XSLPsetlogfile

Purpose

Define an output file to be used to receive messages from Xpress-SLP

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 (<code>Option=0</code>) or (in console mode) to the console as well (<code>Option=1</code>).

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](#), [XSLPsetcbmessageF](#)

XSLPsetparam

Purpose

Set the value of a control parameter by name

Synopsis

```
int XPRS_CC XSLPsetparam(XSLPprob Prob, const char *Param,
    const char *cValue);
```

Arguments

Prob	The current SLP problem.
Param	Name of the control or attribute whose value is to be returned.
cValue	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-SLP 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

Synopsis

```
int XPRS_CC XSLPsetstrcontrol(XSLPprob Prob, int Param,
    const char *cValue);
```

Arguments

Prob	The current SLP problem.
Param	control (SLP or optimizer) whose value is to be returned.
cValue	Character buffer containing the value.

Example

The following example sets the value of the Xpress-SLP control `XSLP_CVNAME` and of the optimizer control `XPRS_MPSSOBJNAME`:

```
XSLPsetstrcontrol(Prob, XSLP_CVNAME, "CharVars");
XSLPsetstrcontrol(Prob, XPRS_MPSSOBJNAME, "_OBJ_");
```

Further information

Both SLP and optimizer controls can be set using this function. If an optimizer control is requested, the return value will be the same as that from `XPRSsetstrcontrol`.

Related topics

[XSLPgetstrcontrol](#), [XSLPsetdblcontrol](#), [XSLPsetintcontrol](#), [XSLPsetstrcontrol](#)

XSLPsetstring

Purpose

Set a value in the Xpress-SLP string table

Synopsis

```
int XPRS_CC XSLPsetstring(XSLPprob Prob, int *Param, const char *cValue);
```

Arguments

Prob	The current SLP problem.
Param	Address of an integer to receive the index of the string in the Xpress-SLP string table.
cValue	Value to be set.

Example

The following example puts the current date and time into the Xpress-SLP string table and later recovers and prints it:

```
int iTime;
char *Buffer[200];
time_t Time;
time(&Time);
XSLPsetstring(Prob, &iTime, ctime(Time));
...
XSLPgetstring(Prob, iTime, Buffer);
printf("\nStarted at %s", Buffer);
```

Further information

`XSLPsetstring` provides a convenient way of passing string information between routines by means of integer indices.

Related topics

[XSLPgetstring](#)

XSLPsetuniqueprefix

Purpose

Find a prefix character string which is different from all the names currently in use within the SLP problem

Synopsis

```
int XPRS_CC XSLPsetuniqueprefix(XSLPprob Prob);
```

Argument

Prob The current SLP problem.

Example

The following example reads a problem from file and then finds a unique prefix so that new names can be added without fear of duplications:

```
char Prefix[20];
XSLPreadprob(Prob, "Matrix", "");
XSLPsetuniqueprefix(Prob);
XSLPgetstrattrib(Prob, XSLP_UNIQUEPREFIX, Prefix);
printf("\nNo names start with %s", Prefix);
```

Further information

The unique prefix may be more than one character in length, and may change if new names are added to the problem. The value of the unique prefix can be obtained from the string attribute `XSLP_UNIQUEPREFIX`.

XSLPsetuserfuncaddress

Purpose

Change the address of a user function

Synopsis

```
int XPRS_CC XSLPsetuserfuncaddress(XSLPprob Prob, int nSLPUF,
    void *Address);
```

Arguments

Prob	The current SLP problem.
nSLPUF	The index of the user function.
Address	The address of the user function.

Example

The following example defines a user function via [XSLPchguserfunc](#) and then re-defines the address.

```
double InternalFunc(double *, int *);
int nUF;

XSLPchguserfunc(Prob, 0, NULL, 023, 1,
    NULL, NULL, NULL);

XSLPgetintattrib(Prob, XSLP_UFS, &nUF);
XSLPaddnames(Prob, XSLP_USERFUNCNAMES, "Func1",
    nUF, nUF);

XSLPsetuserfuncaddress(Prob, nUF, InternalFunc);
```

Note that `InternalFunc` is defined as taking two arguments (`double*` and `int*`). This matches the `ArgType` setting in [XSLPchguserfunc](#). The external function name is `NULL` because it is not required when the address is given.

Further information

`nSLPUF` is an Xpress-SLP index and always counts from 1. It is not affected by the setting of `XPRS_CSTYLE`.

The address of the function is changed to the one provided. `XSLPsetuserfuncaddress` should only be used for functions declared as of type `DLL` or `VB`. Its main use is where a user function is actually internal to the system rather than being provided in an external library. In such a case, the function is initially defined as an external function using [XSLPloaduserfuncs](#), [XSLPadduserfuncs](#) or [XSLPchguserfunc](#) and the address of the function is then provided using `XSLPsetuserfuncaddress`.

Related topics

[XSLPadduserfuncs](#) [XSLPchguserfunc](#), [XSLPchguserfuncaddress](#) [XSLPgetuserfunc](#), [XSLPloaduserfuncs](#)

XSLPsetuserfuncinfo

Purpose

Set up the argument information array for a user function call

Synopsis

```
int XPRS_CC XSLPsetuserfuncinfo(XSLPprob Prob, int *ArgInfo,
    int CallerFlag, int nInput, int nReturn, int nDelta, int nInString,
    int nOutString);
```

Arguments

Prob	The current SLP problem.
ArgInfo	The array to be set up. This must be dimensioned at least XSLP_FUNCINFOSIZE.
CallerFlag	An integer which can be used for any purpose to communicate between the calling and called program. This value will always be zero for user functions which are called directly by Xpress-SLP.
nInput	The number of input values.
nReturn	The number of return values required.
nDelta	The number of sets of partial derivatives required.
nInString	The number of strings contained in the ARGNAME argument to the user function.
nOutString	The number of strings contained in the RETNAME argument to the user function .

Example

The following example sets up the argument information array and then calls the user function ProfitCalc:

```
int ArgInfo[XSLP_FUNCINFOSIZE];
double Values[3];
int iFunc;

XSLPsetuserfuncinfo(Prob, ArgInfo, 99, 3, 1,
    0, 0, 0);
XSLPgetindex(Prob, XSLP_USERFUNCNAMESNOCASE,
    "PackCalcs", &iFunc);
Result = XSLPcalluserfunc(Prob, iFunc, Values,
    ArgInfo, NULL, NULL, NULL, NULL);
```

The function is called with 3 values in Value and expects 1 return value. There are no names expected by the function.

Further information

The total number of values returned will be $(nReturn) * (nDelta + 1)$.

Related topics

[XSLPchgfuncobject](#), [XSLPgetfuncobject](#), [XSLPsetfuncobject](#), [XSLPcalluserfunc](#)

XSLPsetuserfuncobject

Purpose

Set or define one of the objects which can be accessed by the user functions

Synopsis

```
int XPRS_CC XSLPsetuserfuncobject(XSLPprob Prob, int Entity,
    void *Address);
```

Arguments

Prob	The current SLP problem.
Entity	An integer indicating which object is to be defined. The value is interpreted as follows: 0 The <i>Global Function Object</i> ; n > 0 The <i>User Function Object</i> for user function number n; n < 0 The <i>Instance Function Object</i> for user function instance number -n.
Address	The address of the object.

Example

The following example sets the *Global Function Object*. It then sets the *User Function Object* for the function ProfitCalcs.

```
double *GlobObj;
void *ProfitObj;
int iUF;

XSLPsetuserfuncobject(Prob, 0, GlobObj);

if (!XSLPgetindex(Prob, XSLP_USERFUNCNAMESNOCASE,
    "ProfitCalcs", &iUF)) {
    XSLPsetuserfuncobject(Prob, iUF, ProfitObj);
}
```

The function objects can be of any type. The index of the user function is obtained using the case-insensitive search for names. If the name is not found, [XSLPgetindex](#) returns a nonzero value.

Further information

As instance numbers are not normally meaningful, this function should only be used with a negative value of n to reset all *Instance Function Objects* to NULL when a model is being re-optimized within the same program execution.

Related topics

[XSLPchgfuncobject](#), [XSLPchguserfuncobject](#), [XSLPsetfuncobject](#)

XSLPtime

Purpose

Print the current date and time

Synopsis

```
int XPRS_CC XSLPtime(XSLPprob Prob);
```

Argument

Prob The current SLP problem.

Example

The following example prints the date and time before and after reading a problem from file:

```
XSLPtime(Prob);  
XSLPreadprob(Prob, "Matrrix1", "");  
XSLPtime(Prob);
```

Further information

The current date and time are output in accordance with the current settings from `XSLPsetlogfile` and any user message callback function.

Related topics

[XSLPgetdttime](#), [XSLPgettime](#), [XSLPsetcbmessage](#), [XSLPsetcbmessageF](#), [XSLPsetlogfile](#)

XSLPtokencount

Purpose

Count the number of tokens in a free-format character string

Synopsis

```
int XPRS_CC XSLPtokencount(const char *Record);
```

Argument

`Record` The character string to be processed. This must be terminated with a null character.

Return value

The number of tokens (strings separated by one or more spaces) in `Record`.

Example

The following example counts the number of tokens in the string "sin (x + y)":

```
int nToken;  
nToken = XSLPcounttokens("sin ( x + y )");
```

Further information

`Record` should follow the conventions for Extended MPS Format, with each token being separated by one or more spaces from the previous token.

Related topics

[XSLPqparse](#)

XSLPtoVBString

Purpose

Return a string to VB given its address in Xpress-SLP

Synopsis

```
Function XPRS_CC XSLPtoVBString (adrstr As Long) As String
```

Argument

adrstr The address of a string in Xpress-SLP

Return value

Example

The following example shows how `XSLPtoVBString` can be used in a VB message callback:

```
Public Sub messagecb(ByVal prob As Long,
                    ByVal object As Long, ByVal msg As Long,
                    ByVal length As Long, ByVal msgtype As Long)
    Debug.Print("LOG: " & XSLPtoVBString(msg))
End Sub
```

This prints any messages to the immediate window in the VB editor.

Further information

This function is part of the VB interface and is not available (or indeed useful) in the standard Xpress-SLP API.

If `XSLPsetcbmessage` is used from VB, then the message pointer which is returned to the callback function must be converted using `XSLPtoVBString` before it can be used.

Related topics

[XSLPsetcbmessage](#), [XSLPsetcbmessageF](#)

XSLPuprintmemory

Purpose

Print the dimensions and memory allocations for a problem

Synopsis

```
int XPRS_CC XSLPuprintmemory(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", "");
XSLPuprintmemory(Prob);
```

The output is similar to the following:

Arrays and dimensions:

Array	Item Size	Used Items	Max Items	Allocated Memory	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

XSLPuprintmemory 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.

XSLPuserfuncinfo

Purpose

Get or set user function declaration information

Synopsis

```
int XSLP_CC XSLPuserfuncinfo(XSLPprob prob, int iFunc, int InfoType,
    void *Info);
```

Arguments

Prob	The current SLP problem.
iFunc	Index of the user function
InfoType	Type of information to be set or retrieved
Info	Address of information to be set or retrieved

Example

The following example sets the external name of user function number 4 to "ANewFunc":

```
XSLPuserfuncinfo(Prob, 4, XSLP_SETUFNAME, "ANewFunc");
```

Further information

This function allows the setting or retrieving of individual items for a user function. The following constants are provided for user function handling:

XSLP_GETUFNAME	Retrieve the external name of the user function
XSLP_GETUFPARAM1	Retrieve the first string parameter
XSLP_GETUFPARAM2	Retrieve the second string parameter
XSLP_GETUFPARAM3	Retrieve the third string parameter
XSLP_GETUFARGTYPE	Retrieve the argument types
XSLP_GETUFEXETYPE	Retrieve the linkage type
XSLP_SETUFNAME	Set the external name of the user function
XSLP_SETUFPARAM1	Set the first string parameter
XSLP_SETUFPARAM2	Set the second string parameter
XSLP_SETUFPARAM3	Set the third string parameter
XSLP_SETUFARGTYPE	Set the argument types
XSLP_SETUFEXETYPE	Set the linkage type

For information which sets or retrieves character string information, `Info` is the string to be used or a buffer large enough to hold the string to be retrieved.

For other information, `Info` is the address of an integer containing the information or to receive the information.

Related topics

[XSLPadduserfuncs](#), [XSLPchguserfunc](#), [XSLPgetuserfuncs](#), [XSLPloaduserfuncs](#)

XSLPvalidformula

Purpose

Check a formula in internal (parsed or unparsed) format for unknown tokens

Synopsis

```
int XPRS_CC XSLPvalidformula(int *inType, double *inValue, int *nToken,
    char *Name, char *StringTable);
```

Arguments

<code>inType</code>	Array of token types providing the formula.
<code>inValue</code>	Array of values corresponding to the types in <code>inType</code>
<code>nToken</code>	Number of the first invalid token in the formula. A value of zero means that the formula is valid. May be <code>NULL</code> if not required.
<code>Name</code>	Character buffer to hold the name of the first invalid token. May be <code>NULL</code> if not required.
<code>StringTable</code>	Character buffer holding the names of the unidentified tokens (this can be created by <code>XSLPpreparseformula</code>).

Example

The following example pre-parses the formula "sin (x + y)" and then tries to identify the unknown tokens:

```
int n, Index, NewType, Type[20];
double Value[20];
char Strings[200], Name[20];
XSLPpreparseformula(Prob, "sin ( x + y )", NULL,
    Type, Value, Strings, NULL);
for (;;) {
    XSLPvalidformula(&Type[n], &Value[n], &n, Name, Strings);
    if (n == 0) break;
    Index = 0;
    if (Type[n+1] == XSLP_LB) { /* function */
        NewType = XSLP_IFUN;
        XSLPgetindex(Prob, XSLP_INTERNALFUNCNAME$NOCASE,
            Name, &Index);
    }
    else { /* try for column */
        NewType = XSLP_VAR;
        XSLPgetindex(Prob, 2, Name, &Index);
    }
    if (Index) {
        Type[n] = NewType; Value[n] = Index;
    }
    else {
        printf("\nUnidentified token %s",Name);
        break;
    }
}
```

`XSLPpreparseformula` converts the formula into unparsed internal format.

`XSLPvalidformula` then checks forward from the last invalid token and tries to identify it as an internal function (followed by a left bracket) or as a column (otherwise). If it cannot be identified, the checking stops with an error message. Otherwise, the token type and value are updated and the procedure continues.

Related topics

[XSLPprepareformula](#)

XSLPvalidate

Purpose

Validate the feasibility of constraints in a converged 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 $D < XSLP_VALIDATIONTOL_A$

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 $D < MAX(ABS(TPos), ABS(TNeg)) * XSLP_VALIDATIONTOL_R$

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`

XSLPwriteprob

Purpose

Write the current problem to a file in extended MPS or text format

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 "text" format when ".txt" will be appended.								
Flags	The following flags can be used: <table><tr><td>l</td><td>(lower-case "L") write the 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>s</td><td>"scrambled" names (the default is to use the names provided on input);</td></tr><tr><td>t</td><td>write the matrix in "text" (the default is to write extended MPS format).</td></tr></table>	l	(lower-case "L") write the 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);	s	"scrambled" names (the default is to use the names provided on input);	t	write the matrix in "text" (the default is to write extended MPS format).
l	(lower-case "L") write the 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);								
s	"scrambled" names (the default is to use the names provided on input);								
t	write the matrix in "text" (the default is to write extended MPS format).								

Example

The following example reads a problem from file, augments it and writes the augmented (linearized) matrix in text form to file "output.txt":

```
XSLPreadprob(Prob, "Matrix", "");  
XSLPconstruct(Prob);  
XSLPwriteprob(Prob, "output", "lt");
```

Further information

The `t` flag is used to produce a "human-readable" form of the problem. It is similar to the `lp` format of `XPRSwriteprob`, but does not contain all the potential complexities of the Extended MPS Format, so the resulting file cannot be used for input. A quadratic objective is written with its true coefficients (not scaled by 2 as in the equivalent `lp` format).

Related topics

[XSLPreadprob](#)

Chapter 11

Internal Functions

Xpress-SLP 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-SLP internal functions:

ABS	Absolute value	p. 327
ACT	Activity (left hand side) of a row	p. 344
ARCCOS	Arc cosine trigonometric function	p. 320
ARCSIN	Arc sine trigonometric function	p. 321
ARCTAN	Arc tangent trigonometric function	p. 322
COS	Cosine trigonometric function	p. 323
DJ	Reduced cost (DJ) of a column	p. 345
EQ	Equality test	p. 335
EXP	Exponential function (e raised to the power)	p. 328
GE	Greater than or equal test	p. 336
GT	Greater than test	p. 337
IAC	Gasoline blending interaction coefficients	p. 355
IF	Zero/nonzero test	p. 338
INTERP	General-purpose interpolation	p. 356
LE	Less than or equal test	p. 339
LN	Natural logarithm	p. 329
LO	Lower bound of a column	p. 346
LOG, LOG10	Logarithm to base 10	p. 330
LT	Less than test	p. 340
MATRIX	Current matrix entry	p. 347
MAX	Maximum value of an arbitrary number of items	p. 331
MIN	Minimum value of an arbitrary number of items	p. 332

MV	Marginal value of a row	p. 348
NE	Inequality test	p. 341
NOT	Logical inversion	p. 342
PARAM	Value of a numeric attribute or control	p. 349
RHS	Right hand side of a row	p. 350
RHSRANGE	Range (upper limit minus lower limit of the right side) of a row	p. 351
SIN	Sine trigonometric function	p. 324
SLACK	Slack activity of a row	p. 352
SQRT	Square root	p. 333
TAN	Tangent trigonometric function	p. 325
UP	Upper bound of a column	p. 353

11.1 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; or an XV with only one item

Return value

A value in the range 0 to $+\pi$.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
ARCCOS (0.99)
ARCCOS (A)
ARCCOS (B^2)
ARCCOS (SQRT (A) )
ARCCOS (XVA)
ARCCOS (XVB)
```

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; or an XV with only one item

Return value

A value in the range $-\pi / 2$ to $+\pi / 2$.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
ARCSIN (0.99)
ARCSIN (A)
ARCSIN (B^2)
ARCSIN (SQRT (A) )
ARCSIN (XVA)
ARCSIN (XVB)
```

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; or an XV with only one item

Return value

A value in the range $-\pi / 2$ to $+\pi / 2$.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
ARCTAN (99)
ARCTAN (A)
ARCTAN (B^2)
ARCTAN (SQRT (A) )
ARCTAN (XVA)
ARCTAN (XVB)
```

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; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
COS (99)
COS (A)
COS (B^2)
COS (SQRT (A) )
COS (XVA)
COS (XVB)
```

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; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
SIN (99)
SIN (A)
SIN (B^2)
SIN (SQRT (A) )
SIN (XVA)
SIN (XVB)
```

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; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
TAN (99)
TAN (A)
TAN (B^2)
TAN (SQRT (A) )
TAN (XVA)
TAN (XVB)
```

11.2 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; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
ABS (99)
ABS (A)
ABS (B^2)
ABS (SQRT (A) )
ABS (XVA)
ABS (XVB)
```

Further information

ABS is not always differentiable and so alternative modeling approaches should be used where possible.

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; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
EXP (99)
EXP (A)
EXP (B^2)
EXP (SQRT (A) )
EXP (XVA)
EXP (XVB)
```

LN

Purpose

Natural logarithm

Synopsis

LN (value)

Argument

value

One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
LN ( 99 )
LN ( A )
LN ( B ^ 2 )
LN ( SQRT ( A ) )
LN ( XVA )
LN ( XVB )
```

Further information

value must be strictly positive (greater than 1.0E-300).

LOG, LOG10

Purpose

Logarithm to base 10

Synopsis

LOG (value)
LOG10 (value)

Argument

value One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
LOG (99)
LOG10 (99)
LOG (A)
LOG10 (A)
LOG (B^2)
LOG10 (B^2)
LOG (SQRT (A) )
LOG10 (SQRT (A) )
LOG (XVA)
LOG10 (XVA)
LOG (XVB)
LOG10 (XVB)
```

Further information

value must be strictly positive (greater than 1.0E-300).

MAX

Purpose

Maximum value of an arbitrary number of items

Synopsis

`MAX(value1, value2, ...)`

Argument

`value1, ...` Each argument is one of the following: a constant; a variable; a formula evaluating to a single value; or an XV

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVB: = = = B ^ 2
          = = = A * B
```

then the following are all valid uses of the function:

```
MAX(A, 99)
MAX(A, B, 99)
MAX(A, B^2)
MAX(SQRT(A), B)
MAX(XVB)
```

Further information

`MAX` is not always differentiable and so alternative modeling approaches should be used where possible.

If an XV is used as an argument to the function, then all members of the XV will be included.

MIN

Purpose

Minimum value of an arbitrary number of items

Synopsis

`MIN(value1, value2, ...)`

Argument

`value1, ...` Each argument is one of the following: a constant; a variable; a formula evaluating to a single value; or an XV

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVB: = = = B ^ 2
          = = = A * B
```

then the following are all valid uses of the function:

```
MIN(A, 99)
MIN(A, B, 99)
MIN(A, B^2)
MIN(SQRT(A), B)
MIN(XVB)
```

Further information

`MIN` is not always differentiable and so alternative modeling approaches should be used where possible.

If an XV is used as an argument to the function, then all members of the XV will be included.

SQRT

Purpose

Square root

Synopsis

SQRT (value)

Argument

value

One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
SQRT (99)
SQRT (A)
SQRT (B^2)
SQRT (SQRT (A) )
SQRT (XVA)
SQRT (XVB)
```

Further information

value must be non-negative.

11.3 Logical functions

The logical functions all return 0 for "false" and 1 for "true". They are implemented so that complementary functions are never both true or both false.

For example:

exactly one of $EQ(X, Y)$ and $NE(X, Y)$ is true;

exactly one of $LT(X, Y)$ and $GE(X, Y)$ is true;

exactly one of $IF(X)$ and $NOT(X)$ is true;

if $LE(X, Y)$ is true, then exactly one of $LT(X, Y)$ and $EQ(X, Y)$ is true.

Equality tests are carried out using the tolerances $XSLP_EQTOL_A$ and $XSLP_EQTOL_R$. If

$abs(X - Y) < XSLP_EQTOL_A$ or

$abs(X - Y) < abs(X) * XSLP_EQTOL_R$

then X and Y are regarded as equal.

Functions IF and NOT test for zero using tolerance $XSLP_EQTOL_A$.

Because of these tolerances, it is possible that $EQ(X, Y)$ and $EQ(Y, Z)$ are both true, but $EQ(X, Z)$ is false. Where multiple tests of this type are being carried out, they should all test against the same value if possible.

Logical functions are not continuous or differentiable, and should be used with care in coefficients. Alternative modeling approaches should be used where possible.

EQ

Purpose

Equality test

Synopsis

`EQ(value1, value2)`

Arguments

<code>value1</code>	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item
<code>value2</code>	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if `value1` is not equal to `value2` within tolerance;
1 ("true") if `value1` is equal to `value2` within tolerance.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
EQ(A, 99)
EQ(A, B)
EQ(A, B^2)
EQ(XVB, SQRT(A))
EQ(XVA, XVB)
EQ(99, XVB)
```

GE

Purpose

Greater than or equal test

Synopsis

`GE(value1, value2)`

Arguments

<code>value1</code>	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item
<code>value2</code>	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if `value1` is not greater than or equal to `value2` within tolerance;
1 ("true") if `value1` is greater than or equal to `value2` within tolerance.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
GE(A, 99)
GE(A, B)
GE(A, B^2)
GE(XVB, SQRT(A))
GE(XVA, XVB)
GE(99, XVB)
```

GT

Purpose

Greater than test

Synopsis

`GT(value1, value2)`

Arguments

<code>value1</code>	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item
<code>value2</code>	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if `value1` is not greater than `value2` within tolerance;
1 ("true") if `value1` is greater than `value2` within tolerance.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
GT(A, 99)
GT(A, B)
GT(A, B^2)
GT(XVB, SQRT(A))
GT(XVA, XVB)
GT(99, XVB)
```

IF

Purpose

Zero/nonzero test

Synopsis

IF (value)

Argument

value

One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if `value1` is equal to zero within tolerance;
1 ("true") if `value1` is not equal to zero within tolerance.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
IF (99)
IF (B)
IF (XVB)
IF (EQ (XVA, XVB) +EQ (A, B) )
IF (A-99)
```

LE

Purpose

Less than or equal test

Synopsis

LE (value1, value2)

Arguments

value1	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item
value2	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if value1 is not less than or equal to value2 within tolerance;
1 ("true") if value1 is less than or equal to value2 within tolerance.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
LE (A, 99)
LE (A, B)
LE (A, B^2)
LE (XVB, SQRT (A) )
LE (XVA, XVB)
LE (99, XVB)
```

LT

Purpose

Less than test

Synopsis

LT(value1, value2)

Arguments

value1	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item
value2	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if value1 is not less than value2 within tolerance;
1 ("true") if value1 is less than value2 within tolerance.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
LT(A, 99)
LT(A, B)
LT(A, B^2)
LT(XVB, SQRT(A))
LT(XVA, XVB)
LT(99, XVB)
```

NE

Purpose

Inequality test

Synopsis

NE (value1, value2)

Arguments

value1	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item
value2	One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if value1 is equal to value2 within tolerance;
1 ("true") if value1 is not equal to value2 within tolerance.

Example

Given the following matrix items:

```
Column: A
Column: B
XV      : XVA: A
XV      : XVB: = = = B ^ 2
```

then the following are all valid uses of the function:

```
NE (A, 99)
NE (A, B)
NE (A, B^2)
NE (XVB, SQRT (A) )
NE (XVA, XVB)
NE (99, XVB)
```

NOT

Purpose

Logical inversion

Synopsis

NOT (value)

Argument

value

One of the following: a constant; a variable; a formula evaluating to a single value; or an XV with only one item

Return value

0 ("false") if `value1` is not equal to zero within tolerance;

1 ("true") if `value1` is equal to zero within tolerance.

Example

Given the following matrix items:

Column: A

Column: B

XV : XVA: A

XV : XVB: = = = B ^ 2

then the following are all valid uses of the function:

NOT (99)

NOT (B)

NOT (XVB)

NOT (EQ (XVA, XVB) + EQ (A, B))

NOT (A-99)

11.4 Problem-related functions

The problem-related functions allow access to a limited range of problem and solution data. If they are used in formulae for coefficients they will be regarded as constants (their derivatives will be zero).

Row and column indices used as arguments to the functions always count from 1 and do not respect the value of `XPRS_CSTYLE`.

ACT

Purpose

Activity (left hand side) of a row

Synopsis

ACT (RowIndex)

Argument

RowIndex The index of a row

Example

The following formula starts a delayed constraint when the activity of row 99 becomes greater than 5:

```
DC MyRow 0 = GT ( ACT ( 99 ) , 5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the row can be used instead, and will be translated internally into the row index.

DJ

Purpose

Reduced cost (DJ) of a column

Synopsis

`DJ (ColIndex)`

Argument

`ColIndex` The index of a column

Example

The following formula starts a delayed constraint when the DJ of column 99 becomes greater than 5:

```
DC MyRow 0 = GT ( DJ ( 99 ) , 5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the column can be used instead, and will be translated internally into the column index.

LO

Purpose

Lower bound of a column

Synopsis

`LO (ColIndex)`

Argument

`ColIndex` The index of a column

Example

The following formula starts a delayed constraint when the activity of column MyCol (with index 99) is within 5 of its lower bound:

```
DC MyRow 0 = LT ( MyCol - LO ( 99 ) , 5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the column can be used instead, and will be translated internally into the column index.

MATRIX

Purpose

Current matrix entry

Synopsis

`MATRIX (RowIndex, ColIndex)`

Arguments

<code>RowIndex</code>	The index of a row
<code>ColIndex</code>	The index of a column

Example

The following formula starts a delayed constraint when the value of the coefficient in row 99, column 7 is greater than 5:

```
DC MyRow 0 = GT ( MATRIX ( 99 , 7 ) , 5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the names of the row and column can be used instead, and will be translated internally into the corresponding indices.

MV

Purpose

Marginal value of a row

Synopsis

MV (RowIndex)

Argument

RowIndex The index of a row

Example

The following formula starts a delayed constraint when the marginal value of row 99 becomes greater than 5:

```
DC MyRow 0 = GT ( MV ( 99 ) , 5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the row can be used instead, and will be translated internally into the row index.

PARAM

Purpose

Value of a numeric attribute or control

Synopsis

PARAM(value)

Argument

value

One of the following: a constant; a formula evaluating to a constant; or an XV with only one item which is a constant

Example

The following formula starts a delayed constraint when the SLP iteration count is greater than 5:

```
DC MyRow 0 = GT ( PARAM ( 12001 ) , 5 )
```

Further information

`XSLP_ITER` is number 12001 (see the header file `xslp.h` for the full list of parameters and values. The example shows the use of the formula in Extended MPS format; the same information can also be provided in internal parsed or unparsed format.

RHS

Purpose

Right hand side of a row

Synopsis

`RHS (RowIndex)`

Argument

`RowIndex` The index of a row

Example

The following formula starts a delayed constraint when the slack (right hand side minus left hand side) of row 99 becomes greater than 5:

```
DC MyRow 0 = GT ( RHS ( 99 ) - ACT ( 99 ) , 5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the row can be used instead, and will be translated internally into the row index.

RHSRANGE

Purpose

Range (upper limit minus lower limit of the right side) of a row

Synopsis

RHSRANGE (RowIndex)

Argument

RowIndex The index of a row

Example

The following formula starts a delayed constraint when the slack of row 99 becomes greater than half the RHS range:

```
DC MyRow 0 = GT ( ACT ( 99 ) , 0.5 * RHSRANGE ( 99 ) )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the row can be used instead, and will be translated internally into the row index.

SLACK

Purpose

Slack activity of a row

Synopsis

`SLACK (RowIndex)`

Argument

`RowIndex` The index of a row

Example

The following formula starts a delayed constraint when the slack of row 99 becomes less than 0.5:

```
DC MyRow 0 = LT ( SLACK ( 99 ) , 0.5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the row can be used instead, and will be translated internally into the row index.

UP

Purpose

Upper bound of a column

Synopsis

UP (ColIndex)

Argument

ColIndex The index of a column

Example

The following formula starts a delayed constraint when the activity of column MyCol (with index 99) is within 5 of its upper bound:

```
DC MyRow 0 = LT ( UP ( 99 ) - MyCol , 5 )
```

Further information

When Extended MPS format is used for input of a problem from file, the name of the column can be used instead, and will be translated internally into the column index.

11.5 Specialized functions

The specialized functions are designed for use in particular applications, to reduce the need for custom-built user functions. Notes about their use will be found under the individual functions.

IAC

Purpose
Gasoline blending interaction coefficients

Synopsis
 $IAC(X, V_1, \dots, V_n, C_{12}, C_{13}, \dots, C_{1n}, C_{23}, \dots, C_{2n}, \dots, C_{n-1n})$

Arguments

X	Total quantity.
V_i	Quantities of components 1 to n.
C_{ij}	Interaction coefficient between component i and component j (i < j).

Example
Typically X and V_i will be variables (although the V_i could be provided in an XV), and the interaction coefficients C_{ij} are given in an XV. Given the following matrix items:

```
Column: TotalGas
Columns: Comp1, Comp2, Comp3, Comp4
XV      : XVIA: = = 2.2
              = = 1.1
              = = 0
              = = -1
              = = 0
              = = 2
```

then the following formula calculates the interaction adjustment for the blend:

$= IAC (TotalGas , Comp1 , Comp2 , Comp3 , Comp4 , XVIA)$

Further information
IAC is always differentiated using numerical methods.

INTERP

Purpose
General-purpose interpolation

Synopsis
`INTERP (X, X1, Y1, X2, Y2, . . . , Xn, Yn)`

Arguments
`X` X-value to be interpolated.
`Xi, Yi` Pairs of values for the interpolation. The Xi must be in increasing order.

Example
Typically `X` will be a variable and the interpolation pairs (`Xi, Yi`) are given in an `XV`. Given the following matrix items:

```
Column: Total
XV      : XVI: = = 0
          = = 0
          = = 1
          = = 1
          = = 2
          = = 4
          = = 3
          = = 9
```

then the following formula interpolates `X`:

```
= INTERP ( X, XVI )
```

Further information
In the above example, if `X` has a current value of 1.5, then the function will be evaluated as 2.5 (`X` is halfway between `X = 1` and `X = 2`, so the result is halfway between `Y = 1` and `Y = 4`). As can be seen, the points in this case are the squares of the integers, so the function is approximating the square of `X` by interpolation.

Chapter 12

Xpress-SLP Formulae

Xpress-SLP 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.

12.1 Parsed and unparsed formulae

All formulae input into Xpress-SLP 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.

Type	Description	Value
XSLP_COL	column	index of matrix column. This respects the setting of XPRS_CSTYLE.
XSLP_CON	constant	(double) value.
XSLP_CONSTRAINT	constraint	index of constraint. Note that constraints count from 1, so that the index of matrix row n is $n +$ XPRS_CSTYLE.
XSLP_CV	character variable	index of character variable.
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
XSLP_ROW	row	index of matrix row. This respects the setting of XPRS_CSTYLE.
XSLP_STRING	character string	internal index of character string
XSLP_UNKNOWN	unidentified token	internal index of character string
XSLP_VAR	variable	index of variable. Note that variables count from 1, so that the index of matrix column n is $n +$ XPRS_CSTYLE.
XSLP_VARREF	reference to variable	index of variable. Note that variables count from 1, so that the index of matrix column n is $n +$ XPRS_CSTYLE.
XSLP_XV	extended variable array	index of XV
XSLP_UFARGTYPE	requirements and types of argument for a user function	bitmap of types (see below).
XSLP_UFEXETYPE	linkage of a user function	bitmap of linkage information (see below).
XSLP_XVVARTYPE	type of variable in XV	XSLP_VAR or XSLP_XV
XSLP_XVINTINDEX	index of XV item name	index of name in Xpress-SLP string table

Argument types for user function definition are stored as a bit map. Each type is stored in 3 bits: bits 0-2 for argument 1, bits 3-5 for argument 2 and so on. The possible values for each argument are as follows:

- 0 omitted
- 1 NULL
- 2 INTEGER
- 3 DOUBLE
- 4 VARIANT
- 6 CHAR

The linkage type and other function information are stored as a bit map as follows:

- Bits 0-2 type of linkage:
 - 1 = User library or DLL
 - 2 = Excel spreadsheet
 - 3 = Excel macro
 - 5 = MOSEL
 - 6 = VB
 - 7 = COM
- Bits 3-4 re-evaluation flags:
 - 0 = default
 - 1 (Bit 3) = re-evaluation at each SLP iteration
 - 2 (Bit 4) = re-evaluation when independent variables are outside tolerance
- Bits 6-7 derivative flags:
 - 0 = default
 - 1 (Bit 6) = tangential derivatives
 - 2 (Bit 7) = forward derivatives
- Bit 8 calling mechanism:
 - 0 = standard
 - 1 = CDECL (Windows only)
- Bit 24 set if the function is multi-valued
- Bit 28 set if the function is not differentiable

Token types `XSLP_ROW` and `XSLP_COL` are used only when passing formulae *into* Xpress-SLP. These tokens both respect the setting of `XPRS_CSTYLE`. Any formulae recovered from Xpress-SLP will use the `XSLP_CONSTRAINT` and `XSLP_VAR` token types which always count from 1.

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

<code>XSLP_DEL</code>	comma is optional.
<code>XSLP_FUN</code>	implies a following left-bracket, which is not included explicitly.
<code>XSLP_IFUN</code>	implies a following left-bracket, which is not included explicitly.
<code>XSLP_LB</code>	never used.
<code>XSLP_RB</code>	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 `XSLPgetcccoef` — fill in brackets as required to achieve the correct evaluation order. The result may not match the formula as originally provided.

Token type `XSLP_UNKNOWN` is returned by the parsing routines when a string cannot be identified as any other type of token. Token type `XSLP_STRING` is returned by the parsing routine where the token has been specifically identified as being a character string: the only case where this occurs at present is in the names of return arguments from user-defined multi-valued functions. The "value" field for both these token types is an index into the Xpress-SLP string table and can be accessed using the `XSLPgetstring` function.

12.2 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_VAR	index of x
XSLP_OP	XSLP_EXPONENT
XSLP_CON	2
XSLP_OP	XSLP_PLUS
XSLP_CON	4
XSLP_OP	XSLP_MULTIPLY
XSLP_VAR	index of y
XSLP_OP	XSLP_MULTIPLY
XSLP_LB	0
XSLP_VAR	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:

$x^2 \cdot y^4 \cdot z^3 - \cdot +$

and this is then transcribed as follows:

Type	Value
XSLP_VAR	index of x
XSLP_CON	2
XSLP_OP	XSLP_EXPONENT
XSLP_CON	4
XSLP_VAR	index of y
XSLP_OP	XSLP_MULTIPLY
XSLP_VAR	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.

12.3 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_VAR	index of y
XSLP_OP	XSLP_MULTIPLY
XSLP_FUN	index of MyFunc
XSLP_LB	0
XSLP_VAR	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 MyFunc( *`

and this is then transcribed as follows:

Type	Value
XSLP_VAR	index of <i>y</i>
XSLP_RB	0
XSLP_CON	3
XSLP_DEL	XSLP_COMMA
XSLP_VAR	index of <i>z</i>
XSLP_FUN	index of <i>MyFunc</i>
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.

12.4 Example of a formula involving a complicated function

This example uses a function which takes two arguments and returns an array of results, which are identified by name. In the formula, the return value named `VAL1` is being retrieved.

`y * MyFunc(z, 3 : VAL1)`

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

Type	Value
XSLP_VAR	index of <i>y</i>
XSLP_OP	XSLP_MULTIPLY
XSLP_FUN	index of <i>MyFunc</i>
XSLP_LB	0
XSLP_VAR	index of <i>z</i>
XSLP_DEL	XSLP_COMMA
XSLP_CON	3
XSLP_DEL	XSLP_COLON
XSLP_STRING	index of <i>VAL1</i> in string table
XSLP_RB	0
XSLP_EOF	0

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

`y ) VAL1 : 3 , z MyFunc( *`

and this is then transcribed as follows:

Type	Value
XSLP_VAR	index of <i>y</i>
XSLP_RB	0
XSLP_STRING	index of <i>VAL1</i> in string table
XSLP_DEL	XSLP_COLON
XSLP_CON	3
XSLP_DEL	XSLP_COMMA
XSLP_VAR	index of <i>z</i>
XSLP_FUN	index of <i>MyFunc</i>
XSLP_OP	XSLP_MULTIPLY
XSLP_EOF	0

Notice that the function arguments are in reverse order, including the name of the return value and the colon delimiter, and that a right bracket is used as a delimiter to indicate the end of the argument list.

12.5 Example of a formula defining a user function

User function definitions in `XSLPadduserfuncs` and `XSLPloaduserfuncs` are provided through the formula structure. Assume we wish to add the function defined in Extended MPS format as

MyFunc = *Func1* (*DOUBLE, INTEGER*) *MOSEL* = *MyModel* = *MyArray*

We also want to evaluate the function only when its arguments have changed outside tolerances. This also requires function instances. In the definition of a user function, there is no distinction made between parsed and unparsed format: the tokens provide information and are interpreted in the order in which they are encountered. The function definition is as follows:

Type	Value
XSLP_STRING	index of <i>Func1</i> in string table
XSLP_UFARGTYPE	26 (octal 32)
XSLP_UFEXETYPE	21 (Bit 4 set, and Bits 0-2 = 5)
XSLP_STRING	index of <i>MyModel</i> in string table
XSLP_STRING	index of <i>MyArray</i> in string table
XSLP_EOF	0

The string arguments are interpreted in the order in which they appear. Therefore, if any of the function parameters (*param1* to *param3* in Extended MPS format) is required, there must be entries for the internal function name and any preceding function parameters. If the fields are blank, use an XSLP_STRING token with a zero value.

The name of the function itself (*MyFunc* in this case) is provided through the function `XSLPaddnames`.

12.6 Example of a formula defining an XV

An XV (extended variable array) is defined by its individual items. XV definitions in `XSLPaddxvs` and `XSLPloadxvs` are provided through the formula structure. Assume we wish to add the XV defined in Extended MPS format as:

MyXV *x*
MyXV *y* = *VAR1*
MyXV = = = *x* * 2

Then the definition in parsed format is as follows:

Type	Value
XSLP_XVVARTYPE	XSLP_VAR
XSLP_XVVARINDEX	index of <i>x</i>
XSLP_EOF	0
XSLP_XVVARTYPE	XSLP_VAR
XSLP_XVVARINDEX	index of <i>y</i>
XSLP_XVINTINDEX	index of <i>VAR1</i> in string table
XSLP_EOF	0
XSLP_VAR	index of <i>x</i>
XSLP_CON	2
XSLP_OP	XSLP_MULTIPLY
XSLP_EOF	0

Parsed or unparsed format is only relevant where formulae are being provided (as in the third item above).

12.7 Example of a formula defining a DC

A DC (delayed constraint) can be activated when a certain condition is met by the solution of a preceding linear approximation. The condition is described in a formula which evaluates to zero (if the condition is not met) or nonzero (if the condition is met). Assume we wish to add the DCs described in Extended MPS format as follows:

```
DC ROW1 = GT(x, 1)
DC ROW2 = MV(ROW99)
```

Then the definition in parsed format is as follows:

Type	Value
XSLP_RB	0
XSLP_CON	1
XSLP_DEL	XSLP_COMMA
XSLP_VAR	index of x
XSLP_IFUN	index of GT
XSLP_EOF	0
XSLP_RB	0
XSLP_ROW	index (respecting XPRS_CSTYLE) of ROW99
XSLP_IFUN	index of MV
XSLP_EOF	0

12.8 Formula evaluation and derivatives

In many applications, the same function is used in several matrix entries. Indeed, often the only difference between the entries is the sign of the entry or a difference in (constant) scaling factor. Xpress-SLP separates any constant factor from the formula, and stores a non-linear coefficient as *factor * formula*. In this way, when a formula has been evaluated once, its value can be used repeatedly without the need for re-evaluation.

Xpress-SLP needs partial derivatives of all formulae in order to create the linear approximations to the problem. In the absence of any other information, derivatives are calculated numerically, by making small perturbations of the independent variables and re-evaluating the formulae.

Analytic derivatives will be used if `XSLP_DERIVATIVES` is set to 1. The mathematical operators and the internal functions are differentiated automatically. User functions must provide their own derivatives; if they do not, then derivatives for the functions will be evaluated numerically.

Analytic derivatives need more time to set up, but evaluation of the derivatives is then faster particularly for formulae like:

$$\sum_{i=1}^N f(x_i)$$

Chapter 13

User Functions

13.1 Constant Derivatives

If a user function has constant derivatives with respect to one or more of its arguments, then it is possible to arrange that Xpress-SLP bypasses the repeated evaluation of the function when calculating numerical derivatives for such arguments. There is no benefit in using this feature if the function offers analytic derivatives.

There are two ways of providing constant derivative information to Xpress-SLP:

- **Implicit constant derivatives.**
In this case, Xpress-SLP will initially calculate derivatives as normal. However, if it finds for a particular argument that the "upward" numerical derivative and the "downward" numerical derivative around a point are the same within tolerances, then the derivative for the argument will be marked as constant and will not be re-evaluated. The tolerances `XSLP_CDTOL_A` and `XSLP_CDTOL_R` are used to decide constancy.
- **Interrogate for constant derivatives.**
In this case, Xpress-SLP will call the user function in a special way for each of the arguments in turn. The user function must recognize the special nature of the call and return a value indicating whether the derivative is constant. If the derivative is constant, it will be calculated once in the usual way (numerically), and the result will be used unchanged thereafter.

If a function is marked for interrogation for constant derivatives, then Xpress-SLP will issue a series of special calls the first time that derivatives are required. The only difference from a normal call is that the number of derivatives requested (`FunctionInfo[2]`) will be negative; the absolute value of this number is the number of the argument for which information is required (counting from 1). The single value returned by the function (or in the first element of the return array, depending on the type of function) is zero if the derivative is not constant, or nonzero (normally 1) if the derivative is constant.

The following simple example in C shows how interrogation might be handled:

```
double XPRS_CC MyUserFunc(double *InputValues, int *FunctionInfo) {
    int iArg;
    if ( (iArg=FunctionInfo[2]) < 0) { /* interrogation */
        switch (-iArg) {
            case 1: /* constant with respect to first argument */
            case 4: /* constant with respect to fourth argument */
                return 1.0; /* constant derivative */
            default:
                return 0.0; /* not constant derivative */
        }
    }
    /* normal call */
}
```

```

    }
}

/* normal call for evaluation */
return MyCalc(InputValues);
}

```

13.2 Callbacks and user functions

Callbacks and user functions both provide mechanisms for connecting user-written functions to Xpress-SLP. 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-SLP 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-SLP generally allows problem data to be changed between function calls, and also by callbacks called from within an Xpress-SLP 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.

13.3 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: these are discussed at the end of this section.

Xpress-SLP supports two basic forms of user function. The simple form of function returns a single value, and is treated in essentially the same way as a normal mathematical function. The general form of function returns an array of values and may also perform automatic differentiation.

The main difference between the simple and general form of a user function is in the way the value is returned.

- The simple function calculates and returns one value and is declared as such (for example, `double` in C).
- The general function calculates an array of values. It can either return the array itself (and is declared as such: for example, `double *` in C), or it can return the results in one of the function arguments, in which case the function itself returns a single (double precision) status value (and is declared as such: for example `double` in C).

Values are passed to and from the function in a format dependent on the type of the function and the type of the argument.

- *NULL* format provides a place-holder for the argument but it is a null or empty argument which cannot be used to access or return data. This differs from the omitted argument which does not appear at all.
- *INTEGER* format is used only for the Function Information array (the second argument to the function).
- *DOUBLE* format is used for passing and returning all other numeric values
- *CHAR* format is used for passing character information to the function (input and return variable names)
- *VARIANT* format is used for user functions written in Microsoft Excel, COM and VB. All arguments in Xpress-SLP are then of type *VARIANT*, which is the same as the *Variant* type in COM, VB and Excel VBA. In the function source code, the function itself is declared with all its arguments and return value(s) as *Variant*. *VARIANT* is not available for user functions called through other linkage mechanisms.

13.4 Function Declaration in Xpress-SLP

User functions are declared through the `XSLPloaduserfuncs`, `XSLPadduserfuncs` and `XSLPchguserfunc` functions, or in the `SLPDATA` section of the Extended MPS file format in `UF` type records. These declarations define which of the arguments will actually be made available to the function and (by implication) whether the function can perform automatic differentiation. Simple functions and general functions are declared in the same way. Xpress-SLP recognizes the difference because of the way in which the functions are referenced in formulae.

13.4.1 Function declaration in Extended MPS format

In the `SLPDATA` section of Extended MPS format, the full `UF` record format is:

```
UF Function [= Extname] ( InputValues , FunctionInfo ,
InputNames , ReturnNames , Deltas , ReturnArray )
Linkage = Param1 [ [= Param2 ] = Param3 ]
```

The fields are as follows:

Function	The name of the user function. This is used in the formulae within the problem. A function which returns only one value must return it as a double-precision value. A function which returns multiple values must return a double-precision array, or return the values in the <code>ReturnArray</code> argument. In the latter case, the function must return a single double-precision status value.												
Extname	This field is optional. If it is used, then it is the external name of the function or program when it is called. If the field is omitted, then the same name is used for the internal and external function name. If the name matches the name of a character variable, then the value of the character variable will be used instead. This allows the definition of external names which contain spaces.												
InputValues	DOUBLE or VARIANT or NULL. This is the data type for the input argument list. Use NULL or omit the argument if the data is not required.												
FunctionInfo	INTEGER or VARIANT or NULL. This is the data type for the array of function and argument information. Use NULL or omit the argument if the data is not required. Note that this argument is required if function objects are used by the function.												
InputNames	CHAR or VARIANT or NULL. This is the data type for the names of the input arguments. Use NULL or omit the argument if the data is not required.												
ReturnNames	CHAR or VARIANT or NULL. This is the data type for the names of the return arguments. Use NULL or omit the argument if the data is not required.												
Deltas	DOUBLE or VARIANT or NULL. This is the data type for the perturbations (or differentiation flags). Use NULL or omit the argument if the data is not required.												
ReturnArray	DOUBLE or VARIANT or NULL. This is the data type for the array of results from a multi-valued function. Use NULL or omit the argument if the results are returned directly by the function.												
Linkage	This defines the linkage type and calling mechanism. The following are supported: <table> <tr> <td>DLL</td><td>The function is compiled in a user library or DLL. The name of the file is in the <code>Param1</code> field.</td></tr> <tr> <td>XLS</td><td>The function is in an Excel workbook and communicates through a sheet within the workbook. The name of the workbook is in the <code>Param1</code> field and the name of the sheet is in the <code>Param2</code> field. If <code>Extname</code> is non-blank, it is the name of a macro on the workbook which is to be executed after the data is loaded.</td></tr> <tr> <td>XLF</td><td>The function is in an Excel workbook and communicates directly with Xpress-SLP. The name of the workbook is in the <code>Param1</code> field and the name of the sheet containing the function is in the <code>Param2</code> field.</td></tr> <tr> <td>MOSEL</td><td>This can only be used in conjunction with Xpress-Mosel. See the Xpress Mosel User Guide (Xpress-SLP section) for more information.</td></tr> <tr> <td>VB</td><td>This can only be used when the main program is in VB. The function address must be set up using <code>XSLPchguserfuncaddress</code>.</td></tr> <tr> <td>COM</td><td>This is used for a function compiled into an ActiveX DLL. The <code>PROGID</code> (typically of the form <code>file.class</code>) is in the <code>Param1</code> field.</td></tr> </table> Optionally, the type can be suffixed with additional characters, indicating when the function is to be re-evaluated, what sort of numerical derivatives are to be calculated and what sort of calling mechanism is to be used. The possible types for re-evaluation are:	DLL	The function is compiled in a user library or DLL. The name of the file is in the <code>Param1</code> field.	XLS	The function is in an Excel workbook and communicates through a sheet within the workbook. The name of the workbook is in the <code>Param1</code> field and the name of the sheet is in the <code>Param2</code> field. If <code>Extname</code> is non-blank, it is the name of a macro on the workbook which is to be executed after the data is loaded.	XLF	The function is in an Excel workbook and communicates directly with Xpress-SLP. The name of the workbook is in the <code>Param1</code> field and the name of the sheet containing the function is in the <code>Param2</code> field.	MOSEL	This can only be used in conjunction with Xpress-Mosel. See the Xpress Mosel User Guide (Xpress-SLP section) for more information.	VB	This can only be used when the main program is in VB. The function address must be set up using <code>XSLPchguserfuncaddress</code> .	COM	This is used for a function compiled into an ActiveX DLL. The <code>PROGID</code> (typically of the form <code>file.class</code>) is in the <code>Param1</code> field.
DLL	The function is compiled in a user library or DLL. The name of the file is in the <code>Param1</code> field.												
XLS	The function is in an Excel workbook and communicates through a sheet within the workbook. The name of the workbook is in the <code>Param1</code> field and the name of the sheet is in the <code>Param2</code> field. If <code>Extname</code> is non-blank, it is the name of a macro on the workbook which is to be executed after the data is loaded.												
XLF	The function is in an Excel workbook and communicates directly with Xpress-SLP. The name of the workbook is in the <code>Param1</code> field and the name of the sheet containing the function is in the <code>Param2</code> field.												
MOSEL	This can only be used in conjunction with Xpress-Mosel. See the Xpress Mosel User Guide (Xpress-SLP section) for more information.												
VB	This can only be used when the main program is in VB. The function address must be set up using <code>XSLPchguserfuncaddress</code> .												
COM	This is used for a function compiled into an ActiveX DLL. The <code>PROGID</code> (typically of the form <code>file.class</code>) is in the <code>Param1</code> field.												

A	Function is re-evaluated when input variables change outside strict tolerance
R	Function is re-evaluated every time that input variables change
I	Function always generates function instances.
M	Function is multi-valued.
N	Function is non-differentiable
V	Function can be interrogated to provide some constant derivatives
W	Function may have constant derivatives, which can be deduced by the calling program

If no re-evaluation suffix is provided, then re-evaluation will be determined from the setting of `XSLP_FUNCEVAL`, and function instances will be generated only if the function is "complicated". See the section on "More complicated user functions" for further details.

Normally, a user function is identified as multi-valued from the context in which it is used, and so the `M` suffix is not required. It must be used if the user function being defined is not used directly in any formulae.

Any formulae involving a non-differentiable function will always be evaluated using numerical derivatives.

The possible types for numerical derivatives are:

- 1 Forward derivatives
- 2 Tangential derivatives (calculated from forward and backward perturbation)

The suffix for numerical derivatives is not used if the function is defined as calculating its own derivatives. If no suffix is provided, then the method of calculating derivatives will be determined from the setting of `XSLP_FUNCEVAL`. The possible types for the calling mechanism are:

- | | |
|---|---|
| S | STDCALL (the default under Windows) |
| C | CDECL (the alternative mechanism under Windows) |

The setting of the calling mechanism has no effect on platforms other than Windows.

Param1	See Linkage
Param2	See Linkage
Param3	Name of return array for <code>MOSEL</code> linkage

Notes:

1. If an argument is declared as `NULL`, then Xpress-SLP will provide a dummy argument of the correct type, but it will contain no useful information.
2. Arguments can be omitted entirely. This is achieved by leaving the space for the declaration of the argument empty (for example, by having two consecutive commas). In this case, Xpress-SLP will omit the argument altogether. Trailing empty declarations can be omitted (that is, the closing bracket can immediately follow the last required argument).
3. VB, COM, XLS and XLF require `VARIANT` types for their arguments. A declaration of any other type will be treated as `VARIANT` for these linkage types. `VARIANT` cannot be used for other linkage types.

4. Functions which do not perform their own differentiation must declare `Deltas` as `NULL` or omit it altogether.
5. The `Extname`, `Param1`, `Param2` and `Param3` fields can contain the names of character variables (defined on `CV` records). This form is required if the data to go in the field contains spaces. If the data does not contain spaces, the data can be provided directly in the field.

If a function has a constant derivative with respect to any of its variables, Xpress-SLP can save some time by not repeatedly evaluating the function to obtain the same result. Provided that there are no circumstances in which the function might return values which imply derivatives identical to within about $1.0E-08$ over a range of ± 0.0001 or so for a derivative which is *not* constant, then the suffix `W` can be used so that Xpress-SLP will assume that where a derivative appears to be constant within tolerances `XSLP_CDTOL_A` or `XSLP_CDTOL_R` it is actually constant and does not need further re-evaluation. If there are some derivatives which might falsely appear to be constant, then it is better to use the suffix `V` and write the function so that it can be interrogated for constant derivatives.

See for a detailed explanation of constant derivatives.

Examples:

UF MyLog (DOUBLE) DLL = MyFuncs

This declares a simple function called `MyLog` which only needs the input arguments. Because `FunctionInfo` is omitted, the number of arguments is probably fixed, or can be determined from the input argument list itself. The function is compiled as a user function in the library file `MyFuncs` (depending on the platform, the file may have an extension).

UF MyCalc = Simulator (VARIANT , VARIANT) XLS = MyTests.xls = XSLPInOut

This declares a function called `MyCalc` in Xpress-SLP formulae. It is implemented as an Excel macro called `Simulator` in the workbook `MyTests.xls`. Xpress-SLP will place the input data in sheet `XSLPInOut` in columns `A` and `B`; this is because only the first two arguments are declared to be in use. Xpress-SLP will expect the results in column `I` of the same sheet. Note that although the arguments are respectively of type `DOYUBLE` and `INTEGER`, they are both declared as `VARIANT` because the linkage mechanism uses only `VARIANT` types.

UF MyFunc = AdvancedFunction (VARIANT , VARIANT , VARIANT , VARIANT) XLF = MyTests.xls = XSLPFunc

This declares a function called `MyFunc` in Xpress-SLP formulae. It is implemented as an Excel function on sheet `XSLPFunc` in the Excel workbook `MyTests.xls`. It will take values from and return values directly to Xpress-SLP without using a sheet as an intermediary.

UF MyFunc = CFunc (DOUBLE , INTEGER , CHAR , , DOUBLE , DOUBLE) DLL = MyLib

This declares a function called `MyFunc` in Xpress-SLP formulae. It is implemented as the function `CFunc` compiled in the user library `MyLib`. It takes a list of input names as the third argument, so it can identify arguments by name instead of by position. The fourth argument in the declaration is empty, meaning that the `ReturnNames` argument is not used. The fourth argument to the function is therefore the `Deltas` array of perturbations. Because `Deltas` is specified, the function must produce its own array of derivatives if required. It returns the array of results into the array defined by its fifth argument. The function itself will return a single status value.

13.4.2 Function declaration through `XSLPloaduserfuncs` and `XSLPadduserfuncs`

The method for declaring a user function is the same for `XSLPloaduserfuncs` and `XSLPadduserfuncs`. In each case the user function declaration is made using a variant of the parsed formula structure. Given the `UF` record described in the previous section:

`UF Function = Extname ( InputValues , FunctionInfo ,`

```
InputNames , ReturnNames , Deltas , ReturnArray )
Linkage = Param1 = Param2 = Param3
```

the equivalent formula sequence is:

Type	Value
XSLP_STRING	index of <code>Extname</code> in string table
XSLP_UFARGTYPE	bit map representing the number and type of the arguments (see below)
XSLP_UFEXETYPE	bitmap representing the linkage type, calling mechanism, derivative and evaluation options (see below)
XSLP_STRING	index of <code>Param1</code> in string table
XSLP_STRING	index of <code>Param2</code> in string table
XSLP_STRING	index of <code>Param3</code> in string table
XSLP_EOF	0

Notes:

1. The value of the `XSLP_UFARGTYPE` token holds the information for the existence and type of each of the 6 possible arguments. Bits 0-2 represent the first argument (`InputValues`), bits 3-5 represent the second argument (`FunctionInfo`) and so on. Each 3-bit field takes one of the following values, describing the existence and type of the argument:

0	argument is omitted
1	NULL (argument is present but has no information content)
2	INTEGER
3	DOUBLE
4	VARIANT
6	CHAR

2. The value of the `XSLP_UFEXETYPE` token holds the linkage type, the calling mechanism, and the options for evaluation and for calculating derivatives:

Bits 0-2 type of linkage:

1	DLL (User library or DLL)
2	XLS (Excel spreadsheet)
3	XLf (Excel macro)
5	MOSEL
6	VB
7	COM

Bits 3-4 evaluation flags:

0	default
1 (Bit 3)	re-evaluation at each SLP iteration
2 (Bit 4)	re-evaluation when independent variables have changed outside tolerance

Bits 6-7 derivative flags:

0	default
1 (Bit 6)	tangential derivatives
2 (Bit 7)	forward derivatives

Bit 8 calling mechanism:

0	standard
---	----------

1 CDECL (Windows only)
 Bit 24 set if the function is multi-valued
 Bit 28 set if the function is not differentiable

Bits 11-12 constant derivative flags:

0 default: no known constant derivatives
 1 (Bit 11) assume that derivatives which do not change outside the tolerance are constant
 2 (Bit 12) interrogate function for constant derivatives

The following manifest constants are provided for setting these bits:

Setting bit 11 XSLP_DEDUCECONSTDERIVS

Setting bit 12 XSLP_SOMECONSTDERIVS

See *Constant Derivatives* for a detailed explanation of constant derivatives.

3. The string arguments are interpreted in the order in which they appear. Therefore, if any of the function parameters `Param1` to `Param3` is required, there must be entries for the internal function name and any preceding function parameters. If the fields are blank, use an `XSLP_STRING` token with a zero value.
4. The name of the function itself (`Function` in this case) is provided through the function `XSLPaddnames`.

13.4.3 Function declaration through XSLPchguserfunc

Functions can be declared individually using `XSLPchguserfunc`. The function information is passed in separate variables, rather than in an array of tokens. Given the `UF` record described earlier in Extended MPS format:

```
UF Function = Extname ( InputValues , FunctionInfo ,
  InputNames , ReturnNames , Deltas , ReturnArray )
Linkage = Param1 = Param2 = Param3
```

the equivalent declaration is:

```
XSLPchguserfunc(Prob, 0, Extname, &ArgType, &ExeType,
  Param1, Param2, Param3)
```

where: `Extname`, `Param1`, `Param2` and `Param3` are character strings; `ArgType` and `ExeType` are integers.

An unused character string can be represented by an empty string or a `NULL` argument.

`ArgType` and `ExeType` are bitmaps with the same meaning as in the previous section.

Using zero as the second argument to `XSLPchguserfunc` forces the creation of a new user function definition. A positive integer will *change* the definition of an existing user function. In that case, a `NULL` argument means "no change".

13.4.4 Function declaration through SLPDATA in Mosel

In Mosel, a user function is declared to Xpress-SLP using the `SLPDATA` function which mirrors the Extended MPS format declaration for file-based definitions.

```
SLPDATA(UF:string, Function:string, Extname:string, ArgList:string,
  ArgType:string [,Param1:string [,Param2:string [,Param3:string] ] ] )
```

Arguments:

UF	string containing UF, indicating the SLPDATA type.
Function	name of the function (as used within a <code>Func()</code> expression)
Extname	name of the function to be used when it is called. This may be different from <code>Function</code> (for example, it may be decorated or have a special prefix).
ArgList	list of the argument types to the function, as described in Extended MPS format . Effectively, it is the same as the list of argument types within the brackets in an Extended MPS format declaration: for example "DOUBLE, INTEGER". The argument types must match exactly the declaration of the function in its native language.
ArgType	the function type as described in Extended MPS format .
Param1-3	optional strings giving additional parameter information as required by the particular function type. Details are in Extended MPS format .

13.5 User Function declaration in native languages

This section describes how to declare a user function in C, Fortran and so on. 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-SLP.

13.5.1 User function declaration in C

The `XPRS_CC` calling convention (equivalent to `__stdcall` under Windows) must be used for the function. For example:

```
type XPRS_CC MyFunc(double *InputValues, int *FunctionInfo,
                   char *InputNames, char *ReturnNames
                   double *Deltas, double *ReturnArray);
```

where *type* is *double* or *double** depending on the nature of the function.

In C++, the function should be declared as having a standard C-style linkage. For example, with Microsoft C++ under Windows:

```
extern "C" type _declspec(dllexport) XPRS_CC
MyFunc(double *InputValues, int *FunctionInfo,
       char *InputNames, char *ReturnNames
       double *Deltas, double *ReturnArray);
```

If the function is placed in a library, the function name may need to be externalized. If the compiler adds "decoration" to the name of the function, the function may also need to be given an alias which is the original name. For example, with the Microsoft compiler, a definition file can be used, containing the following items:

```
EXPORTS
MyFunc=_MyFunc@12
```

where the name after the equals sign is the original function name preceded by an underscore and followed by the @ sign and the number of bytes in the arguments. As all arguments in

Xpress-SLP external function calls are pointers, each argument represents 4 bytes on a 32-bit platform, and 8 bytes on a 64-bit platform.

A user function can be included in the executable program which calls Xpress-SLP. In such a case, the user function is declared as usual, but the address of the program is provided using `XSLPchguserfuncaddress` or `XSLPsetuserfuncaddress`. The same technique can also be used when the function has been loaded by the main program and, again, its address is already known.

The `InputNames` and `ReturnNames` arrays, if used, contain a sequence of character strings which are the names, each terminated by a null character.

Any argument omitted from the declaration in Xpress-SLP will be omitted from the function call.

Any argument declared in Xpress-SLP as of type `NULL` will generally be passed as a null pointer to the program.

13.5.2 User function declaration in Fortran

```
FUNCTION MyFunc(InputValues, FunctionInfo, InputNames,  
1 ReturnNames, Deltas, ReturnArray)  
  INTEGER*4 FunctionInfo(XSLP_FUNCINFOSIZE)  
  REAL*8 InputValues(1), Deltas(1), ReturnArray(1)  
  CHARACTER*4 InputNames, ReturnNames  
  REAL*8 MyFunc
```

Fortran functions receive all their arguments by reference but return a single value. A multi-valued Fortran function must use the `ReturnArray` argument (see "General Function Form" below) and must then return zero for success or 1 for failure.

Depending on the compiler, it may be necessary to describe an alias for the function and/or define a specific calling linkage from Xpress-SLP (for example, use the `CDECL` calling mechanism).

A user function can be included in the executable program which calls Xpress-SLP. In such a case, the user function is declared as usual, but the address of the program is provided using `XSLPchguserfuncaddress` or `XSLPsetuserfuncaddress`. The same technique can also be used when the function has been loaded by the main program and, again, its address is already known.

The `InputNames` and `ReturnNames` arrays, if used, contain a sequence of character strings which are the names, each terminated by a null character.

Any argument omitted from the declaration in Xpress-SLP will be omitted from the function call.

Any argument declared in Xpress-SLP as of type `NULL` will generally be passed as a null pointer to the program and cannot be used.

13.5.3 User function declaration in Excel (spreadsheet)

A user function written in formulae in a spreadsheet does not have a declaration as such. Instead, the values of the arguments supplied are placed in the sheet named in the Xpress-SLP declaration as follows:

Column A	<code>InputValues</code>
Column B	<code>FunctionInfo</code>
Column C	<code>InputNames</code>
Column D	<code>ReturnNames</code>
Column E	<code>Deltas</code>

The results are returned in the same sheet as follows:

Column I	Return values
Column J	Derivatives w.r.t. first required variable
Column K	Derivatives w.r.t. second required variable
...	

An Excel macro can also be executed as part of the calculation. If one is required, its name is given as `Extname` in the Xpress-SLP declaration of the user function.

Any argument omitted from the declaration in Xpress-SLP will be omitted from the function call.

Any argument declared in Xpress-SLP as of type `NULL` or omitted from the declaration will leave an empty column.

13.5.4 User function declaration in VBA (Excel macro)

All arguments to VBA functions are passed as arrays of type `Variant`. This includes integer or double precision arrays, which are handled as `Variant` arrays of integers or doubles. The following style of function declaration should be used:

```
Function MyFunc ( InputValues() as Variant, FunctionInfo() as Variant, _
    InputNames() as Variant, ReturnNames as Variant, _
    Deltas as Variant(), ReturnArray as Variant()) as Variant
```

For compatibility with earlier versions of Xpress-SLP, a return type of `Double` (or `Double()` for a multi-valued function) is also accepted. The return should be set to the value or to the array of values. For example:

```
Dim myDouble as Double
...
MyFunc = myDouble
```

or

```
Dim myDouble(10) as Double
...
MyFunc = myDouble
```

The return type is always `Variant`, regardless of whether the function returns one value or an array of values. The return should be set to the value or to the array of values as described in the VBA (Excel) section above.

All arrays are indexed from zero.

Any argument omitted from the declaration in Xpress-SLP will be omitted from the function call.

Any argument declared in Xpress-SLP as of type `NULL` will generally be passed as an empty `Variant`.

13.5.5 User function declaration in Visual Basic

All arguments to VB functions are passed as arrays of type `Variant`. This includes integer or double precision arrays, which are handled as `Variant` arrays of integers or doubles. The following style of function declaration should be used:

```
Public Function MyFunc ( InputValues() as Variant, FunctionInfo() as Variant,
    InputNames() as Variant, ReturnNames as Variant, _
    Deltas as Variant(), ReturnArray as Variant()) as Variant
```

The return type is always `Variant`, regardless of whether the function returns one value or an

array of values. The return should be set to the value or to the array of values as described in the VBA (Excel) section above.

All arrays are indexed from zero.

The address of a VB function is always defined to Xpress-SLP by using `XSLPchguserfuncaddress`.

Any argument omitted from the declaration in Xpress-SLP will be omitted from the function call.

Any argument declared in Xpress-SLP as of type `NULL` will generally be passed as an empty `Variant`.

13.5.6 User function declaration in COM

This example uses Visual Basic. All arguments to COM functions are passed as arrays of type `Variant`. This includes integer or double precision arrays, which are handled as `Variant` arrays of integers or doubles. The function must be stored in a class module, whose name will be needed to make up the `PROGID` for the function. The `PROGID` is typically of the form `file.class` where `file` is the name of the ActiveX DLL which has been created, and `class` is the name of the class module in which the function has been stored. If you are not sure of the name, check the registry. The following style of function declaration should be used:

```
Public Function MyFunc ( InputValues() as Variant, FunctionInfo() as Variant,
                        InputNames() as Variant, ReturnNames as Variant, _
                        Deltas as Variant(), ReturnArray as Variant()) as Variant
```

The return type is always `Variant`, regardless of whether the function returns one value or an array of values. The return should be set to the value or to the array of values as described in the VBA (Excel) section above.

All arrays are indexed from zero.

Any argument omitted from the declaration in Xpress-SLP will be omitted from the function call.

Any argument declared in Xpress-SLP as of type `NULL` will generally be passed as an empty `Variant`.

13.5.7 User function declaration in MOSEL

A simple function taking one or more input values and returning a single result can be declared in Mosel using the following form:

```
function MyFunc (InputValues:array(aRange:range) of real, Num:integer) : real
```

where `Num` will hold the number of values in the array `InputValues`.

The single result is placed in the reserved `returned variable`.

If the function returns more than one value, or calculates derivatives, then the full form of the function is used:

```
function MyFunc (InputValues:array(vRange:range) of real,
                        FunctionInfo:(array(fRange:range) of integer,
                        InputNames:(array(iRange:range) of string,
                        ReturnNames:(array(rRange:range) of string,
                        Deltas:(array(dRange:range) of real,
                        ReturnArray:(array(aRange:range) of real) : real
```

The `SLPDATA` declaration of the function references an array (the *transfer array*) which is a string array containing the names of the arrays used as arguments to the function.

The results are placed in `ReturnArray` and the function should return zero for success or 1 for failure.

For more details about user functions in Mosel, see the Xpress Mosel SLP Reference Manual.

13.6 Simple functions and general functions

A *simple function* is one which returns a single value calculated from its arguments, and does not provide derivatives. A *general function* returns more than one value, because it calculates an array of results, or because it calculates derivatives, or both.

Because of restrictions in the various types of linkage, not all types of function can be declared and used in all languages. Any limitations are described in the appropriate sections.

For simplicity, the functions will be described using only examples in C. Implementation in other languages follows the same rules.

13.6.1 Simple user functions

A simple user function returns only one value and does not calculate derivatives. It therefore does not use the `ReturnNames`, `Deltas` or `ReturnArray` arguments.

The full form of the declaration is:

```
double XPRS_CC MyFunc(double *InputValues, int *FunctionInfo,
                      char *InputNames);
```

`FunctionInfo` can be omitted if the number of arguments is not required, and access to problem information and function objects is not required.

`InputNames` can be omitted if the input values are identified by position and not by name (see "Programming Techniques for User Functions" below).

The function supplies its single result as the return value of the function.

There is no provision for indicating that an error has occurred, so the function must always be able to calculate a value.

13.6.2 General user functions returning an array of values through a reference

General user functions calculate more than one value, and the results are returned as an array. In the first form of a general function, the values are supplied by returning the address of an array which holds the values. See the notes below for restrictions on the use of this method.

The full form of the declaration is:

```
double * XPRS_CC MyFunc(double *InputValues, int *FunctionInfo,
                        char *InputNames, char *ReturnNames,
                        double *Deltas);
```

`FunctionInfo` can be omitted if the number of arguments is not required, no derivatives are being calculated, the number of return values is fixed, and access to problem information and function objects is not required. However, it is recommended that `FunctionInfo` is always included.

`InputNames` can be omitted if the input values are identified by position and not by name (see "Programming Techniques for User Functions" below).

`ReturnNames` can be omitted if the return values are identified by position and not by name (see

"Programming Techniques for User Functions" below).

`Deltas` must be omitted if no derivatives are calculated.

The function supplies the address of an array of results. This array must be available after the function has returned to its caller, and so is normally a static array. This may mean that the function cannot be called from a multi-threaded optimization, or where multiple instances of the function are required, because the single copy of the array may be overwritten by another call to the function. An alternative method is to use a *function object* which refers to an array specific to the thread or problem being optimized.

`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. The total number of nonzero entries in `Deltas` is given in `FunctionInfo[2]`. In particular, if it is zero, then no derivatives are required at all.

When no derivatives are calculated, the array of return values simply contains the results (in the order specified by `ReturnNames` if used).

When derivatives are calculated, the array contains the values and the derivatives as follows (`DVi` 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.

This form must be used by user functions which are called through OLE automation (VBA (Excel) and COM) because they cannot directly access the memory areas of the main program.

This form cannot be used by Fortran programs because Fortran functions can only return a single value, not an array.

This form cannot be used by Mosel programs because Mosel functions can only return a single value, not an array.

13.6.3 General user functions returning an array of values through an argument

General user functions calculate more than one value, and the results are returned as an array. In the second form of a general function, the values are supplied by returning the values in an array provided as an argument to the function by the calling program. See the notes below for restrictions on the use of this method.

The full form of the declaration is:

```
double XPRS_CC MyFunc(double *InputValues, int *FunctionInfo,
                     char *InputNames, char *ReturnNames
                     double *Deltas, double *ReturnArray);
```

`FunctionInfo` can be omitted if the number of arguments is not required, no derivatives are being calculated, the number of return values is fixed, and access to problem information and function objects is not required. However, it is recommended that `FunctionInfo` is always included.

`InputNames` can be omitted if the input values are identified by position and not by name (see "Programming Techniques for User Functions" below).

`ReturnNames` can be omitted if the return values are identified by position and not by name (see "Programming Techniques for User Functions" below).

`Deltas` must be omitted if no derivatives are calculated.

The function must supply the results in the array `ReturnArray`. This array is guaranteed to be large enough to hold all the values requested by the calling program. No guarantee is given that the results will be retained between function calls.

`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. The total number of nonzero entries in `Deltas` is given in `FunctionInfo[2]`. In particular, if it is zero, then no derivatives are required at all.

When no derivatives are calculated, the array of return values simply contains the results (in the order specified by `ReturnNames` if used).

When derivatives are calculated, the array contains the values and the derivatives as follows (`DVi` 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 function 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.

This form must be not used by user functions which are called through OLE automation (VBA (Excel) and COM) because they cannot directly access the memory areas (in particular `ReturnArray`) in the main program.

This form must be used by Fortran programs because Fortran functions can only return a single value, not an array. An array of values must therefore be returned through `ReturnArray`.

This form must be used by Mosel programs because Mosel functions can only return a single value, not an array. An array of values must therefore be returned through `ReturnArray`.

13.7 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.

13.7.1 FunctionInfo

The array `FunctionInfo` is primarily used to provide the sizes of the arrays used as arguments to the functions, and to indicate how many derivatives are required.

In particular:

`FunctionInfo[0]` holds the number of input values supplied

`FunctionInfo[1]` holds the number of return values required

`FunctionInfo[2]` holds the number of sets of derivatives to be calculated.

In addition, it contains problem-specific information which allows the user function to access problem data such as control parameters and attributes, matrix elements and solution values. It also holds information about function objects and function instances.

13.7.2 InputNames

The function may have the potential to take a very large number of input values but in practice, within a particular problem, not all of them are used. For example, a function representing the model of a distillation unit may have input values relating to external air temperature and pressure which are not known or which cannot be controlled by the optimization. In general, therefore, these will take default values except for very specialized studies.

Although it would be possible to require that every function call had every input value specified, it would be wasteful in processing time to do so. In such cases, it is worth considering using named input variables, so that only those which are not at default values are included. The user function then picks up the input values by name, and assigns default values to the remainder.

`InputNames` is an array of character strings which contains the names of the input variables. The order of the input values is then determined by the order in `InputNames`. This may be different for each instance of the function (that is, for each different formula in which it appears) and so it is necessary for the function to check the order of the input values. If *function instances* are used, then it may be necessary to check only when the function instance is called for the first time, provided that the order can be stored for future calls to the same instance.

Unless the user function is being called directly from a program, `InputNames` can only be used with input values defined in XVs, so that names can be assigned to the values.

13.7.3 ReturnNames

The function may have the potential to calculate a very large number of results but in practice, within a particular problem, not all of them are used. For example, a detailed model of a process unit might calculate yields and qualities of streams, but also internal flow rates and catalyst usage which are not required for a basic planning problem (although they are very important for

detailed engineering investigations).

Although it would be possible to calculate every value and pass it back to the calling function every time, it could be wasteful in processing time to do so. In such cases, it is worth considering using named return values, so that only those which are actually required are included. The user function then identifies which values are required and only passes those values to its caller (possibly, therefore, omitting some of the calculations in the process).

`ReturnNames` is an array of character strings which contains the names of the return variables. The order of the values is then determined by the order in `ReturnNames`. This order may be different for different instances of the function (that is, for different formulae in which it is used). If the function does use named return values, it must check the order. If *function instances* are used for the function, then it may be necessary to check the order only when the function instance is called for the first time, if the order can be stored for subsequent use.

If the user function is being called by Xpress-SLP to calculate values during matrix generation or optimization, the list of return values required is created dynamically and the names will appear in the order in which they are first encountered. It is possible, therefore, that changes in the structure of a problem may change the order in which the names appear.

13.7.4 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.

13.7.5 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 `FunctionInfo` array holds details of the number of input values supplied, the number of return values required (`nRet`) and the number of sets derivatives required (`nDeriv`). The total number of values (and hence the minimum size of the array) is $nRet * (nDeriv + 1)$. Xpress-SLP 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-SLP 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 value of 1 or greater is an error code which will cause any formula evaluation to stop and will normally interrupt any optimization or other procedure. A value of -1 asks Xpress-SLP to estimate the function values from the last calculation of the values and partial derivatives. This will produce an error if there is no such set of values.

13.7.6 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-SLP 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.

13.7.7 Function Instances

Xpress-SLP 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:

$$f(x) + f(y) + g(x, y : 1)$$

$$f(y) * f(x) * g(x, y : 2)$$

$$f(z)$$

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-SLP 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. A function of this type can still be made to generate instances by defining its `ExeType` as creating instances (set bit 9 when using the normal library functions, or use the "I" suffix when using file-based input through `XSLPreadprob` or when using `SLPDATA` in Mosel).

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.

There are reasons for not wanting to use function instances:

- When the function is fast. It may be as fast to recalculate the value as to work out if evaluation is required.
- When the function is discontinuous. Small changes are estimated by using derivatives. These behave badly across a discontinuity and so it is usually better to evaluate the derivative of a formula by using the whole formula, rather than to calculate it from estimates of the derivatives of each term.

- Function instances do use more memory. Each instance holds a full copy of the last input and output values, and a full set of first-order derivatives. However, the only time when function instances are optional is when there is only one return value, so the extra space is not normally significant.

13.7.8 Function Objects

Normally, a user function is effectively a free-standing program: that is, it requires only its argument list in order to calculate its result(s). However, there may be circumstances where a user function requires access to additional data, as in the following examples:

1. The function is actually a simulator which needs access to specific (named) external files. In this case, the function needs to access a list of file names (or file handles if the files have been opened externally).
2. The function uses named input or output values and, having established the order once, needs to save the order for future calls. In this case, the function needs to use an array which is external to the function, so that it is not destroyed when the function exits.
3. The function returns an array of results and so the array must remain accessible after the function has returned. In this case, the function needs to use an array which is external to the function, so that it is not destroyed when the function exits.
4. The function determines whether it needs to re-evaluate its results when the values of the arguments have not changed significantly, and so it needs to keep a copy of the previous input and output values. In this case, the function needs to use an array which is external to the function, so that it is not destroyed when the function exits.
5. The function has to perform an initialization the first time it is called. In this case, the function needs to keep a reference to indicate whether it has been called before. It may be that a single initialization is required for the function, or it may be that it has to be initialized separately for each instance.

There is a potential difference between examples (3) and (4) above. In example (3), the array is needed only because Xpress-SLP will pick up the values when the function has returned and so the array still needs to exist. However, once the values have been obtained, the array is no longer required, and so the next call to the same function can use the same array. In example (4), the argument values are really required for each instance of the function: for example, if $f(x)$ and $f(y)$ are both used in formulae, where $f()$ is a user function and x and y are distinct variables, then it only makes sense to compare input argument values for $f(x)$ (that is, the value of x) against the previous value for x ; it does not make sense to compare against the previous value for y . In this case, a separate array is needed for each function instance.

Xpress-SLP provides three levels of user function object. These are:

- The *Global Function Object*. There is only one of these for each problem, which is accessible to all user functions.
- The *User Function Object*. There is one of these for each defined user function.
- The *Instance Function Object*. There is one of these for each instance of a function.

The library functions `XSLPsetuserfuncobject`, `XSLPchguserfuncobject` and `XSLPgetuserfuncobject` can be used to set, change and retrieve the values from a program or function which has access to the Xpress-SLP problem pointer.

The library functions `XSLPsetfuncobject`, `XSLPchgfuncobject` and `XSLPgetfuncobject` can be used by a user function to set, change or retrieve the *Global Function Object*, the *User Function Object* for the function, and the *Instance Function Object* for the instance of the function.

`XSLPgetfuncobject` can also be used to obtain the Xpress-SLP and Xpress Optimizer problem pointers. These can then be used to obtain any problem data, or to execute any allowable library function from within the user function.

Example:

A function which uses input or return names is regarded as a complicated function, and will therefore generate function instances. All the calls for a particular instance have the same set of inputs in the same order. It is therefore necessary to work out the order of the names only once, as long as the information can be retained for subsequent use. Because each instance may have a different order, as well as different variables, for its inputs, the information should be retained separately for each instance.

The following example shows the use of the *Instance Function Object* to retain the order of input values

```
NOTE
1  typedef struct tagMyStruct {
    int InputFromArg[5];
  } MyStruct;

static char *MyNames[] = {"SUL", "RVP", "ARO", "OLE", "BEN"};
static double Defaults[] = {0, 8, 4, 1, 0.5};

double XPRS_CC MyUserFunc(double *InputValues, int *FunctionInfo,
                          char *InputNames) {
    MyStruct *InstanceObject;
    void *Object;
    char *NextName;

    int i, iArg, nArg;
    double Inputs[5], Results[10];

2  XSLPgetfuncobject(FunctionInfo,XSLP_INSTANCEFUNCOBJECT,&Object);
3  if (Object == NULL) {
    Object = calloc(1,sizeof(MyStruct));
4  XSLPsetfuncobject(FunctionInfo,XSLP_INSTANCEFUNCOBJECT,Object);
    InstanceObject = (MyStruct *) Object;
    NextName = InputNames;
    nArg = FunctionInfo[0];
5  for (iArg = 1;iArg<=nArg;iArg++) {
    for (i=0;i<5;i++) {
        if (strcmp(NextName,MyNames[i])) continue;
        InstanceObject->InputFromArg[i] = iArg;
        break;
    }
    NextName = &NextName[strlen(NextName)+1];
    }
    InstanceObject = (MyStruct *) Object;
6  if (InstanceObject == NULL) {
7      XSLPgetfuncobject(FunctionInfo,XSLP_XSLPPROBLEM,&Object);
8      XSLPsetfunctionerror(Prob);
      return(1);
    }
9  for (i=0;i<5;i++) {
    iArg=InstanceObject->InputFromArg[i];
    if (iArg) Inputs[i] = InputValues[iArg-1];
    else Inputs[i] = Defaults[i];
    }
    MyCalc(Inputs, Results);
    .....
}
```

Notes:

1. A structure for the instance function object is defined. This is a convenient way of starting, because it is easy to expand it if more information (such as results) needs to be retained.
2. `XSLPgetfuncobject` recovers the instance function object reference from the `FunctionInfo` data.
3. On the first call to the function, the object is `NULL`.
4. After the object has been created, its address is stored as the instance function object.
5. The names in `InputNames` are in a continuous sequence, each separated from the next by a null character. This section tests each name against the ordered list of internal names. When there is a match, the correspondence is stored in the `InputFromArg` array. A more sophisticated version might fault erroneous or duplicate input names.
6. If `InstanceObject` is `NULL` then the initialization must have failed in some way. Depending on the circumstances, the user function may be able to proceed, or it may have to terminate in error. We will assume that it has to terminate.
7. `XSLPgetfuncobject` recovers the Xpress-SLP problem.
8. `XSLPsetfunctionerror` sets the error flag for the problem which will stop the optimization.
9. If the initialization was successful, the correspondence in `InputFromArg` is now available on each call to the function, because on subsequent calls, `Object` is not `NULL` and contains the address of the object for this particular instance.

If there are different instances for this function, or if several problems are in use simultaneously, each distinct call to the function will have its own object.

A similar method can be used to set up and retain a correspondence between the calculated results and those requested by the calling program.

The *User Function Object* can be used in a similar way, but there is only one such object for each function (not for each instance), so it is only appropriate for saving information which does not have to be kept separate at an instance level. One particular use for the *User Function Object* is to provide a return array which is not destroyed after the user function returns (an alternative is to use the `ReturnArray` argument to the function).

Note that one or more arrays may be allocated dynamically by each function using this type of approach. It may be necessary to release the memory if the problem is destroyed before the main program terminates. There is no built-in mechanism for this, because Xpress-SLP cannot know how the objects are structured. However, there is a specific callback (`XSLPsetcbdestroy`) which is called when a problem is about to be destroyed. As a simple example, if each non-null object is the address of an allocated array, and there are no other arrays that need to be freed, the following code fragment will free the memory:

```
int i, n;
void *Object;
XSLPgetintattrib(Prob, XSLP_UFINSTANCES, &n);
for (i=1;i<=n;i++) {
    XSLPgetuserfuncobject(Prob, -i, &Object);
    if (Object) free(Object);
    XSLPsetuserfuncobject(Prob, -i, NULL);
}
```

When used in the "destroy" callback, it is not necessary to set the instance function object to `NULL`. However, if an object is being freed at some other time, then it should be reset to `NULL` so that any subsequent call that requires it will not try to use an unallocated area of memory.

13.7.9 Calling user functions

A user function written in a particular language (such as C) can be called directly from another function written in the same language, using the normal calling mechanism. All that is required is for the calling routine to provide the arguments in the form expected by the user function.

Xpress-SLP provides a set of functions for calling between different languages so that, for example, it is possible for a program written in Mosel to call a user function written in C. Not all combinations of language are possible. The following table shows which are available:

User function	Calling program			
	Mosel	C/Fortran	VB	VBA (Excel)
Mosel	1	3	3	3
C/Fortran	1	1	1	1
VB	X	X	1	X
VBA (Excel macro)	2	2	2	2
Excel spreadsheet	2	2	2	2
COM	2	2	2	2

1: User functions available with full functionality

2: User functions available, but with reduced functionality

3: User functions available if Mosel model is executed from main program

X: User functions not available.

In general, those user functions which are called using OLE automation (Excel macro, Excel spreadsheet and COM) do not have the full functionality of user functions as described below, because the calling mechanism works with a copy of the data from the calling program rather than the original. Mosel user functions can only be called from problems which are created in the same Mosel model; however, because Mosel can itself be called from another program, Mosel functions may still be accessible to programs written in other languages.

`XSLPcalluserfunc` provides the mechanism for calling user functions. The user function is declared to Xpress-SLP as described earlier, so that its location, linkage and arguments are defined. In this section, we shall use three example user functions, defined in Extended MPS format as follows:

```
UF MyRealFunc ( DOUBLE , INTEGER ) .....
UF MyArrayFunc ( DOUBLE , INTEGER ) DLLM .....
UF MyRetArrayFunc ( DOUBLE , INTEGER , , , DOUBLE ) .....
```

These all take as arguments an array of input values and the `FunctionInfo` array. `MyArrayFunc` is declared as multi-valued (using the suffix `M` on the linkage). `MyRetArrayFunc` returns its results in `ReturnArray`; thus usually means that it is multi-valued, or calculates its own derivatives.

```
double Values[100];
double ReturnArray[200];
integer FunctionInfo[XSLP_FUNCINFOSIZE];
integer RealFunc, ArrayFunc, RetArrayFunc;
double ReturnValue;
```

The calling program has to provide its own arrays for the function calls, which must be sufficient to hold the largest amount of data required for any call. In particular, `ReturnArray` may need to allow space for derivatives.

`FunctionInfo` should always be declared as shown.

```
XSLPgetIndex(Prob, XSLP_USERFUNCNAMES, "MyRealFunc", RealFunc);
```

```
XSLPgetIndex(Prob, XSLP_USERFUNCNAMES, "MyArrayFunc", ArrayFunc);
XSLPgetIndex(Prob, XSLP_USERFUNCNAMES, "MyRetArrayFunc", RetArrayFunc);
```

As `XSLPcalluserfunc` needs the function number, we get this for each function by using the function `XSLPgetIndex`. If you are not sure of the upper- or lower-case, then use `XSLP_USERFUNCNAMESNOCASE` instead. If the functions are set up using library functions, the function indices can be obtained at that time.

```

...
/*... set up Values array .....*/
...
XSLPsetuserfuncinfo(Prob,ArgInfo,1,n,1,0,0,0);
```

The input data for the function call is set up. The contents of the input array `Values` obviously depend on the nature of the function being called, so we do not include them here. The function information array `FunctionInfo` must be set up. `XSLPsetuserfuncinfo` will fill in the array with the items shown. The arguments after `FunctionInfo` are:

- `CallerFlag`. This is always zero when the function is called directly by Xpress-SLP, and so if set nonzero it indicates a call from the user application; its value can be used for any purpose in the calling and called functions.
- The number of input variables: this is the number of elements used in the input array `Values`.
- The number of return values required for each calculation.
- The number of sets of partial derivatives required.
- The number of items in the array of input argument names.
- The number of items in the array of return value names.

This structure actually allows more flexibility than is used when the function is called directly by Xpress-SLP because, for example, there is no requirement for the number of input names to be the same as the number of input arguments. However, such usage is beyond the scope of this manual.

```
ReturnValue = XSLPcalluserfunc(Prob,RealFunc,Values,FunctionInfo,
                               NULL,NULL,NULL,NULL);
```

`XSLPcalluserfunc` calls the function using the appropriate linkage and calling mechanism. The arguments to `XSLPcalluserfunc` are:

- The Xpress-SLP problem.
- The index of the function being called.
- Six arguments corresponding to the six possible arguments to a user function. If the user function requires an argument, then the corresponding argument in the call must contain the appropriate data in the correct format. If the user function does not require an argument, then it can be `NULL` in the call (in any case, it will be omitted from the call). The `FunctionInfo` argument is always required for function calls using `XSLPcalluserfunc`.

`ReturnValue` will contain the single value returned by the user function.

```
ReturnValue = XSLPcalluserfunc(Prob,ArrayFunc,Values,FunctionInfo,
                               NULL,NULL,NULL,NULL);
```

This time, `ReturnValue` will contain the first value in the array of results returned by the function. This is because the function is multi-valued and there is nowhere for the other values to go.

Multi-valued functions must be called using the `ReturnArray` argument. Even if the user function itself does not recognize it, `XSLPcalluserfunc` does, and will transfer the results into it.

```
ReturnValue = XSLPcalluserfunc(Prob,ArrayFunc,Values,FunctionInfo,  
                              NULL,NULL,NULL,ReturnArray);
```

The difference between this call and the previous one is the presence of the additional argument `ReturnArray`. This will be used to hold all the values returned by the function. The function will behave in exactly the same way as in the previous example, and `ReturnValue` will also be the same, but `ReturnArray` will be filled in with the values from the function.

```
ReturnValue = XSLPcalluserfunc(Prob,RetArrayFunc,Values,FunctionInfo,  
                              NULL,NULL,NULL,ReturnArray);
```

As `MyRetArrayFunc` is defined as returning its results in an array, the `ReturnArray` argument is a required argument for the function anyway. In this case, `ReturnValue` is the value returned by the function, which indicates success (zero), failure (1) or not calculated (-1).

13.8 Function Derivatives

Xpress-SLP 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_FUNCNEVAL` 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_FUNCNEVAL` 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 manifest constants are provided for setting these bits:

Setting bit 3	<code>XSLP_RECALC</code>
Setting bit 4	<code>XSLP_TOLCALC</code>
Setting bit 5	<code>XSLP_ALLCALCS</code>
Setting bit 6	<code>XSLP_2DERIVATIVE</code>
Setting bit 7	<code>XSLP_1DERIVATIVE</code>
Setting bit 8	<code>XSLP_ALLDERIVATIVES</code>

A function can make its own determination of whether to re-evaluate. If the function has already calculated and returned a full set of values and partial derivatives, then it can request Xpress-SLP to estimate the values required from those already provided.

The function must be defined as using the `ReturnArray` argument, so that the return value from the function itself is a double precision status value as follows:

- | | |
|----|---|
| 0 | normal return. The function has calculated the values and they are in <code>ReturnArray</code> . |
| 1 | error return. The function has encountered an unrecoverable error. The values in <code>ReturnArray</code> are ignored and the optimization will normally terminate. |
| -1 | no calculation. Xpress-SLP should recalculate the values from the previous results. The values in <code>ReturnArray</code> are ignored. |

Chapter 14

Management of zero placeholder entries

14.1 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. 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 – $(R1, \delta X)$, $(R1, \delta Y)$ and $(R1, =)$ – are known as *placeholders*.

14.2 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.

14.3 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.

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**.

Chapter 15

Special Types of Problem

15.1 Nonlinear objectives

Xpress-SLP 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.

15.2 Quadratic Programming

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). It is possible to solve quadratic problems using SLP, but it is not usually the best way. The reason is that the solution to a 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 is 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-SLP, 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. There is an overhead in using the QP optimizer because it cannot make use of the basis from the previous solution. However, this is usually more than

offset by the saving in SLP iterations. If the objective is not semi-definite, the QP optimizer may not give a solution; 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 `XPRSloadqglobal`). Xpress-SLP 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.

15.3 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-SLP offers a set of components to implement MINLP using Mixed Integer Successive Linear Programming (MISLP).

15.3.1 Approaches to MISLP

Essentially, there are three ways to approach MISLP:

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 within MIP. In this, MIP 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.
3. 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 (3) 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.

Xpress-SLP normally uses the approach outlined in (2). Other approaches can be used by changing the value of the control parameter `XSLP_MIPALGORITHM`. 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.

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

15.3.2 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.

15.3.3 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.

15.3.4 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

15.3.5 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.

Chapter 16

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-SLP error* and *Xpress-SLP 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 `XSLPuprintmemory` 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 XSLPuprintmemory 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-SLP 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, VB 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-SLP function func which is referencing the SLP item (row, column variable, XV, 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-SLP items count from 1, and are not sensitive to the setting of XPRS_CSTYLE.
- 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-12021** *row #num is defined more than once as an SLP delayed constraint*
This message is produced by `XSLPadddc` or `XSLPloaddc` if the same row appears more than once in the list, or has already been defined as a delayed constraint. Although `XSLPchgdc` is less efficient, it can be used to set the properties of an SLP delayed constraint 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-SLP. 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`.

W-12036 *the number of items in XV #num cannot be increased*

This message is produced by `XSLPchgxxv` if the number of `XVitems` specified is larger than the current number. `XSLPchgxxv` can only reduce the number of items; use `XSLPchgxxvitem` to add new items.

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-SLP 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-SLP 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-12039 *incorrect OLE version*

This message is produced if a user function is specified using an OLE linkage (Excel or COM) but the OLE version is not compatible with the version used by Xpress-SLP. If this error is encountered, please contact your local Xpress support office.

E-12040 *unable to initialize OLE - code num*

This message is produced if the OLE initialization failed. The initialization error code is printed in hexadecimal. Consult the appropriate OLE documentation to establish the cause of the error.

- E-12041** *unable to open Excel/COM - code num*
This message is printed if the initialization of Excel or COM failed after OLE was initialized successfully. The error code is printed in hexadecimal. Consult the appropriate documentation to establish the cause of the error.
- E-12042** *OLE/Excel/COM error: msg*
This message is produced if OLE automation produces an error during transfer of data to or from Excel or COM. The message text gives more information about the specific error.
- E-12084** *Xpress-SLP has not been initialized*
An attempt has been made to use Xpress-SLP 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-SLP has not been licensed for use here*
Either Xpress-SLP 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-SLP 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-SLP 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`, `VB` or `COM` and is run on a non-Windows platform.
- E-12111** *Xpress-SLP error: unidentified section in REVISE: name*
This message is produced by `XSLPrevise` if it encounters a section other than `NAME`, `ROWS`, `COLUMNS`, `RHS`, `RANGES`, `BOUNDS`, `ENDATA`, `MODIFY`, `BEFORE` or `AFTER`.
- E-12112** *Xpress-SLP error: unidentified row type in REVISE: name*
This message is produced by `XSLPrevise` if it encounters a row type other than `L`, `G`, `E` or `N`.
- E-12113** *Xpress-SLP error: unidentified row in REVISE: name*
This message is produced by `XSLPrevise` if it encounters a row name which is not already in the problem. `XSLPprevise` cannot add new rows to the problem.
- E-12114** *Xpress-SLP error: unidentified bound type in REVISE: name*
This message is produced by `XSLPprevise` if it encounters a bound type that is not `LO`, `UP`, `FX` or `FR`. It is not possible to change bounds for global variables using `XSLPprevise`.
- E-12115** *Xpress-SLP error: unidentified column in REVISE: name*
This message is produced by `XSLPprevise` if it encounters a column name which is not already in the problem. `XSLPprevise` cannot add new columns to the problem.
- E-12121** *Xpress-SLP 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-SLP error: augmented problem not set up*
The message is produced by `XSLPvalidate` if an attempt is made to validate the problem without a preceding call to `XSLPconstruct`. In fact, unless a solution to the linearized problem is available, `XSLPvalidate` will not be able to give useful results.

- E-12125** *Xpress-SLP 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 `XSLPsetfunctionerror`).
- W-12142** *Xpress-SLP warning: invalid record: `text`*
This error is produced by `XSLPreadprob` 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-SLP error: incompatible arguments in user function `func`*
This message is produced if a user function is called by `XSLPcalluserfunc` but the function call does not provide the arguments required by the function.
- E-12148** *Xpress-SLP error: user function `func` should return an array not a single value*
This message is produced if a user function is defined within Xpress-SLP as returning an array, but the function is returning a single value. This message is produced only when it is possible to identify the type of value being returned by the function (for example, the value from an Excel macro).
- E-12158** *Xpress-SLP 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 (`XSLPgetparam` or `XSLPsetparam` where the name is incorrect).
- E-12159** *Xpress-SLP 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-SLP 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

Chapter 17

Files used by Xpress-SLP

Most of the data used by Xpress-SLP 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.mat</i>	Created by: XSLPwriteprob This is the matrix file in extended MPS format. The name is user-defined. The extension <i>.mat</i> is appended automatically.
<i>NAME.txt</i>	Created by: XSLPwriteprob This is the matrix file in human-readable "text". The name is user-defined. The extension <i>.txt</i> is appended automatically.
<i>PROBNAME.sgb</i>	Created by: XSLPglobal This is the SLP part of the global save file (the linear part is in <i>probname.glb</i>). Used by XSLPglobal .
<i>PROBNAME.svx</i>	Created by: XSLPsave This is the SLP part of the save file (the linear part is in <i>probname.svf</i>). Used by XSLPrestore .
<i>slpXXX.mat</i>	Created by: XSLPreadprob When a matrix is read, it is split into linear and nonlinear parts, in two temporary files beginning with "slp" and with unique names. These are stored in a temporary directory, as defined by one of the environment variables <code>TMPDIR</code> , <code>TMP</code> , <code>TEMP</code> , <code>tmpdir</code> , <code>tmp</code> , or <code>temp</code> if any are defined. If more than one is defined, the first is used. If none is defined, the directory <code>\temp</code> or <code>\tmp</code> (under Windows) or <code>/tmp</code> or <code>/var/tmp</code> (under Unix) will be used if it exists. If none exist, the current directory will be used.

Index

Symbols

= column, 5

A

ABS, 327

Absolute tolerance record

Tx, 8

ACT, 344

ARCCOS, 320

ARCSIN, 321

ARCTAN, 322

Attributes, Problem, 34

Augmentation, 25

B

BOUNDS section in file, 5

C

Callbacks and user functions, 365

Callbacks in MISLP, 395

Calling user functions, 385

Cascading, 15

Character Variable record, 5

Closure convergence tolerance, 19

Coefficients

and terms, 1

COLUMNS section in file, 4

Control parameters, 52

Convergence

closure, 19

delta, 20

extended convergence continuation, 24

impact, 21

matrix, 20

slack impact, 21

static objective (1), 22

static objective (2), 22

static objective (3), 23

user-defined, 22

Convergence criteria, 17

COS, 323

Counting, 130

CV record in SLPDATA, 5

D

DC record in SLPDATA, 6

Delayed Constraint

add, 142

change, 162

Extended MPS record, 6

get, 196

load, 236

Delta convergence tolerance, 20

Derivatives

returning from user function, 380

user function, 387

Determining Row record, 6

DJ, 345

DR record in SLPDATA, 6

E

E-12001, 397

E-12002, 397

E-12003, 397

E-12004, 397

E-12005, 397

E-12006, 398

E-12007, 398

E-12008, 398

E-12009, 398

E-12010, 398

E-12011, 398

E-12012, 398

E-12013, 398

E-12014, 398

E-12015, 398

E-12016, 398

E-12017, 398

E-12018, 399

E-12019, 399

E-12020, 399

E-12021, 399

E-12022, 399

E-12025, 399

E-12026, 399

E-12027, 399

E-12029, 400

E-12030, 400

E-12031, 400

E-12032, 400

E-12033, 400

E-12037, 400

E-12038, 400

E-12039, 400

E-12040, 400

E-12041, 401

E-12042, 401

E-12084, 401

E-12085, 401

E-12105, 401

E-12107, 401

- E-12111, 401
- E-12112, 401
- E-12113, 401
- E-12114, 401
- E-12115, 401
- E-12121, 401
- E-12124, 401
- E-12125, 402
- E-12147, 402
- E-12148, 402
- E-12158, 402
- E-12159, 402
- E-12160, 402
- EC record in SLPDATA, 7
- Enforced Constraint record, 7
- EQ, 335
- Equals column, 5
- Error Messages, 397
- Error vectors, penalty, 29
- EXP, 328
- Extended convergence continuation tolerance, 24
- Extended MPS file format, 3

F

- Files used by Xpress-SLP, 403
- Fixing values of SLP variables in MISLP, 394
- Formula
 - Initial Value record, 7
- Formulae, 3, 357
- Function object, 382
- Functions, internal, 317
- Functions, library, 130
- Functions, user, 365

G

- GE, 336
- Global Function Object, 382
- GT, 337

H

- Handling Infeasibilities, 14
- History, 33

I

- IAC, 355
- IF, 338
- Impact convergence tolerance, 21
- Implicit variable, 4
- Infeasibilities, handling, 14
- Initial Value formula, 7
- Initial Value record, 7
- Instance
 - user function, 381
- Instance Function Object, 382
- Internal Functions, 317
- INTERP, 356
- Iterating at each node in MISLP, 394
- IV record in SLPDATA, 7

L

- LE, 339

- Library functions, 130, 132
- LN, 329
- LO, 346
- LOG, 330
- LOG10, 330
- LT, 340

M

- MATRIX, 347
- Matrix convergence tolerance, 20
- Matrix Name Generation, 30
- Matrix Structures, 25
- MAX, 331
- MIN, 332
- MINLP, 393
- MISLP
 - Callbacks, 395
 - Fixing or relaxing values of SLP variables, 394
 - Iterating at each node, 394
 - Termination criteria at each node, 395
- Mixed Integer Non-Linear Programming, 393
- MV, 348

N

- Name Generation, 30
- NE, 341
- Nonlinear objectives, 392
- Nonlinear problems, 1
- NOT, 342

O

- Object
 - user function, 382
- Objectives, nonlinear, 392
- Objectives, quadratic, 392

P

- PARAM, 349
- Parsed formula format, 357
- Penalty error vectors, 29
- Pointer (reference) attribute, 50
- Problem attributes, 34
- Problem pointer, 131

Q

- Quadratic objectives, 392

R

- Relative tolerance record Rx, 8
- Relaxing values of SLP variables in MISLP, 394
- RHS, 350
- RHSRANGE, 351
- Row weight
 - Extended MPS record, 9
- Rx record in SLPDATA, 8

S

- SB record in SLPDATA, 9
- Sequential Linear Programming, see Successive Linear Programming
- SIN, 324

SLACK, 352
 Slack impact convergence tolerance, 21
 SLP problem pointer, 131
 SLP variable, 2
 SLPDATA
 CV record, 5
 DC record, 6
 DR record, 6
 EC record, 7
 IV record, 7
 Rx record, 8
 SB record, 9
 Tx record, 8
 UF record, 9
 WT record, 9
 XV record, 10
 SLPDATA section in file, 5
 Solution Process, 12
 Special Types of Problem, see Problem, special types
 Mixed Integer Non-Linear Programming, 393
 Nonlinear objectives, 392
 Quadratic objectives, 392
 SQRT, 333
 Static objective (1) convergence tolerance, 22
 Static objective (2) convergence tolerance, 22
 Static objective (3) convergence tolerance, 23
 Statistics, Xpress-SLP, 31
 Step Bound record, 9
 Structures, SLP matrix, 25
 Successive Linear Programming, 1

T
 TAN, 325
 Termination criteria at each node in MISLP, 395
 Terms
 and coefficients, 1
 Tolerance record
 Rx, 8
 Tolerance record Tx, 8
 Tolerances, convergence, 17
 Tx record in SLPDATA, 8

U
 UF record in SLPDATA, 9
 Unparsed formula format, 357
 UP, 353
 User function, 365
 calling, 385
 declaration in COM, 375
 declaration in Excel macro (VBA), 374
 declaration in Excel spreadsheet, 373
 declaration in Extended MPS format, 366
 declaration in MOSEL, 375
 declaration in native languages, 372
 declaration in SLPDATA section, 366
 declaration in Visual Basic, 374
 declaration in XSLPchguserfunc, 371
 declaration in XSLPloaduserfuncs, 369
 Deltas, 380
 FunctionInfo, 379
 general, returning array by reference, 376
 general, returning array through argument, 377
 InputNames, 379
 instance, 381
 object, 382
 programming techniques, 379
 ReturnArray, 380
 returning derivatives, 380
 ReturnNames, 379
 simple, 376
 User function Derivatives, 387
 User function interface, 365
 User Function Object, 382
 User Function record, 9
 User Functions, 365
 User-defined convergence, 22

V
 Values of SLP variables in MISLP, fixing or relaxing, 394

Variable
 implicit, 4
 SLP, 2

W
 W-12023, 399
 W-12024, 399
 W-12028, 399
 W-12034, 400
 W-12036, 400
 W-12142, 402
 WT record in SLPDATA, 9

X
 Xpress-SLP problem pointer, 131
 Xpress-SLP Statistics, 31
 XSLP_ALGORITHM, 89
 XSLP_ATOL_A, 59
 XSLP_ATOL_R, 59
 XSLP_AUGMENTATION, 90
 XSLP_AUTOSAVE, 92
 XSLP_BARLIMIT, 92
 XSLP_CASCADE, 92
 XSLP_CASCADENLIMIT, 93
 XSLP_CASCADETOL_PA, 59
 XSLP_CASCADETOL_PR, 60
 XSLP_CDTOL_A, 60
 XSLP_CDTOL_R, 60
 XSLP_COEFFICIENTS, 40
 XSLP_CONTROL, 93
 XSLP_CTOL, 61
 XSLP_CURRENTDELTACOST, 37
 XSLP_CURRENTERRORCOST, 37
 XSLP_CVNAME, 124
 XSLP_CVS, 40
 XSLP_DAMP, 61
 XSLP_DAMPEXPAND, 61
 XSLP_DAMPMAX, 62

XSLP_DAMPMIN, 62
XSLP_DAMP SHRINK, 63
XSLP_DAMPSTART, 94
XSLP_DCLIMIT, 94
XSLP_DCLOG, 95
XSLP_DECOMPOSE, 95
XSLP_DECOMPOSEPASSLIMIT, 96
XSLP_DEFAULTIV, 63
XSLP_DEFAULTSTEPBOUND, 63
XSLP_DELAYUPDATEROWS, 95
XSLP_DELTA_A, 64
XSLP_DELTA_R, 64
XSLP_DELTA_X, 64
XSLP_DELTA_Z, 65
XSLP_DELTACOST, 65
XSLP_DELTACOSTFACTOR, 65
XSLP_DELTAFORMAT, 124
XSLP_DELTAMAXCOST, 66
XSLP_DELTAOFFSET, 96
XSLP_DELTAS, 40
XSLP_DELTAZLIMIT, 97
XSLP_DERIVATIVES, 97
XSLP_DJTOL, 66
XSLP_ECFCHECK, 98
XSLP_ECFCOUNT, 40
XSLP_ECFTOL_A, 66
XSLP_ECFTOL_R, 67
XSLP_EQTOL_A, 67
XSLP_EQTOL_R, 67
XSLP_EQUALS COLUMN, 40
XSLP_ERRORCOST, 68
XSLP_ERRORCOSTFACTOR, 68
XSLP_ERRORCOSTS, 37
XSLP_ERRORMAXCOST, 68
XSLP_ERROROFFSET, 98
XSLP_ERROR TOL_A, 69
XSLP_ERROR TOL_P, 69
XSLP_ESCALATION, 69
XSLP_ETOL_A, 70
XSLP_ETOL_R, 70
XSLP_EVALUATE, 99
XSLP_EVTOL_A, 70
XSLP_EVTOL_R, 71
XSLP_EXCELVISIBLE, 99
XSLP_EXPAND, 71
XSLP_EXTRACVS, 99
XSLP_EXTRAUFS, 100
XSLP_EXTRAXVITEMS, 100
XSLP_EXTRAXVS, 101
XSLP_FUNCEVAL, 101
XSLP_GLOBALFUNCOBJECT, 50
XSLP_GRANULARITY, 71
XSLP_IFS, 41
XSLP_IMPLICITVARIABLES, 41
XSLP_INFEASLIMIT, 102
XSLP_INFINITY, 72
XSLP_INTERNALFUNCCALLS, 41
XSLP_ITER, 41
XSLP_ITERLIMIT, 102
XSLP_ITOL_A, 72
XSLP_ITOL_R, 73
XSLP_IVNAME, 124
XSLP_LOG, 102
XSLP_MAXTIME, 103
XSLP_MAXWEIGHT, 73
XSLP_MEM_CALCSTACK, 119
XSLP_MEM_COEF, 120
XSLP_MEM_COL, 120
XSLP_MEM_CVAR, 120
XSLP_MEM_DERIVATIVES, 120
XSLP_MEM_EXCELDDOUBLE, 120
XSLP_MEM_FORMULA, 120
XSLP_MEM_FORMULAHASH, 121
XSLP_MEM_FORMULAVALUE, 121
XSLP_MEM_ITERLOG, 121
XSLP_MEM_RETURNARRAY, 121
XSLP_MEM_ROW, 121
XSLP_MEM_STACK, 121
XSLP_MEM_STRING, 122
XSLP_MEM_TOL, 122
XSLP_MEM_UF, 122
XSLP_MEM_VALSTACK, 122
XSLP_MEM_VAR, 122
XSLP_MEM_XF, 122
XSLP_MEM_XFNAMES, 123
XSLP_MEM_XFVALUE, 123
XSLP_MEM_XROW, 123
XSLP_MEM_XV, 123
XSLP_MEM_XVITEM, 123
XSLP_MEMORYFACTOR, 74
XSLP_MINORVERSION, 41
XSLP_MINUSDELTAFORMAT, 125
XSLP_MINUSERERRORFORMAT, 125
XSLP_MINUSPENALTYERRORS, 42
XSLP_MINWEIGHT, 74
XSLP_MIPALGORITHM, 103
XSLP_MIPCUTOFF_A, 74
XSLP_MIPCUTOFF_R, 75
XSLP_MIPCUTOFFCOUNT, 104
XSLP_MIPCUTOFFLIMIT, 104
XSLP_MIPDEFAULTALGORITHM, 105
XSLP_MIPERRORTOL_A, 75
XSLP_MIPERRORTOL_R, 76
XSLP_MIPFIXSTEPBOUNDS, 105
XSLP_MIPITER, 42
XSLP_MIPITERLIMIT, 106
XSLP_MIPLOG, 106
XSLP_MIPNODES, 42
XSLP_MIPOCOUNT, 106
XSLP_MIPOTOL_A, 76
XSLP_MIPOTOL_R, 76
XSLP_MIPPROBLEM, 50
XSLP_MIPRELAXSTEPBOUNDS, 107
XSLP_MTOL_A, 77
XSLP_MTOL_R, 77
XSLP_MVTOL, 78
XSLP_NONLINEARCONSTRAINTS, 42
XSLP_OBJSENSE, 37, 79
XSLP_OBJTOPENALTYCOST, 79
XSLP_OBJVAL, 38

XSLP_OCOUNT, 107
XSLP_OTOL_A, 80
XSLP_OTOL_R, 80
XSLP_PENALTYCOLFORMAT, 125
XSLP_PENALTYDELTACOLUMN, 42
XSLP_PENALTYDELTAROW, 43
XSLP_PENALTYDELTAS, 43
XSLP_PENALTYDELTATOTAL, 38
XSLP_PENALTYDELTAVALUE, 38
XSLP_PENALTYERRORCOLUMN, 43
XSLP_PENALTYERRORROW, 43
XSLP_PENALTYERRORS, 43
XSLP_PENALTYERRORTOTAL, 38
XSLP_PENALTYERRORVALUE, 38
XSLP_PENALTYINFOSTART, 108
XSLP_PENALTYROWFORMAT, 126
XSLP_PLUSDELTAFORMAT, 126
XSLP_PLUSEXACTFORMAT, 127
XSLP_PLUSPENALTYERRORS, 44
XSLP_PRESOLVE, 108
XSLP_PRESOLVEDELETEDDELTA, 44
XSLP_PRESOLVEFIXEDCOEF, 44
XSLP_PRESOLVEFIXEDDDR, 44
XSLP_PRESOLVEFIXEDNZCOL, 45
XSLP_PRESOLVEFIXEDSLPVAR, 45
XSLP_PRESOLVEFIXEDZCOL, 45
XSLP_PRESOLVEPASSES, 45
XSLP_PRESOLVEPASSLIMIT, 108
XSLP_PRESOLVETIGHTENED, 46
XSLP_PRESOLVEZERO, 81
XSLP_SAMECOUNT, 109
XSLP_SAMEDAMP, 109
XSLP_SBLOROWFORMAT, 127
XSLP_SBNAME, 127
XSLP_SBROWOFFSET, 110
XSLP_SBSTART, 110
XSLP_SBUPROWFORMAT, 128
XSLP_SBXCONVERGED, 46
XSLP_SCALE, 110
XSLP_SCALECOUNT, 111
XSLP_SHRINK, 81
XSLP_SLPLOG, 111
XSLP_STATUS, 46
XSLP_STOL_A, 81
XSLP_STOL_R, 82
XSLP_STOPOUTOFRANGE, 112
XSLP_TIMEPRINT, 112
XSLP_TOLNAME, 128
XSLP_TOLSETS, 47
XSLP_UCCONSTRAINEDCOUNT, 47
XSLP_UFINSTANCES, 47
XSLP_UFS, 47
XSLP_UNCONVERGED, 47
XSLP_UNFINISHEDLIMIT, 112
XSLP_UNIQUEPREFIX, 51
XSLP_UPDATEFORMAT, 128
XSLP_UPDATEOFFSET, 113
XSLP_USEDERIVATIVES, 48
XSLP_USERFUNCCALLS, 48
XSLP_VALIDATIONINDEX_A, 39
XSLP_VALIDATIONINDEX_R, 39
XSLP_VALIDATIONTOL_A, 82
XSLP_VALIDATIONTOL_R, 83
XSLP_VARIABLES, 48
XSLP_VCOUNT, 113
XSLP_VERSION, 48
XSLP_VERSIONDATE, 51
XSLP_VLIMIT, 114
XSLP_VSOLINDEX, 39
XSLP_VTOL_A, 83
XSLP_VTOL_R, 84
XSLP_WCOUNT, 114
XSLP_WTOL_A, 84
XSLP_WTOL_R, 85
XSLP_XCOUNT, 115
XSLP_XLIMIT, 116
XSLP_XPRSPROBLEM, 50
XSLP_XSLPPROBLEM, 50
XSLP_XTOL_A, 86
XSLP_XTOL_R, 87
XSLP_XVS, 48
XSLP_ZERO, 87
XSLP_ZEROCRITERION, 117
XSLP_ZEROCRITERIONCOUNT, 117
XSLP_ZEROCRITERIONSTART, 118
XSLP_ZEROESRESET, 49
XSLP_ZEROESRETAINED, 49
XSLP_ZEROESTOTAL, 49
XSLP_ATOL, 20
XSLP_CTOL, 20
XSLP_ITOL, 21
XSLP_MTOL, 20
XSLP_OCOUNT, 23
XSLP_OLIMIT, 23
XSLP_OTOL, 23
XSLP_STOL, 21
XSLP_VCOUNT, 22
XSLP_VLIMIT, 22
XSLP_VTOL, 22
XSLP_WCOUNT, 24
XSLP_WTOL, 24
XSLP_XCOUNT, 23
XSLP_XLIMIT, 23
XSLP_XTOL, 23
XSLPaddcoefs, 139
XSLPaddcvars, 141
XSLPadddcs, 142
XSLPadddfs, 144
XSLPadddivfs, 145
XSLPaddnames, 147
XSLPaddtolsets, 148
XSLPadduserfuncs, 149
XSLPaddvars, 151
XSLPaddxvs, 153
XSLPcalluserfunc, 155
XSLPcascade, 156
XSLPcascadeorder, 157
XSLPchgcoef, 158
XSLPchgcoef, 159
XSLPchgcvvar, 161

- XSLPchgdc, 162
- XSLPchgdf, 164
- XSLPchgfuncobject, 165
- XSLPchgivf, 166
- XSLPchgrow, 167
- XSLPchgrowwt, 168
- XSLPchgtolset, 169
- XSLPchguserfunc, 170
- XSLPchguserfuncaddress, 172
- XSLPchguserfuncobject, 173
- XSLPchgvar, 174
- XSLPchg xv, 176
- XSLPchg xvitem, 177
- XSLPconstruct, 179
- XSLPcopycallbacks, 180
- XSLPcopycontrols, 181
- XSLPcopyprob, 182
- XSLPcreateprob, 183
- XSLPdecompose, 184
- XSLPdestroyprob, 185
- XSLPevaluatecoef, 186
- XSLPevaluateformula, 187
- XSLPformatvalue, 188
- XSLPfree, 189
- XSLPgetbanner, 190
- XSLPgetcccoef, 191
- XSLPgetcoef, 192
- XSLPgetcvar, 193
- XSLPgetdblattr, 194
- XSLPgetdblcontrol, 195
- XSLPgetdc, 196
- XSLPgetdf, 197
- XSLPgetdtime, 198
- XSLPgetfuncinfo, 199
- XSLPgetfuncinfoV, 200
- XSLPgetfuncobject, 201
- XSLPgetfuncobjectV, 202
- XSLPgetindex, 203
- XSLPgetintattrib, 204
- XSLPgetintcontrol, 205
- XSLPgetivf, 206
- XSLPgetlasterror, 207
- XSLPgetmessagetype, 208
- XSLPgetnames, 209
- XSLPgetparam, 210
- XSLPgetptrattrib, 211
- XSLPgetrow, 212
- XSLPgetrowwt, 213
- XSLPgetslpsol, 214
- XSLPgetstrattrib, 215
- XSLPgetstrcontrol, 216
- XSLPgetstring, 217
- XSLPgettime, 218
- XSLPgettolset, 219
- XSLPgetuserfunc, 220
- XSLPgetuserfuncaddress, 222
- XSLPgetuserfuncobject, 223
- XSLPgetvar, 224
- XSLPgetversion, 226
- XSLPget xv, 227

- XSLPget xvitem, 228
- XSLPglobal, 230
- XSLPinit, 231
- XSLPitemname, 232
- XSLPload... functions, 132**
- XSLPloadcoefs, 233
- XSLPloadcvars, 235
- XSLPloaddds, 236
- XSLPloaddfs, 238
- XSLPloaddivfs, 239
- XSLPloadtolsets, 241
- XSLPloaduserfuncs, 242
- XSLPloadvars, 244
- XSLPloadxvs, 246
- XSLPmaxim, 248
- XSLPminim, 249
- XSLPopt, 250
- XSLPparsecformula, 251
- XSLPparseformula, 252
- XSLPpparseformula, 253
- XSLPpresolve, 254
- XSLPprintmsg, 255
- XSLPprob, 131**
- XSLPqparse, 256
- XSLPreadprob, 257
- XSLPpremaxim, 258
- XSLPpreminim, 259
- XSLPrestore, 260
- XSLPrevise, 261
- XSLProwinfo, 262
- XSLPsave, 263
- XSLPsaveas, 264
- XSLPscaling, 265
- XSLPsetcbcascadeend, 266
- XSLPsetcbcascadestart, 267
- XSLPsetcbcascadevar, 268
- XSLPsetcbcascadevarF, 270
- XSLPsetcbconstruct, 272
- XSLPsetcbdestroy, 274
- XSLPsetcbformula, 275
- XSLPsetcbintsol, 277
- XSLPsetcbiterend, 278
- XSLPsetcbiterstart, 279
- XSLPsetcbitervar, 280
- XSLPsetcbitervarF, 282
- XSLPsetcbmessage, 284
- XSLPsetcbmessageF, 286
- XSLPsetcboptnode, 288
- XSLPsetcbprenode, 289
- XSLPsetcbslpend, 291
- XSLPsetcbslpnode, 292
- XSLPsetcbslpstart, 293
- XSLPsetdblcontrol, 294
- XSLPsetdefaultcontrol, 295
- XSLPsetdefaults, 296
- XSLPsetfuncobject, 297
- XSLPsetfunctionerror, 298
- XSLPsetintcontrol, 299
- XSLPsetlogfile, 300
- XSLPsetparam, 301

- XSLPsetstrcontrol, [302](#)
- XSLPsetstring, [303](#)
- XSLPsetuniqueprefix, [304](#)
- XSLPsetuserfuncaddress, [305](#)
- XSLPsetuserfuncinfo, [306](#)
- XSLPsetuserfuncobject, [307](#)
- XSLPtime, [308](#)
- XSLPtokencount, [309](#)
- XSLPtoVBString, [310](#)
- XSLPuprintmemory, [311](#)
- XSLPuserfuncinfo, [312](#)
- XSLPvalidate, [315](#)
- XSLPvalidformula, [313](#)
- XSLPwriteprob, [316](#)
- XV record in SLPDATA, [10](#)