

Modeling Language GNU MathProg

Language Reference

for GLPK Version 4.45

(DRAFT, December 2010)

The GLPK package is part of the GNU Project released under the aegis of GNU.

Copyright © 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010
Andrew Makhorin, Department for Applied Informatics, Moscow Aviation In-
stitute, Moscow, Russia. All rights reserved.

Free Software Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA 02110-
1301, USA.

Permission is granted to make and distribute verbatim copies of this manual
provided the copyright notice and this permission notice are preserved on all
copies.

Permission is granted to copy and distribute modified versions of this manual
under the conditions for verbatim copying, provided also that the entire resulting
derived work is distributed under the terms of a permission notice identical to
this one.

Permission is granted to copy and distribute translations of this manual into
another language, under the above conditions for modified versions.

Contents

1	Introduction	6
1.1	Linear programming problem	6
1.2	Model objects	7
1.3	Structure of model description	8
2	Coding model description	9
2.1	Symbolic names	9
2.2	Numeric literals	10
2.3	String literals	10
2.4	Keywords	10
2.5	Delimiters	11
2.6	Comments	11
3	Expressions	12
3.1	Numeric expressions	12
3.1.1	Numeric literals	13
3.1.2	Dummy indices	13
3.1.3	Unsubscripted parameters	13
3.1.4	Subscripted parameters	13
3.1.5	Function references	13
3.1.6	Iterated expressions	14
3.1.7	Conditional expressions	15
3.1.8	Parenthesized expressions	15
3.1.9	Arithmetic operators	16
3.1.10	Hierarchy of operations	16
3.2	Symbolic expressions	17
3.2.1	Function references	17
3.2.2	Symbolic operators	18
3.2.3	Hierarchy of operations	18
3.3	Indexing expressions and dummy indices	18
3.4	Set expressions	22
3.4.1	Literal sets	22
3.4.2	Unsubscripted sets	22
3.4.3	Subscripted sets	23
3.4.4	“Arithmetic” sets	23
3.4.5	Indexing expressions	23
3.4.6	Iterated expressions	23
3.4.7	Conditional expressions	24
3.4.8	Parenthesized expressions	24
3.4.9	Set operators	24
3.4.10	Hierarchy of operations	25
3.5	Logical expressions	25
3.5.1	Numeric expressions	26
3.5.2	Relational operators	26

3.5.3	Iterated expressions	26
3.5.4	Parenthesized expressions	27
3.5.5	Logical operators	27
3.5.6	Hierarchy of operations	28
3.6	Linear expressions	28
3.6.1	Unsubscripted variables	28
3.6.2	Subscripted variables	29
3.6.3	Iterated expressions	29
3.6.4	Conditional expressions	29
3.6.5	Parenthesized expressions	29
3.6.6	Arithmetic operators	30
3.6.7	Hierarchy of operations	30
4	Statements	31
4.1	Set statement	31
4.2	Parameter statement	32
4.3	Variable statement	34
4.4	Constraint statement	36
4.5	Objective statement	37
4.6	Solve statement	38
4.7	Check statement	39
4.8	Display statement	39
4.9	Printf statement	40
4.10	For statement	41
4.11	Table statement	42
4.11.1	Table structure	43
4.11.2	Reading data from input table	43
4.11.3	Writing data to output table	44
5	Model data	45
5.1	Coding data section	46
5.2	Set data block	46
5.2.1	Assign data record	48
5.2.2	Slice data record	48
5.2.3	Simple data record	48
5.2.4	Matrix data record	48
5.2.5	Transposed matrix data record	49
5.3	Parameter data block	49
5.3.1	Assign data record	51
5.3.2	Slice data record	51
5.3.3	Plain data record	52
5.3.4	Tabular data record	52
5.3.5	Transposed tabular data record	53
5.3.6	Tabbing data format	53
A	Using suffixes	54

B	Date and time functions	55
B.1	Obtaining current calendar time	55
B.2	Converting character string to calendar time	55
B.3	Converting calendar time to character string	56
C	Table drivers	60
C.1	CSV table driver	60
C.2	xBASE table driver	61
C.3	ODBC table driver	62
C.4	MySQL table driver	64
D	Solving models with glpsol	66
E	Example model description	67
E.1	Model description written in MathProg	67
E.2	Generated LP problem instance	68
E.3	Optimal LP solution	68
	Acknowledgment	70

$-\infty < x < +\infty$	Free (unbounded) variable
$l \leq x < +\infty$	Variable with lower bound
$-\infty < x \leq u$	Variable with upper bound
$l \leq x \leq u$	Double-bounded variable
$l = x = u$	Fixed variable
$-\infty < \sum a_j x_j < +\infty$	Free (unbounded) linear form
$L \leq \sum a_j x_j < +\infty$	Inequality constraint “greater than or equal to”
$-\infty < \sum a_j x_j \leq U$	Inequality constraint “less than or equal to”
$L \leq \sum a_j x_j \leq U$	Double-bounded inequality constraint
$L = \sum a_j x_j = U$	Equality constraint

In addition to pure LP problems MathProg also allows mixed integer linear programming (MIP) problems, where some or all variables are restricted to be integer or binary.

1.2 Model objects

In MathProg the model is described in terms of sets, parameters, variables, constraints, and objectives, which are called *model objects*.

The user introduces particular model objects using the language statements. Each model object is provided with a symbolic name that uniquely identifies the object and is intended for referencing purposes.

Model objects, including sets, can be multidimensional arrays built over indexing sets. Formally, n -dimensional array A is the mapping:

$$A : \Delta \rightarrow \Xi, \quad (1.4)$$

where $\Delta \subseteq S_1 \times \dots \times S_n$ is a subset of the Cartesian product of indexing sets, Ξ is a set of array members. In MathProg the set Δ is called the *subscript domain*. Its members are n -tuples (i_1, \dots, i_n) , where $i_1 \in S_1, \dots, i_n \in S_n$.

If $n = 0$, the Cartesian product above has exactly one member (namely, 0-tuple), so it is convenient to think scalar objects as 0-dimensional arrays having one member.

The type of array members is determined by the type of corresponding model object as follows:

Model object	Array member
Set	Elemental plain set
Parameter	Number or symbol
Variable	Elemental variable
Constraint	Elemental constraint
Objective	Elemental objective

In order to refer to a particular object member the object should be provided with *subscripts*. For example, if a is a 2-dimensional parameter defined over $I \times J$, a reference to its particular member can be written as $a[i, j]$, where $i \in I$ and $j \in J$. It is understood that scalar objects being 0-dimensional need no subscripts.

1.3 Structure of model description

It is sometimes desirable to write a model which, at various points, may require different data for each problem instance to be solved using that model. For this reason in MathProg the model description consists of two parts: the *model section* and the *data section*.

The model section is a main part of the model description that contains declarations of model objects and is common for all problems based on the corresponding model.

The data section is an optional part of the model description that contains data specific for a particular problem instance.

Depending on what is more convenient the model and data sections can be placed either in one file or in two separate files. The latter feature allows having arbitrary number of different data sections to be used with the same model section.

2 Coding model description

The model description is coded in plain text format using ASCII character set. Characters valid in the model description are the following:

- alphabetic characters:
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
a b c d e f g h i j k l m n o p q r s t u v w x y z _
- numeric characters:
0 1 2 3 4 5 6 7 8 9
- special characters:
! " # & ' () * + , - . / : ; < = > [] ^ { | }
- white-space characters:
SP HT CR NL VT FF

Within string literals and comments any ASCII characters (except control characters) are valid.

White-space characters are non-significant. They can be used freely between lexical units to improve readability of the model description. They are also used to separate lexical units from each other if there is no other way to do that.

Syntactically model description is a sequence of lexical units in the following categories:

- symbolic names;
- numeric literals;
- string literals;
- keywords;
- delimiters;
- comments.

The lexical units of the language are discussed below.

2.1 Symbolic names

A *symbolic name* consists of alphabetic and numeric characters, the first of which must be alphabetic. All symbolic names are distinct (case sensitive).

Examples

```
alpha123  
This_is_a_name  
_P123_abc_321
```

Symbolic names are used to identify model objects (sets, parameters, variables, constraints, objectives) and dummy indices.

All symbolic names (except names of dummy indices) must be unique, i.e. the model description must have no objects with identical names. Symbolic names of dummy indices must be unique within the scope, where they are valid.

2.2 Numeric literals

A *numeric literal* has the form $xxEsyy$, where xx is a number with optional decimal point, s is the sign + or -, yy is a decimal exponent. The letter E is case insensitive and can be coded as e.

Examples

```
123
3.14159
56.E+5
.78
123.456e-7
```

Numeric literals are used to represent numeric quantities. They have obvious fixed meaning.

2.3 String literals

A *string literal* is a sequence of arbitrary characters enclosed either in single quotes or in double quotes. Both these forms are equivalent.

If the single quote is part of a string literal enclosed in single quotes, it must be coded twice. Analogously, if the double quote is part of a string literal enclosed in double quotes, it must be coded twice.

Examples

```
'This is a string'
"This is another string"
'1 + 2 = 3'
'That''s all'
"She said: ""No"""
```

String literals are used to represent symbolic quantities.

2.4 Keywords

A *keyword* is a sequence of alphabetic characters and possibly some special characters.

All keywords fall into two categories: *reserved keywords*, which cannot be used as symbolic names, and *non-reserved keywords*, which being recognized by context can be used as symbolic names.

The reserved keywords are the following:

<code>and</code>	<code>else</code>	<code>mod</code>	<code>union</code>
<code>by</code>	<code>if</code>	<code>not</code>	<code>within</code>
<code>cross</code>	<code>in</code>	<code>or</code>	
<code>diff</code>	<code>inter</code>	<code>syndiff</code>	
<code>div</code>	<code>less</code>	<code>then</code>	

Non-reserved keywords are described in following sections.

All the keywords have fixed meaning, which will be explained on discussion of corresponding syntactic constructions, where the keywords are used.

2.5 Delimiters

A *delimiter* is either a single special character or a sequence of two special characters as follows:

<code>+</code>	<code>^</code>	<code>==</code>	<code>!</code>	<code>:</code>	<code>)</code>
<code>-</code>	<code>&</code>	<code>>=</code>	<code>&&</code>	<code>;</code>	<code>[</code>
<code>*</code>	<code><</code>	<code>></code>	<code> </code>	<code>:=</code>	<code> </code>
<code>/</code>	<code><=</code>	<code><></code>	<code>.</code>	<code>..</code>	<code>{</code>
<code>**</code>	<code>=</code>	<code>!=</code>	<code>,</code>	<code>(</code>	<code>}</code>

If the delimiter consists of two characters, there must be no spaces between the characters.

All the delimiters have fixed meaning, which will be explained on discussion corresponding syntactic constructions, where the delimiters are used.

2.6 Comments

For documenting purposes the model description can be provided with *comments*, which may have two different forms. The first form is a *single-line comment*, which begins with the character `#` and extends until end of line. The second form is a *comment sequence*, which is a sequence of any characters enclosed within `/*` and `*/`.

Examples

```
param n := 10; # This is a comment
/* This is another comment */
```

Comments are ignored by the model translator and can appear anywhere in the model description, where white-space characters are allowed.

3 Expressions

An *expression* is a rule for computing a value. In model description expressions are used as constituents of certain statements.

In general case expressions consist of operands and operators.

Depending on the type of the resultant value all expressions fall into the following categories:

- numeric expressions;
- symbolic expressions;
- indexing expressions;
- set expressions;
- logical expressions;
- linear expressions.

3.1 Numeric expressions

A *numeric expression* is a rule for computing a single numeric value represented as a floating-point number.

The primary numeric expression may be a numeric literal, dummy index, unsubscripted parameter, subscripted parameter, built-in function reference, iterated numeric expression, conditional numeric expression, or another numeric expression enclosed in parentheses.

Examples

1.23	(numeric literal)
j	(dummy index)
time	(unsubscripted parameter)
a['May 2003',j+1]	(subscripted parameter)
abs(b[i,j])	(function reference)
sum{i in S diff T} alpha[i] * b[i,j]	(iterated expression)
if i in I then 2 * p else q[i+1]	(conditional expression)
(b[i,j] + .5 * c)	(parenthesized expression)

More general numeric expressions containing two or more primary numeric expressions may be constructed by using certain arithmetic operators.

Examples

```
j+1
2 * a[i-1,j+1] - b[i,j]
sum{j in J} a[i,j] * x[j] + sum{k in K} b[i,k] * x[k]
(if i in I then 2 * p else q[i+1]) / (a[i,j] + 1.5)
```

3.1.1 Numeric literals

If the primary numeric expression is a numeric literal, the resultant value is obvious.

3.1.2 Dummy indices

If the primary numeric expression is a dummy index, the resultant value is current value assigned to that dummy index.

3.1.3 Unsubscripted parameters

If the primary numeric expression is an unsubscripted parameter (which must be 0-dimensional), the resultant value is the value of that parameter.

3.1.4 Subscripted parameters

The primary numeric expression, which refers to a subscripted parameter, has the following syntactic form:

$$name[i_1, i_2, \dots, i_n]$$

where *name* is the symbolic name of the parameter, i_1, i_2, \dots, i_n are subscripts.

Each subscript must be a numeric or symbolic expression. The number of subscripts in the subscript list must be the same as the dimension of the parameter with which the subscript list is associated.

Actual values of subscript expressions are used to identify a particular member of the parameter that determines the resultant value of the primary expression.

3.1.5 Function references

In MathProg there exist the following built-in functions which may be used in numeric expressions:

<code>abs(<i>x</i>)</code>	$ x $, absolute value of x
<code>atan(<i>x</i>)</code>	$\arctan x$, principal value of the arc tangent of x (in radians)
<code>atan(<i>y</i>, <i>x</i>)</code>	$\arctan y/x$, principal value of the arc tangent of y/x (in radians). In this case the signs of both arguments y and x are used to determine the quadrant of the resultant value
<code>card(<i>X</i>)</code>	$ X $, cardinality (the number of elements) of set X
<code>ceil(<i>x</i>)</code>	$\lceil x \rceil$, smallest integer not less than x ("ceiling of x ")
<code>cos(<i>x</i>)</code>	$\cos x$, cosine of x (in radians)
<code>exp(<i>x</i>)</code>	e^x , base- e exponential of x
<code>floor(<i>x</i>)</code>	$\lfloor x \rfloor$, largest integer not greater than x ("floor of x ")

<code>gmtime()</code>	the number of seconds elapsed since 00:00:00 Jan 1, 1970, Coordinated Universal Time (for details see Subsection B.1, page 55)
<code>length(s)</code>	$ s $, length of character string s
<code>log(x)</code>	$\log x$, natural logarithm of x
<code>log10(x)</code>	$\log_{10} x$, common (decimal) logarithm of x
<code>max(x₁, x₂, ..., x_n)</code>	the largest of values x_1, x_2, \dots, x_n
<code>min(x₁, x₂, ..., x_n)</code>	the smallest of values x_1, x_2, \dots, x_n
<code>round(x)</code>	rounding x to nearest integer
<code>round(x, n)</code>	rounding x to n fractional decimal digits
<code>sin(x)</code>	$\sin x$, sine of x (in radians)
<code>sqrt(x)</code>	\sqrt{x} , non-negative square root of x
<code>str2time(s, f)</code>	converting character string s to calendar time (for details see Subsection B.2, page 55)
<code>trunc(x)</code>	truncating x to nearest integer
<code>trunc(x, n)</code>	truncating x to n fractional decimal digits
<code>Irand224()</code>	generating pseudo-random integer uniformly distributed in $[0, 2^{24})$
<code>Uniform01()</code>	generating pseudo-random number uniformly distributed in $[0, 1)$
<code>Uniform(a, b)</code>	generating pseudo-random number uniformly distributed in $[a, b)$
<code>Normal01()</code>	generating Gaussian pseudo-random variate with $\mu = 0$ and $\sigma = 1$
<code>Normal(μ, σ)</code>	generating Gaussian pseudo-random variate with given μ and σ

Arguments of all built-in functions, except `card`, `length`, and `str2time`, must be numeric expressions. The argument of `card` must be a set expression. The argument of `length` and both arguments of `str2time` must be symbolic expressions.

The resultant value of the numeric expression, which is a function reference, is the result of applying the function to its argument(s).

Note that each pseudo-random generator function has a latent argument (i.e. some internal state), which is changed whenever the function has been applied. Thus, if the function is applied repeatedly even to identical arguments, due to the side effect different resultant values are always produced.

3.1.6 Iterated expressions

An *iterated numeric expression* is a primary numeric expression, which has the following syntactic form:

$$\textit{iterated-operator indexing-expression integrand}$$

where *iterated-operator* is the symbolic name of the iterated operator to be performed (see below), *indexing-expression* is an indexing expression which intro-

duces dummy indices and controls iterating, *integrand* is a numeric expression that participates in the operation.

In MathProg there exist four iterated operators, which may be used in numeric expressions:

sum	summation	$\sum_{(i_1, \dots, i_n) \in \Delta} f(i_1, \dots, i_n)$
prod	production	$\prod_{(i_1, \dots, i_n) \in \Delta} f(i_1, \dots, i_n)$
min	minimum	$\min_{(i_1, \dots, i_n) \in \Delta} f(i_1, \dots, i_n)$
max	maximum	$\max_{(i_1, \dots, i_n) \in \Delta} f(i_1, \dots, i_n)$

where i_1, \dots, i_n are dummy indices introduced in the indexing expression, Δ is the domain, a set of n -tuples specified by the indexing expression which defines particular values assigned to the dummy indices on performing the iterated operation, $f(i_1, \dots, i_n)$ is the integrand, a numeric expression whose resultant value depends on the dummy indices.

The resultant value of an iterated numeric expression is the result of applying of the iterated operator to its integrand over all n -tuples contained in the domain.

3.1.7 Conditional expressions

A *conditional numeric expression* is a primary numeric expression, which has one of the following two syntactic forms:

if b **then** x **else** y
if b **then** x

where b is an logical expression, x and y are numeric expressions.

The resultant value of the conditional expression depends on the value of the logical expression that follows the keyword **if**. If it takes on the value *true*, the value of the conditional expression is the value of the expression that follows the keyword **then**. Otherwise, if the logical expression takes on the value *false*, the value of the conditional expression is the value of the expression that follows the keyword *else*. If the second, reduced form of the conditional expression is used and the logical expression takes on the value *false*, the resultant value of the conditional expression is zero.

3.1.8 Parenthesized expressions

Any numeric expression may be enclosed in parentheses that syntactically makes it a primary numeric expression.

Parentheses may be used in numeric expressions, as in algebra, to specify the desired order in which operations are to be performed. Where parentheses are used, the expression within the parentheses is evaluated before the resultant value is used.

The resultant value of the parenthesized expression is the same as the value of the expression enclosed within parentheses.

3.1.9 Arithmetic operators

In MathProg there exist the following arithmetic operators, which may be used in numeric expressions:

<code>+ x</code>	unary plus
<code>- x</code>	unary minus
<code>x + y</code>	addition
<code>x - y</code>	subtraction
<code>x less y</code>	positive difference (if $x < y$ then 0 else $x - y$)
<code>x * y</code>	multiplication
<code>x / y</code>	division
<code>x div y</code>	quotient of exact division
<code>x mod y</code>	remainder of exact division
<code>x ** y, x ^ y</code>	exponentiation (raising to power)

where x and y are numeric expressions.

If the expression includes more than one arithmetic operator, all operators are performed from left to right according to the hierarchy of operations (see below) with the only exception that the exponentiaion operators are performed from right to left.

The resultant value of the expression, which contains arithmetic operators, is the result of applying the operators to their operands.

3.1.10 Hierarchy of operations

The following list shows the hierarchy of operations in numeric expressions:

Operation	Hierarchy
Evaluation of functions (<code>abs</code> , <code>ceil</code> , etc.)	1st
Exponentiation (<code>**</code> , <code>^</code>)	2nd
Unary plus and minus (<code>+</code> , <code>-</code>)	3rd
Multiplication and division (<code>*</code> , <code>/</code> , <code>div</code> , <code>mod</code>)	4th
Iterated operations (<code>sum</code> , <code>prod</code> , <code>min</code> , <code>max</code>)	5th
Addition and subtraction (<code>+</code> , <code>-</code> , <code>less</code>)	6th
Conditional evaluation (<code>if ... then ... else</code>)	7th

This hierarchy is used to determine which of two consecutive operations is performed first. If the first operator is higher than or equal to the second, the first operation is performed. If it is not, the second operator is compared to the third, etc. When the end of the expression is reached, all of the remaining operations are performed in the reverse order.

3.2 Symbolic expressions

A *symbolic expression* is a rule for computing a single symbolic value represented as a character string.

The primary symbolic expression may be a string literal, dummy index, unsubscripted parameter, subscripted parameter, built-in function reference, conditional symbolic expression, or another symbolic expression enclosed in parentheses.

It is also allowed to use a numeric expression as the primary symbolic expression, in which case the resultant value of the numeric expression is automatically converted to the symbolic type.

Examples

<code>'May 2003'</code>	(string literal)
<code>j</code>	(dummy index)
<code>p</code>	(unsubscripted parameter)
<code>s['abc',j+1]</code>	(subscripted parameter)
<code>substr(name[i],k+1,3)</code>	(function reference)
<code>if i in I then s[i,j] else t[i+1]</code>	(conditional expression)
<code>((10 * b[i,j]) & '.bis')</code>	(parenthesized expression)

More general symbolic expressions containing two or more primary symbolic expressions may be constructed by using the concatenation operator.

Examples

```
'abc[' & i & ',' & j & ']'  
"from " & city[i] & " to " & city[j]
```

The principles of evaluation of symbolic expressions are completely analogous to the ones given for numeric expressions (see above).

3.2.1 Function references

In MathProg there exist the following built-in functions which may be used in symbolic expressions:

<code>substr(<i>s</i>, <i>x</i>)</code>	substring of <i>s</i> starting from position <i>x</i>
<code>substr(<i>s</i>, <i>x</i>, <i>y</i>)</code>	substring of <i>s</i> starting from position <i>x</i> and having length <i>y</i>
<code>time2str(<i>t</i>, <i>f</i>)</code>	converting calendar time to character string (for details see Subsection B.3 , page 56)

The first argument of `substr` must be a symbolic expression while its second and optional third arguments must be numeric expressions.

The first argument of `time2str` must be a numeric expression, and its second argument must be a symbolic expression.

The resultant value of the symbolic expression, which is a function reference, is the result of applying the function to its arguments.

3.2.2 Symbolic operators

Currently in MathProg there exists the only symbolic operator:

$$s \ \& \ t$$

where s and t are symbolic expressions. This operator means concatenation of its two symbolic operands, which are character strings.

3.2.3 Hierarchy of operations

The following list shows the hierarchy of operations in symbolic expressions:

Operation	Hierarchy
Evaluation of numeric operations	1st-7th
Concatenation (&)	8th
Conditional evaluation (if ... then ... else)	7th

This hierarchy has the same meaning as was explained above for numeric expressions (see Subsection 3.1.10, page 16).

3.3 Indexing expressions and dummy indices

An *indexing expression* is an auxiliary construction, which specifies a plain set of n -tuples and introduces dummy indices. It has two syntactic forms:

$$\{ \text{entry}_1, \text{entry}_2, \dots, \text{entry}_m \}$$
$$\{ \text{entry}_1, \text{entry}_2, \dots, \text{entry}_m : \text{predicate} \}$$

where $\text{entry}_1, \text{entry}_2, \dots, \text{entry}_m$ are indexing entries, *predicate* is a logical expression that specifies an optional predicate (logical condition).

Each *indexing entry* in the indexing expression has one of the following three forms:

$$i \ \text{in} \ S$$
$$(i_1, i_2, \dots, i_n) \ \text{in} \ S$$
$$S$$

where i_1, i_2, \dots, i_n are indices, S is a set expression (discussed in the next section) that specifies the basic set.

The number of indices in the indexing entry must be the same as the dimension of the basic set S , i.e. if S consists of 1-tuples, the first form must be used, and if S consists of n -tuples, where $n > 1$, the second form must be used.

If the first form of the indexing entry is used, the index i can be a dummy index only (see below). If the second form is used, the indices i_1, i_2, \dots, i_n can be either dummy indices or some numeric or symbolic expressions, where at least one index must be a dummy index. The third, reduced form of the indexing entry has the same effect as if there were i (if S is 1-dimensional) or i_1, i_2, \dots, i_n (if S is n -dimensional) all specified as dummy indices.

A *dummy index* is an auxiliary model object, which acts like an individual variable. Values assigned to dummy indices are components of n -tuples from basic sets, i.e. some numeric and symbolic quantities.

For referencing purposes dummy indices can be provided with symbolic names. However, unlike other model objects (sets, parameters, etc.) dummy indices need not be explicitly declared. Each *undeclared* symbolic name being used in the indexing position of an indexing entry is recognized as the symbolic name of corresponding dummy index.

Symbolic names of dummy indices are valid only within the scope of the indexing expression, where the dummy indices were introduced. Beyond the scope the dummy indices are completely inaccessible, so the same symbolic names may be used for other purposes, in particular, to represent dummy indices in other indexing expressions.

The scope of indexing expression, where implicit declarations of dummy indices are valid, depends on the context, in which the indexing expression is used:

1. If the indexing expression is used in iterated operator, its scope extends until the end of the integrand.
2. If the indexing expression is used as a primary set expression, its scope extends until the end of that indexing expression.
3. If the indexing expression is used to define the subscript domain in declarations of some model objects, its scope extends until the end of the corresponding statement.

The indexing mechanism implemented by means of indexing expressions is best explained by some examples discussed below.

Let there be given three sets:

$$A = \{4, 7, 9\},$$

$$B = \{(1, Jan), (1, Feb), (2, Mar), (2, Apr), (3, May), (3, Jun)\},$$

$$C = \{a, b, c\},$$

where A and C consist of 1-tuples (singlets), B consists of 2-tuples (doublets). Consider the following indexing expression:

$$\{i \text{ in } A, (j, k) \text{ in } B, l \text{ in } C\}$$

where i , j , k , and l are dummy indices.

Although MathProg is not a procedural language, for any indexing expression an equivalent algorithmic description can be given. In particular, the algorithmic description of the indexing expression above could look like follows:

```

for all  $i \in A$  do
  for all  $(j, k) \in B$  do
    for all  $l \in C$  do
      action;

```

where the dummy indices i, j, k, l are consecutively assigned corresponding components of n -tuples from the basic sets A, B, C , and *action* is some action that depends on the context, where the indexing expression is used. For example, if the action were printing current values of dummy indices, the printout would look like follows:

$$\begin{array}{llll} i = 4 & j = 1 & k = Jan & l = a \\ i = 4 & j = 1 & k = Jan & l = b \\ i = 4 & j = 1 & k = Jan & l = c \\ i = 4 & j = 1 & k = Feb & l = a \\ i = 4 & j = 1 & k = Feb & l = b \\ & \dots & \dots & \dots \\ i = 9 & j = 3 & k = Jun & l = b \\ i = 9 & j = 3 & k = Jun & l = c \end{array}$$

Let the example indexing expression be used in the following iterated operation:

$$\text{sum}\{i \text{ in } A, (j,k) \text{ in } B, l \text{ in } C\} p[i,j,k,l]$$

where p is a 4-dimensional numeric parameter or some numeric expression whose resultant value depends on i, j, k , and l . In this case the action is summation, so the resultant value of the primary numeric expression is:

$$\sum_{i \in A, (j,k) \in B, l \in C} (p_{ijkl}).$$

Now let the example indexing expression be used as a primary set expression. In this case the action is gathering all 4-tuples (quadruplets) of the form (i, j, k, l) in one set, so the resultant value of such operation is simply the Cartesian product of the basic sets:

$$A \times B \times C = \{(i, j, k, l) : i \in A, (j, k) \in B, l \in C\}.$$

Note that in this case the same indexing expression might be written in the reduced form:

$$\{A, B, C\}$$

because the dummy indices i, j, k , and l are not referenced and therefore their symbolic names need not be specified.

Finally, let the example indexing expression be used as the subscript domain in the declaration of a 4-dimensional model object, say, a numeric parameter:

$$\text{param } p\{i \text{ in } A, (j,k) \text{ in } B, l \text{ in } C\} \dots;$$

In this case the action is generating the parameter members, where each member has the form $p[i, j, k, l]$.

As was said above, some indices in the second form of indexing entries may be numeric or symbolic expressions, not only dummy indices. In this case resultant

values of such expressions play role of some logical conditions to select only that n -tuples from the Cartesian product of basic sets that satisfy these conditions.

Consider, for example, the following indexing expression:

$$\{i \text{ in } A, (i-1, k) \text{ in } B, l \text{ in } C\}$$

where i, k, l are dummy indices, and $i-1$ is a numeric expression. The algorithmic description of this indexing expression is the following:

```

for all  $i \in A$  do
  for all  $(j, k) \in B$  and  $j = i - 1$  do
    for all  $l \in C$  do
      action;

```

Thus, if this indexing expression were used as a primary set expression, the resultant set would be the following:

$$\{(4, \text{May}, a), (4, \text{May}, b), (4, \text{May}, c), (4, \text{Jun}, a), (4, \text{Jun}, b), (4, \text{Jun}, c)\}.$$

Should note that in this case the resultant set consists of 3-tuples, not of 4-tuples, because in the indexing expression there is no dummy index that corresponds to the first component of 2-tuples from the set B .

The general rule is: the number of components of n -tuples defined by an indexing expression is the same as the number of dummy indices in that expression, where the correspondence between dummy indices and components on n -tuples in the resultant set is positional, i.e. the first dummy index corresponds to the first component, the second dummy index corresponds to the second component, etc.

In some cases it is needed to select a subset from the Cartesian product of some sets. This may be attained by using an optional logical predicate, which is specified in the indexing expression.

Consider, for example, the following indexing expression:

$$\{i \text{ in } A, (j, k) \text{ in } B, l \text{ in } C: i \leq 5 \text{ and } k \neq \text{'Mar'}\}$$

where the logical expression following the colon is a predicate. The algorithmic description of this indexing expression is the following:

```

for all  $i \in A$  do
  for all  $(j, k) \in B$  do
    for all  $l \in C$  do
      if  $i \leq 5$  and  $l \neq \text{'Mar'}$  then
        action;

```

Thus, if this indexing expression were used as a primary set expression, the resultant set would be the following:

$$\{(4, 1, \text{Jan}, a), (4, 1, \text{Feb}, a), (4, 2, \text{Apr}, a), \dots, (4, 3, \text{Jun}, c)\}.$$

If no predicate is specified in the indexing expression, one, which takes on the value *true*, is assumed.

3.4 Set expressions

A *set expression* is a rule for computing an elemental set, i.e. a collection of n -tuples, where components of n -tuples are numeric and symbolic quantities.

The primary set expression may be a literal set, unsubscripted set, subscripted set, “arithmetic” set, indexing expression, iterated set expression, conditional set expression, or another set expression enclosed in parentheses.

Examples

<code>{(123, 'aa'), (i, 'bb'), (j-1, 'cc')}</code>	(literal set)
<code>I</code>	(unsubscripted set)
<code>S[i-1, j+1]</code>	(subscripted set)
<code>1..t-1 by 2</code>	(“arithmetic” set)
<code>{t in 1..T, (t+1, j) in S: (t, j) in F}</code>	(indexing expression)
<code>setof{i in I, j in J}(i+1, j-1)</code>	(iterated expression)
<code>if i < j then S[i] else F diff S[j]</code>	(conditional expression)
<code>(1..10 union 21..30)</code>	(parenthesized expression)

More general set expressions containing two or more primary set expressions may be constructed by using certain set operators.

Examples

```
(A union B) inter (I cross J)
1..10 cross (if i < j then {'a', 'b', 'c'} else {'d', 'e', 'f'})
```

3.4.1 Literal sets

A *literal set* is a primary set expression, which has the following two syntactic forms:

$$\{e_1, e_2, \dots, e_m\}$$

$$\{(e_{11}, \dots, e_{1n}), (e_{21}, \dots, e_{2n}), \dots, (e_{m1}, \dots, e_{mn})\}$$

where $e_1, \dots, e_m, e_{11}, \dots, e_{mn}$ are numeric or symbolic expressions.

If the first form is used, the resultant set consists of 1-tuples (singlets) enumerated within the curly braces. It is allowed to specify an empty set as $\{ \}$, which has no 1-tuples. If the second form is used, the resultant set consists of n -tuples enumerated within the curly braces, where a particular n -tuple consists of corresponding components enumerated within the parentheses. All n -tuples must have the same number of components.

3.4.2 Unsubscripted sets

If the primary set expression is an unsubscripted set (which must be 0-dimensional), the resultant set is an elemental set associated with the corresponding set object.

3.4.3 Subscripted sets

The primary set expression, which refers to a subscripted set, has the following syntactic form:

$$name[i_1, i_2, \dots, i_n]$$

where $name$ is the symbolic name of the set object, i_1, i_2, \dots, i_n are subscripts.

Each subscript must be a numeric or symbolic expression. The number of subscripts in the subscript list must be the same as the dimension of the set object with which the subscript list is associated.

Actual values of subscript expressions are used to identify a particular member of the set object that determines the resultant set.

3.4.4 “Arithmetic” sets

The primary set expression, which is an “arithmetic” set, has the following two syntactic forms:

$$t_0 \dots t_1 \text{ by } \delta t$$
$$t_0 \dots t_1$$

where t_0, t_1 , and δt are numeric expressions (the value of δt must not be zero). The second form is equivalent to the first form, where $\delta t = 1$.

If $\delta t > 0$, the resultant set is determined as follows:

$$\{t : \exists k \in \mathcal{Z}(t = t_0 + k\delta t, t_0 \leq t \leq t_1)\}.$$

Otherwise, if $\delta t < 0$, the resultant set is determined as follows:

$$\{t : \exists k \in \mathcal{Z}(t = t_0 + k\delta t, t_1 \leq t \leq t_0)\}.$$

3.4.5 Indexing expressions

If the primary set expression is an indexing expression, the resultant set is determined as described above in Subsection 3.3, page 18.

3.4.6 Iterated expressions

An *iterated set expression* is a primary set expression, which has the following syntactic form:

$$\text{setof } indexing-expression \text{ integrand}$$

where *indexing-expression* is an indexing expression, which introduces dummy indices and controls iterating, *integrand* is either a single numeric or symbolic expression or a list of numeric and symbolic expressions separated by commas and enclosed in parentheses.

If the integrand is a single numeric or symbolic expression, the resultant set consists of 1-tuples and is determined as follows:

$$\{x : (i_1, \dots, i_n) \in \Delta\},$$

where x is a value of the integrand, i_1, \dots, i_n are dummy indices introduced in the indexing expression, Δ is the domain, a set of n -tuples specified by the indexing expression, which defines particular values assigned to the dummy indices on performing the iterated operation.

If the integrand is a list containing m numeric and symbolic expressions, the resultant set consists of m -tuples and is determined as follows:

$$\{(x_1, \dots, x_m) : (i_1, \dots, i_n) \in \Delta\},$$

where x_1, \dots, x_m are values of the expressions in the integrand list, i_1, \dots, i_n and Δ have the same meaning as above.

3.4.7 Conditional expressions

A *conditional set expression* is a primary set expression that has the following syntactic form:

if b **then** X **else** Y

where b is an logical expression, X and Y are set expressions, which must define sets of the same dimension.

The resultant value of the conditional expression depends on the value of the logical expression that follows the keyword **if**. If it takes on the value *true*, the resultant set is the value of the expression that follows the keyword **then**. Otherwise, if the logical expression takes on the value *false*, the resultant set is the value of the expression that follows the keyword **else**.

3.4.8 Parenthesized expressions

Any set expression may be enclosed in parentheses that syntactically makes it a primary set expression.

Parentheses may be used in set expressions, as in algebra, to specify the desired order in which operations are to be performed. Where parentheses are used, the expression within the parentheses is evaluated before the resultant value is used.

The resultant value of the parenthesized expression is the same as the value of the expression enclosed within parentheses.

3.4.9 Set operators

In MathProg there exist the following set operators, which may be used in set expressions:

X union Y	union $X \cup Y$
X diff Y	difference $X \setminus Y$
X symdiff Y	symmetric difference $X \oplus Y$
X inter Y	intersection $X \cap Y$
X cross Y	cross (Cartesian) product $X \times Y$

where X and Y are set expressions, which must define sets of the identical dimension (except the Cartesian product).

If the expression includes more than one set operator, all operators are performed from left to right according to the hierarchy of operations (see below).

The resultant value of the expression, which contains set operators, is the result of applying the operators to their operands.

The dimension of the resultant set, i.e. the dimension of n -tuples, of which the resultant set consists of, is the same as the dimension of the operands, except the Cartesian product, where the dimension of the resultant set is the sum of the dimensions of its operands.

3.4.10 Hierarchy of operations

The following list shows the hierarchy of operations in set expressions:

Operation	Hierarchy
Evaluation of numeric operations	1st-7th
Evaluation of symbolic operations	8th-9th
Evaluation of iterated or “arithmetic” set (<code>setof</code> , ..)	10th
Cartesian product (<code>cross</code>)	11th
Intersection (<code>inter</code>)	12th
Union and difference (<code>union</code> , <code>diff</code> , <code>syndiff</code>)	13th
Conditional evaluation (<code>if ... then ... else</code>)	14th

This hierarchy has the same meaning as was explained above for numeric expressions (see Subsection 3.1.10, page 16).

3.5 Logical expressions

A *logical expression* is a rule for computing a single logical value, which can be either *true* or *false*.

The primary logical expression may be a numeric expression, relational expression, iterated logical expression, or another logical expression enclosed in parentheses.

Examples

<code>i+1</code>	(numeric expression)
<code>a[i,j] < 1.5</code>	(relational expression)
<code>s[i+1,j-1] <> 'Mar'</code>	(relational expression)
<code>(i+1,'Jan') not in I cross J</code>	(relational expression)
<code>S union T within A[i] inter B[j]</code>	(relational expression)
<code>forall{i in I, j in J} a[i,j] < .5 * b</code>	(iterated expression)
<code>(a[i,j] < 1.5 or b[i] >= a[i,j])</code>	(parenthesized expression)

More general logical expressions containing two or more primary logical expressions may be constructed by using certain logical operators.

Examples

not (a[i,j] < 1.5 or b[i] >= a[i,j]) and (i,j) in S
(i,j) in S or (i,j) not in T diff U

3.5.1 Numeric expressions

The resultant value of the primary logical expression, which is a numeric expression, is *true*, if the resultant value of the numeric expression is non-zero. Otherwise the resultant value of the logical expression is *false*.

3.5.2 Relational operators

In MathProg there exist the following relational operators, which may be used in logical expressions:

$x < y$	test on $x < y$
$x \leq y$	test on $x \leq y$
$x = y, x == y$	test on $x = y$
$x \geq y$	test on $x \geq y$
$x > y$	test on $x > y$
$x \langle \rangle y, x \neq y$	test on $x \neq y$
x in Y	test on $x \in Y$
(x_1, \dots, x_n) in Y	test on $(x_1, \dots, x_n) \in Y$
x not in $Y, x !$ in Y	test on $x \notin Y$
(x_1, \dots, x_n) not in $Y, (x_1, \dots, x_n) !$ in Y	test on $(x_1, \dots, x_n) \notin Y$
X within Y	test on $X \subseteq Y$
X not within $Y, X !$ within Y	test on $X \not\subseteq Y$

where x, x_1, \dots, x_n, y are numeric or symbolic expressions, X and Y are set expression.

Notes:

1. In the operations **in**, **not in**, and **!in** the number of components in the first operands must be the same as the dimension of the second operand.

2. In the operations **within**, **not within**, and **!within** both operands must have identical dimension.

All the relational operators listed above have their conventional mathematical meaning. The resultant value is *true*, if corresponding relation is satisfied for its operands, otherwise *false*. (Note that symbolic values are ordered lexicographically, and any numeric value precedes any symbolic value.)

3.5.3 Iterated expressions

An *iterated logical expression* is a primary logical expression, which has the following syntactic form:

iterated-operator indexing-expression integrand

where *iterated-operator* is the symbolic name of the iterated operator to be performed (see below), *indexing-expression* is an indexing expression which introduces dummy indices and controls iterating, *integrand* is a numeric expression that participates in the operation.

In MathProg there exist two iterated operators, which may be used in logical expressions:

forall	\forall -quantification	$\forall(i_1, \dots, i_n) \in \Delta[f(i_1, \dots, i_n)],$
exists	\exists -quantification	$\exists(i_1, \dots, i_n) \in \Delta[f(i_1, \dots, i_n)],$

where i_1, \dots, i_n are dummy indices introduced in the indexing expression, Δ is the domain, a set of n -tuples specified by the indexing expression which defines particular values assigned to the dummy indices on performing the iterated operation, $f(i_1, \dots, i_n)$ is the integrand, a logical expression whose resultant value depends on the dummy indices.

For \forall -quantification the resultant value of the iterated logical expression is *true*, if the value of the integrand is *true* for all n -tuples contained in the domain, otherwise *false*.

For \exists -quantification the resultant value of the iterated logical expression is *false*, if the value of the integrand is *false* for all n -tuples contained in the domain, otherwise *true*.

3.5.4 Parenthesized expressions

Any logical expression may be enclosed in parentheses that syntactically makes it a primary logical expression.

Parentheses may be used in logical expressions, as in algebra, to specify the desired order in which operations are to be performed. Where parentheses are used, the expression within the parentheses is evaluated before the resultant value is used.

The resultant value of the parenthesized expression is the same as the value of the expression enclosed within parentheses.

3.5.5 Logical operators

In MathProg there exist the following logical operators, which may be used in logical expressions:

not x , $!x$	negation $\neg x$
x and y , $x \ \&\& \ y$	conjunction (logical “and”) $x \ \& \ y$
x or y , $x \ \ y$	disjunction (logical “or”) $x \ \vee \ y$

where x and y are logical expressions.

If the expression includes more than one logical operator, all operators are performed from left to right according to the hierarchy of the operations (see below). The resultant value of the expression, which contains logical operators, is the result of applying the operators to their operands.

3.5.6 Hierarchy of operations

The following list shows the hierarchy of operations in logical expressions:

Operation	Hierarchy
Evaluation of numeric operations	1st-7th
Evaluation of symbolic operations	8th-9th
Evaluation of set operations	10th-14th
Relational operations (<, <=, etc.)	15th
Negation (not , !)	16th
Conjunction (and , &&)	17th
\forall - and \exists -quantification (forall , exists)	18th
Disjunction (or ,)	19th

This hierarchy has the same meaning as was explained above for numeric expressions (see Subsection 3.1.10, page 16).

3.6 Linear expressions

An *linear expression* is a rule for computing so called a *linear form* or simply a *formula*, which is a linear (or affine) function of elemental variables.

The primary linear expression may be an unsubscripted variable, subscripted variable, iterated linear expression, conditional linear expression, or another linear expression enclosed in parentheses.

It is also allowed to use a numeric expression as the primary linear expression, in which case the resultant value of the numeric expression is automatically converted to a formula that includes the constant term only.

Examples

<code>z</code>	(unsubscripted variable)
<code>x[i,j]</code>	(subscripted variable)
<code>sum{j in J} (a[i] * x[i,j] + 3 * y)</code>	(iterated expression)
<code>if i in I then x[i,j] else 1.5 * z + 3</code>	(conditional expression)
<code>(a[i,j] * x[i,j] + y[i-1] + .1)</code>	(parenthesized expression)

More general linear expressions containing two or more primary linear expressions may be constructed by using certain arithmetic operators.

Examples

```
2 * x[i-1,j+1] + 3.5 * y[k] + .5 * z
(- x[i,j] + 3.5 * y[k]) / sum{t in T} abs(d[i,j,t])
```

3.6.1 Unsubscripted variables

If the primary linear expression is an unsubscripted variable (which must be 0-dimensional), the resultant formula is that unsubscripted variable.

3.6.2 Subscripted variables

The primary linear expression, which refers to a subscripted variable, has the following syntactic form:

$$name[i_1, i_2, \dots, i_n]$$

where *name* is the symbolic name of the model variable, i_1, i_2, \dots, i_n are subscripts.

Each subscript must be a numeric or symbolic expression. The number of subscripts in the subscript list must be the same as the dimension of the model variable with which the subscript list is associated.

Actual values of the subscript expressions are used to identify a particular member of the model variable that determines the resultant formula, which is an elemental variable associated with corresponding member.

3.6.3 Iterated expressions

An *iterated linear expression* is a primary linear expression, which has the following syntactic form:

$$\text{sum } indexing-expression \text{ integrand}$$

where *indexing-expression* is an indexing expression, which introduces dummy indices and controls iterating, *integrand* is a linear expression that participates in the operation.

The iterated linear expression is evaluated exactly in the same way as the iterated numeric expression (see Subsection 3.1.6, page 14) with exception that the integrand participated in the summation is a formula, not a numeric value.

3.6.4 Conditional expressions

A *conditional linear expression* is a primary linear expression, which has one of the following two syntactic forms:

$$\text{if } b \text{ then } f \text{ else } g$$

$$\text{if } b \text{ then } f$$

where *b* is an logical expression, *f* and *g* are linear expressions.

The conditional linear expression is evaluated exactly in the same way as the conditional numeric expression (see Subsection 3.1.7, page 15) with exception that operands participated in the operation are formulae, not numeric values.

3.6.5 Parenthesized expressions

Any linear expression may be enclosed in parentheses that syntactically makes it a primary linear expression.

Parentheses may be used in linear expressions, as in algebra, to specify the desired order in which operations are to be performed. Where parentheses are

used, the expression within the parentheses is evaluated before the resultant formula is used.

The resultant value of the parenthesized expression is the same as the value of the expression enclosed within parentheses.

3.6.6 Arithmetic operators

In MathProg there exists the following arithmetic operators, which may be used in linear expressions:

$+ f$	unary plus
$- f$	unary minus
$f + g$	addition
$f - g$	subtraction
$x * f, f * x$	multiplication
f / x	division

where f and g are linear expressions, x is a numeric expression (more precisely, a linear expression containing only the constant term).

If the expression includes more than one arithmetic operator, all operators are performed from left to right according to the hierarchy of operations (see below). The resultant value of the expression, which contains arithmetic operators, is the result of applying the operators to their operands.

3.6.7 Hierarchy of operations

The hierarchy of arithmetic operations used in linear expressions is the same as for numeric expressions (see Subsection [3.1.10](#), page [16](#)).

4 Statements

Statements are basic units of the model description. In MathProg all statements are divided into two categories: declaration statements and functional statements.

Declaration statements (set statement, parameter statement, variable statement, constraint statement, and objective statement) are used to declare model objects of certain kinds and define certain properties of such objects.

Functional statements (solve statement, check statement, display statement, printf statement, loop statement) are intended for performing some specific actions.

Note that declaration statements may follow in arbitrary order, which does not affect the result of translation. However, any model object must be declared before it is referenced in other statements.

4.1 Set statement

```
set name alias domain , attrib , ... , attrib ;
```

Where: *name* is a symbolic name of the set;

alias is an optional string literal, which specifies an alias of the set;

domain is an optional indexing expression, which specifies a subscript domain of the set;

attrib, ..., attrib are optional attributes of the set. (Commae preceding attributes may be omitted.)

Optional attributes:

dimen *n* specifies the dimension of *n*-tuples, which the set consists of;

within *expression*
specifies a superset which restricts the set or all its members (elemental sets) to be within that superset;

:= *expression*
specifies an elemental set assigned to the set or its members;

default *expression*
specifies an elemental set assigned to the set or its members whenever no appropriate data are available in the data section.

Examples

```
set V;  
set E within V cross V;  
set step{s in 1..maxiter} dimen 2 := if s = 1 then E else  
  step[s-1] union setof{k in V, (i,k) in step[s-1], (k,j)  
  in step[s-1]}(i,j);  
set A{i in I, j in J}, within B[i+1] cross C[j-1], within  
  D diff E, default {'abc',123}, (321,'cba');
```

The set statement declares a set. If the subscript domain is not specified, the set is a simple set, otherwise it is an array of elemental sets.

The **dimen** attribute specifies the dimension of n -tuples, which the set (if it is a simple set) or its members (if the set is an array of elemental sets) consist of, where n must be unsigned integer from 1 to 20. At most one **dimen** attribute can be specified. If the **dimen** attribute is not specified, the dimension of n -tuples is implicitly determined by other attributes (for example, if there is a set expression that follows **:=** or the keyword **default**, the dimension of n -tuples of corresponding elemental set is used). If no dimension information is available, **dimen 1** is assumed.

The **within** attribute specifies a set expression whose resultant value is a superset used to restrict the set (if it is a simple set) or its members (if the set is an array of elemental sets) to be within that superset. Arbitrary number of **within** attributes may be specified in the same set statement.

The assign (**:=**) attribute specifies a set expression used to evaluate elemental set(s) assigned to the set (if it is a simple set) or its members (if the set is an array of elemental sets). If the assign attribute is specified, the set is *computable* and therefore needs no data to be provided in the data section. If the assign attribute is not specified, the set must be provided with data in the data section. At most one assign or default attribute can be specified for the same set.

The **default** attribute specifies a set expression used to evaluate elemental set(s) assigned to the set (if it is a simple set) or its members (if the set is an array of elemental sets) whenever no appropriate data are available in the data section. If neither assign nor default attribute is specified, missing data will cause an error.

4.2 Parameter statement

<pre>param name alias domain , attrib , ... , attrib ;</pre>
--

Where: *name* is a symbolic name of the parameter;

alias is an optional string literal, which specifies an alias of the parameter;

domain is an optional indexing expression, which specifies a subscript domain of the parameter;

attrib, ..., attrib are optional attributes of the parameter. (Commae preceding attributes may be omitted.)

Optional attributes:

integer specifies that the parameter is integer;
binary specifies that the parameter is binary;
symbolic specifies that the parameter is symbolic;

relation expression

(where *relation* is one of: <, <=, =, ==, >=, >, <>, !=)
 specifies a condition that restricts the parameter or its members to satisfy that condition;

in expression

specifies a superset that restricts the parameter or its members to be in that superset;

:= expression

specifies a value assigned to the parameter or its members;

default expression

specifies a value assigned to the parameter or its members whenever no appropriate data are available in the data section.

Examples

```
param units{raw, prd} >= 0;
param profit{prd, 1..T+1};
param N := 20, integer, >= 0, <= 100;
param comb 'n choose k' {n in 0..N, k in 0..n} :=
  if k = 0 or k = n then 1 else comb[n-1,k-1] + comb[n-1,k];
param p{i in I, j in J}, integer, >= 0, <= i+j,
  in A[i] symdiff B[j], in C[i,j], default 0.5 * (i + j);
param month symbolic default 'May' in {'Mar', 'Apr', 'May'};
```

The parameter statement declares a parameter. If a subscript domain is not specified, the parameter is a simple (scalar) parameter, otherwise it is a *n*-dimensional array.

The type attributes **integer**, **binary**, and **symbolic** qualify the type of values that can be assigned to the parameter as shown below:

Type attribute	Assigned values
(not specified)	Any numeric values
integer	Only integer numeric values
binary	Either 0 or 1
symbolic	Any numeric and symbolic values

The `symbolic` attribute cannot be specified along with other type attributes. Being specified it must precede all other attributes.

The `condition` attribute specifies an optional condition that restricts values assigned to the parameter to satisfy that condition. This attribute has the following syntactic forms:

<code>< v</code>	check for $x < v$
<code><= v</code>	check for $x \leq v$
<code>= v, == v</code>	check for $x = v$
<code>>= v</code>	check for $x \geq v$
<code>> v</code>	check for $x \geq v$
<code><> v, != v</code>	check for $x \neq v$

where x is a value assigned to the parameter, v is the resultant value of a numeric or symbolic expression specified in the condition attribute. Arbitrary number of condition attributes can be specified for the same parameter. If a value being assigned to the parameter during model evaluation violates at least one of specified conditions, an error is raised. (Note that symbolic values are ordered lexicographically, and any numeric value precedes any symbolic value.)

The `in` attribute is similar to the condition attribute and specifies a set expression whose resultant value is a superset used to restrict numeric or symbolic values assigned to the parameter to be in that superset. Arbitrary number of the `in` attributes can be specified for the same parameter. If a value being assigned to the parameter during model evaluation is not in at least one of specified supersets, an error is raised.

The `assign` (`:=`) attribute specifies a numeric or symbolic expression used to compute a value assigned to the parameter (if it is a simple parameter) or its member (if the parameter is an array). If the `assign` attribute is specified, the parameter is *computable* and therefore needs no data to be provided in the data section. If the `assign` attribute is not specified, the parameter must be provided with data in the data section. At most one `assign` or `default` attribute can be specified for the same parameter.

The `default` attribute specifies a numeric or symbolic expression used to compute a value assigned to the parameter or its member whenever no appropriate data are available in the data section. If neither `assign` nor `default` attribute is specified, missing data will cause an error.

4.3 Variable statement

```
var name alias domain , attrib , ... , attrib ;
```

Where: `name` is a symbolic name of the variable;

`alias` is an optional string literal, which specifies an alias of the variable;

domain is an optional indexing expression, which specifies a subscript domain of the variable;

attrib, ..., *attrib* are optional attributes of the variable. (Commae preceding attributes may be omitted.)

Optional attributes:

integer restricts the variable to be integer;

binary restricts the variable to be binary;

>= *expression*
specifies an lower bound of the variable;

<= *expression*
specifies an upper bound of the variable;

= *expression*
specifies a fixed value of the variable;

Examples

```
var x >= 0;
var y{I,J};
var make{p in prd}, integer, >= commit[p], <= market[p];
var store{raw, 1..T+1} >= 0;
var z{i in I, j in J} >= i+j;
```

The variable statement declares a variable. If a subscript domain is not specified, the variable is a simple (scalar) variable, otherwise it is a n -dimensional array of elemental variables.

Elemental variable(s) associated with the model variable (if it is a simple variable) or its members (if it is an array) correspond to the variables in the LP/MIP problem formulation (see Subsection 1.1, page 6). Note that only elemental variables actually referenced in some constraints and/or objectives are included in the LP/MIP problem instance to be generated.

The type attributes **integer** and **binary** restrict the variable to be integer or binary, respectively. If no type attribute is specified, the variable is continuous. If all variables in the model are continuous, the corresponding problem is of LP class. If there is at least one integer or binary variable, the problem is of MIP class.

The lower bound (**>=**) attribute specifies a numeric expression for computing an lower bound of the variable. At most one lower bound can be specified. By default all variables (except binary ones) have no lower bound, so if a variable is required to be non-negative, its zero lower bound should be explicitly specified.

The upper bound (**<=**) attribute specifies a numeric expression for computing an upper bound of the variable. At most one upper bound attribute can be specified.

The fixed value (=) attribute specifies a numeric expression for computing a value, at which the variable is fixed. This attribute cannot be specified along with the bound attributes.

4.4 Constraint statement

```
s.t. name alias domain : expression , = expression ;
s.t. name alias domain : expression , <= expression ;
s.t. name alias domain : expression , >= expression ;
s.t. name alias domain : expression , <= expression , <= expression ;
s.t. name alias domain : expression , >= expression , >= expression ;
```

Where: *name* is a symbolic name of the constraint;

alias is an optional string literal, which specifies an alias of the constraint;

domain is an optional indexing expression, which specifies a subscript domain of the constraint;

expression is a linear expression used to compute a component of the constraint. (Commae following expressions may be omitted.)

Note: The keyword `s.t.` may be written as `subject to` or as `subj to`, or may be omitted at all.

Examples

```
s.t. r: x + y + z, >= 0, <= 1;
limit{t in 1..T}: sum{j in prd} make[j,t] <= max_prd;
subject to balance{i in raw, t in 1..T}: store[i,t+1] -
    store[i,t] - sum{j in prd} units[i,j] * make[j,t];
subject to rlim 'regular-time limit' {t in time}:
sum{p in prd} pt[p] * rprd[p,t] <= 1.3 * dpp[t] * crews[t];
```

The constraint statement declares a constraint. If a subscript domain is not specified, the constraint is a simple (scalar) constraint, otherwise it is a n -dimensional array of elemental constraints.

Elemental constraint(s) associated with the model constraint (if it is a simple constraint) or its members (if it is an array) correspond to the linear constraints in the LP/MIP problem formulation (see Subsection 1.1, page 6).

If the constraint has the form of equality or single inequality, i.e. includes two expressions, one of which follows the colon and other follows the relation sign =, <=, or >=, both expressions in the statement can be linear expressions. If the

constraint has the form of double inequality, i.e. includes three expressions, the middle expression can be a linear expression while the leftmost and rightmost ones can be only numeric expressions.

Generating the model is, roughly speaking, generating its constraints, which are always evaluated for the entire subscript domain. Evaluation of the constraints leads, in turn, to evaluation of other model objects such as sets, parameters, and variables.

Constructing an actual linear constraint included in the problem instance, which (constraint) corresponds to a particular elemental constraint, is performed as follows.

If the constraint has the form of equality or single inequality, evaluation of both linear expressions gives two resultant linear forms:

$$\begin{aligned} f &= a_1x_1 + a_2x_2 + \dots + a_nx_n + a_0, \\ g &= b_1x_1 + a_2x_2 + \dots + a_nx_n + b_0, \end{aligned}$$

where x_1, x_2, \dots, x_n are elemental variables; $a_1, a_2, \dots, a_n, b_1, b_2, \dots, b_n$ are numeric coefficients; a_0 and b_0 are constant terms. Then all linear terms of f and g are carried to the left-hand side, and the constant terms are carried to the right-hand side, that gives the final elemental constraint in the standard form:

$$(a_1 - b_1)x_1 + (a_2 - b_2)x_2 + \dots + (a_n - b_n)x_n \left\{ \begin{array}{l} = \\ \leq \\ \geq \end{array} \right\} b_0 - a_0.$$

If the constraint has the form of double inequality, evaluation of the middle linear expression gives the resultant linear form:

$$f = a_1x_1 + a_2x_2 + \dots + a_nx_n + a_0,$$

and evaluation of the leftmost and rightmost numeric expressions gives two numeric values l and u , respectively. Then the constant term of the linear form is carried to both left-hand and right-handsides that gives the final elemental constraint in the standard form:

$$l - a_0 \leq a_1x_1 + a_2x_2 + \dots + a_nx_n \leq u - a_0.$$

4.5 Objective statement

```

minimize name alias domain : expression ;
maximize name alias domain : expression ;

```

Where: *name* is a symbolic name of the objective;

alias is an optional string literal, which specifies an alias of the objective;

domain is an optional indexing expression, which specifies a subscript domain of the objective;

expression is a linear expression used to compute the linear form of the objective.

Examples

```
minimize obj: x + 1.5 * (y + z);  
maximize total_profit: sum{p in prd} profit[p] * make[p];
```

The objective statement declares an objective. If a subscript domain is not specified, the objective is a simple (scalar) objective. Otherwise it is a n -dimensional array of elemental objectives.

Elemental objective(s) associated with the model objective (if it is a simple objective) or its members (if it is an array) correspond to general linear constraints in the LP/MIP problem formulation (see Subsection 1.1, page 6). However, unlike constraints the corresponding linear forms are free (unbounded).

Constructing an actual linear constraint included in the problem instance, which (constraint) corresponds to a particular elemental constraint, is performed as follows. The linear expression specified in the objective statement is evaluated that, gives the resultant linear form:

$$f = a_1x_1 + a_2x_2 + \dots + a_nx_n + a_0,$$

where x_1, x_2, \dots, x_n are elemental variables; a_1, a_2, \dots, a_n are numeric coefficients; a_0 is the constant term. Then the linear form is used to construct the final elemental constraint in the standard form:

$$-\infty < a_1x_1 + a_2x_2 + \dots + a_nx_n + a_0 < +\infty.$$

As a rule the model description contains only one objective statement that defines the objective function used in the problem instance. However, it is allowed to declare arbitrary number of objectives, in which case the actual objective function is the first objective encountered in the model description. Other objectives are also included in the problem instance, but they do not affect the objective function.

4.6 Solve statement

```
solve ;
```

Note: The solve statement is optional and can be used only once. If no solve statement is used, one is assumed at the end of the model section.

The solve statement causes the model to be solved, that means computing numeric values of all model variables. This allows using variables in statements below the solve statement in the same way as if they were numeric parameters.

Note that the variable, constraint, and objective statements cannot be used below the solve statement, i.e. all principal components of the model must be declared above the solve statement.

4.7 Check statement

```
check domain : expression ;
```

Where: *domain* is an optional indexing expression, which specifies a subscript domain of the check statement;

expression is a logical expression which specifies the logical condition to be checked. (The colon preceding *expression* may be omitted.)

Examples

```
check: x + y <= 1 and x >= 0 and y >= 0;
check sum{i in ORIG} supply[i] = sum{j in DEST} demand[j];
check{i in I, j in 1..10}: S[i,j] in U[i] union V[j];
```

The check statement allows checking the resultant value of a logical expression specified in the statement. If the value is *false*, an error is reported.

If the subscript domain is not specified, the check is performed only once. Specifying the subscript domain allows performing multiple checks for every *n*-tuple in the domain set. In the latter case the logical expression may include dummy indices introduced in corresponding indexing expression.

4.8 Display statement

```
display domain : item , ... , item ;
```

Where: *domain* is an optional indexing expression, which specifies a subscript domain of the check statement;

item, ..., item are items to be displayed. (The colon preceding the first item may be omitted.)

Examples

```
display: 'x =', x, 'y =', y, 'z =', z;
display sqrt(x ** 2 + y ** 2 + z ** 2);
display{i in I, j in J}: i, j, a[i,j], b[i,j];
```

The display statement evaluates all items specified in the statement and writes their values to the terminal in plain text format.

If a subscript domain is not specified, items are evaluated and then displayed only once. Specifying the subscript domain causes items to be evaluated and displayed for every n -tuple in the domain set. In the latter case items may include dummy indices introduced in corresponding indexing expression.

An item to be displayed can be a model object (set, parameter, variable, constraint, objective) or an expression.

If the item is a computable object (i.e. a set or parameter provided with the assign attribute), the object is evaluated over the entire domain and then its content (i.e. the content of the object array) is displayed. Otherwise, if the item is not a computable object, only its current content (i.e. members actually generated during the model evaluation) is displayed.

If the item is an expression, the expression is evaluated and its resultant value is displayed.

4.9 Printf statement

```
printf domain : format , expression , ... , expression ;  
printf domain : format , expression , ... , expression > filename ;  
printf domain : format , expression , ... , expression >> filename ;
```

Where: *domain* is an optional indexing expression, which specifies a subscript domain of the printf statement;

format is a symbolic expression whose value specifies a format control string. (The colon preceding the format expression may be omitted.)

expression, ..., expression are zero or more expressions whose values have to be formatted and printed. Each expression must be of numeric, symbolic, or logical type.

filename is a symbolic expression whose value specifies a name of a text file, to which the output is redirected. The flag > means creating a new empty file while the flag >> means appending the output to an existing file. If no file name is specified, the output is written to the terminal.

Examples

```
printf 'Hello, world!\n';  
printf: "x = %.3f; y = %.3f; z = %.3f\n",  
  x, y, z > "result.txt";  
printf{i in I, j in J}: "flow from %s to %s is %d\n",  
  i, j, x[i,j] >> result_file & ".txt";
```

```

printf{i in I} 'total flow from %s is %g\n',
  i, sum{j in J} x[i,j];
printf{k in K} "x[%s] = " & (if x[k] < 0 then "?" else "%g"),
  k, x[k];

```

The printf statement is similar to the display statement, however, it allows formatting data to be written.

If a subscript domain is not specified, the printf statement is executed only once. Specifying a subscript domain causes executing the printf statement for every n -tuple in the domain set. In the latter case the format and expression may include dummy indices introduced in corresponding indexing expression.

The format control string is a value of the symbolic expression *format* specified in the printf statement. It is composed of zero or more directives as follows: ordinary characters (not %), which are copied unchanged to the output stream, and conversion specifications, each of which causes evaluating corresponding expression specified in the printf statement, formatting it, and writing its resultant value to the output stream.

Conversion specifications that may be used in the format control string are the following: **d**, **i**, **f**, **F**, **e**, **E**, **g**, **G**, and **s**. These specifications have the same syntax and semantics as in the C programming language.

4.10 For statement

```

for domain : statement ;
for domain : { statement ... statement } ;

```

Where: *domain* is an indexing expression which specifies a subscript domain of the for statement. (The colon following the indexing expression may be omitted.)

statement is a statement, which should be executed under control of the for statement;

statement, ..., statement is a sequence of statements (enclosed in curly braces), which should be executed under control of the for statement.

Note: Only the following statements can be used within the for statement: check, display, printf, and another for.

Examples

```

for {(i,j) in E: i != j}
{ printf "flow from %s to %s is %g\n", i, j, x[i,j];
  check x[i,j] >= 0;
}

```

```

for {i in 1..n}
{ for {j in 1..n} printf " %s", if x[i,j] then "Q" else ".";
  printf("\n");
}
for {1..72} printf("*");

```

The for statement causes a statement or a sequence of statements specified as part of the for statement to be executed for every n -tuple in the domain set. Thus, statements within the for statement may include dummy indices introduced in corresponding indexing expression.

4.11 Table statement

```

table name alias IN driver arg ... arg :
    set <- [ fld , ... , fld ] , par ~ fld , ... , par ~ fld ;

table name alias domain OUT driver arg ... arg :
    expr ~ fld , ... , expr ~ fld ;

```

Where:

- name* is a symbolic name of the table;
- alias* is an optional string literal, which specifies an alias of the table;
- domain* is an indexing expression, which specifies a subscript domain of the (output) table;
- IN means reading data from the input table;
- OUT means writing data to the output table;
- driver* is a symbolic expression, which specifies the driver used to access the table (for details see Section C, page 60);
- arg* is an optional symbolic expression, which is an argument passed to the table driver. This symbolic expression must not include dummy indices specified in the domain;
- set* is the name of an optional simple set called *control set*. It can be omitted along with the delimiter <-;
- fld* is a field name. Within square brackets at least one field should be specified. The field name following a parameter name or expression is optional and can be omitted along with the delimiter ~, in which case the name of corresponding model object is used as the field name;
- par* is a symbolic name of a model parameter;
- expr* is a numeric or symbolic expression.

Examples

```
table data IN "CSV" "data.csv":  
  S <- [FROM,TO], d~DISTANCE, c~COST;  
table result{(f,t) in S} OUT "CSV" "result.csv":  
  f~FROM, t~TO, x[f,t]~FLOW;
```

The table statement allows reading data from a table into model objects such as sets and (non-scalar) parameters as well as writing data from the model to a table.

4.11.1 Table structure

A *data table* is an (unordered) set of *records*, where each record consists of the same number of *fields*, and each field is provided with a unique symbolic name called the *field name*. For example:

	First field	Second field	. . .	Last field
Table header	→ FROM	TO	DISTANCE	COST
First record	→ Seattle	New-York	2.5	0.12
Second record	→ Seattle	Chicago	1.7	0.08
		Seattle	1.8	0.09
. . .		San-Diego	2.5	0.15
		San-Diego	1.8	0.10
Last record	→ San-Diego	Topeka	1.4	0.07

4.11.2 Reading data from input table

The input table statement causes reading data from the specified table record by record.

Once a next record has been read, numeric or symbolic values of fields, whose names are enclosed in square brackets in the table statement, are gathered into n -tuple, and if the control set is specified in the table statement, this n -tuple is added to it. Besides, a numeric or symbolic value of each field associated with a model parameter is assigned to the parameter member identified by subscripts, which are components of the n -tuple just read.

For example, the following input table statement:

```
table data IN "...": S <- [FROM,TO], d~DISTANCE, c~COST;
```

causes reading values of four fields named FROM, TO, DISTANCE, and COST from each record of the specified table. Values of fields FROM and TO give a pair (f, t) , which is added to the control set S. The value of field DISTANCE is assigned to parameter member $d[f, t]$, and the value of field COST is assigned to parameter member $c[f, t]$.

Note that the input table may contain extra fields whose names are not specified in the table statement, in which case values of these fields on reading the table are ignored.

4.11.3 Writing data to output table

The output table statement causes writing data to the specified table. Note that some drivers (namely, CSV and xBASE) destroy the output table before writing data, i.e. delete all its existing records.

Each n -tuple in the specified domain set generates one record written to the output table. Values of fields are numeric or symbolic values of corresponding expressions specified in the table statement. These expressions are evaluated for each n -tuple in the domain set and, thus, may include dummy indices introduced in the corresponding indexing expression.

For example, the following output table statement:

```
table result{(f,t) in S} OUT "...": f~FROM, t~TO, x[f,t]~FLOW;
```

causes writing records, by one record for each pair (f, t) in set S , to the output table, where each record consists of three fields named `FROM`, `TO`, and `FLOW`. The values written to fields `FROM` and `TO` are current values of dummy indices `f` and `t`, and the value written to field `FLOW` is a value of member $x[f, t]$ of corresponding subscripted parameter or variable.

5 Model data

Model data include elemental sets, which are “values” of model sets, and numeric and symbolic values of model parameters.

In MathProg there are two different ways to saturate model sets and parameters with data. One way is simply providing necessary data using the assign attribute. However, in many cases it is more practical to separate the model itself and particular data needed for the model. For the latter reason in MathProg there is another way, when the model description is divided into two parts: model section and data section.

A *model section* is a main part of the model description that contains declarations of all model objects and is common for all problems based on that model.

A *data section* is an optional part of the model description that contains model data specific for a particular problem.

In MathProg model and data sections can be placed either in one text file or in two separate text files.

1. If both model and data sections are placed in one file, the file is composed as follows:

```
statement;  
statement;  
.  
.  
.  
statement;  
data;  
data block;  
data block;  
.  
.  
.  
data block;  
end;
```

2. If the model and data sections are placed in two separate files, the files are composed as follows:

```
statement;  
statement;  
.  
.  
.  
statement;  
end;
```

Model file

```
data;  
data block;  
data block;  
.  
.  
.  
data block;  
end;
```

Data file

Note: If the data section is placed in a separate file, the keyword **data** is optional and may be omitted along with the semicolon that follows it.

5.1 Coding data section

The *data section* is a sequence of data blocks in various formats, which are discussed in following subsections. The order, in which data blocks follow in the data section, may be arbitrary, not necessarily the same, in which corresponding model objects follow in the model section.

The rules of coding the data section are commonly the same as the rules of coding the model description (see Subsection 2, page 9), i.e. data blocks are composed from basic lexical units such as symbolic names, numeric and string literals, keywords, delimiters, and comments. However, for the sake of convenience and improving readability there is one deviation from the common rule: if a string literal consists of only alphanumeric characters (including the underscore character), the signs + and -, and/or the decimal point, it may be coded without bordering by (single or double) quotes.

All numeric and symbolic material provided in the data section is coded in the form of numbers and symbols, i.e. unlike the model section no expressions are allowed in the data section. Nevertheless, the signs + and - can precede numeric literals to allow coding signed numeric quantities, in which case there must be no white-space characters between the sign and following numeric literal (if there is at least one white-space, the sign and following numeric literal are recognized as two different lexical units).

5.2 Set data block

```
set name , record , ... , record ;  
set name [ symbol , ... , symbol ] , record , ... , record ;
```

Where: *name* is a symbolic name of the set;

symbol, ..., *symbol* are subscripts, which specify a particular member of the set (if the set is an array, i.e. a set of sets);

record, ..., *record* are data records.

Note: Commas preceding data records may be omitted.

Data records:

`:=` is a non-significant data record, which may be used freely to improve readability;

`(slice)` specifies a slice;

`simple-data` specifies set data in the simple format;

`: matrix-data` specifies set data in the matrix format;

(**tr**) : *matrix-data*
 specifies set data in the transposed matrix format. (In this case the colon following the keyword (**tr**) may be omitted.)

Examples

```

set month := Jan Feb Mar Apr May Jun;
set month "Jan", "Feb", "Mar", "Apr", "May", "Jun";
set A[3,Mar] := (1,2) (2,3) (4,2) (3,1) (2,2) (4,4) (3,4);
set A[3,'Mar'] := 1 2 2 3 4 2 3 1 2 2 4 4 2 4;
set A[3,'Mar'] : 1 2 3 4 :=
    1 - + - -
    2 - + + -
    3 + - - +
    4 - + - + ;

set B := (1,2,3) (1,3,2) (2,3,1) (2,1,3) (1,2,2) (1,1,1) (2,1,1);
set B := (*,*,*) 1 2 3, 1 3 2, 2 3 1, 2 1 3, 1 2 2, 1 1 1, 2 1 1;
set B := (1,*,2) 3 2 (2,*,1) 3 1 (1,2,3) (2,1,3) (1,1,1);
set B := (1,*,*) : 1 2 3 :=
    1 + - -
    2 - + +
    3 - + -
    (2,*,*) : 1 2 3 :=
    1 + - +
    2 - - -
    3 + - - ;

```

(In these examples `month` is a simple set of singlets, `A` is a 2-dimensional array of doublets, and `B` is a simple set of triplets. Data blocks for the same set are equivalent in the sense that they specify the same data in different formats.)

The *set data block* is used to specify a complete elemental set, which is assigned to a set (if it is a simple set) or one of its members (if the set is an array of sets).²

Data blocks can be specified only for non-computable sets, i.e. for sets, which have no assign (`:=`) attribute in the corresponding set statements.

If the set is a simple set, only its symbolic name should be specified in the header of the data block. Otherwise, if the set is a n -dimensional array, its symbolic name should be provided with a complete list of subscripts separated by commae and enclosed in square brackets to specify a particular member of the set array. The number of subscripts must be the same as the dimension of the set array, where each subscript must be a number or symbol.

An elemental set defined in the set data block is coded as a sequence of data records described below.³

²There is another way to specify data for a simple set along with data for parameters. This feature is discussed in the next subsection.

³*Data record* is simply a technical term. It does not mean that data records have any special formatting.

5.2.1 Assign data record

The *assign* ($:=$) *data record* is a non-significant element. It may be used for improving readability of data blocks.

5.2.2 Slice data record

The *slice data record* is a control record, which specifies a *slice* of the elemental set defined in the data block. It has the following syntactic form:

$$(s_1 , s_2 , \dots , s_n)$$

where s_1, s_2, \dots, s_n are components of the slice.

Each component of the slice can be a number or symbol or the asterisk (*). The number of components in the slice must be the same as the dimension of n -tuples in the elemental set to be defined. For instance, if the elemental set contains 4-tuples (quadruplets), the slice must have four components. The number of asterisks in the slice is called the *slice dimension*.

The effect of using slices is the following. If a m -dimensional slice (i.e. a slice having m asterisks) is specified in the data block, all subsequent data records must specify tuples of the dimension m . Whenever a m -tuple is encountered, each asterisk in the slice is replaced by corresponding components of the m -tuple that gives the resultant n -tuple, which is included in the elemental set to be defined. For example, if the slice $(a, *, 1, 2, *)$ is in effect, and 2-tuple $(3, b)$ is encountered in a subsequent data record, the resultant 5-tuple included in the elemental set is $(a, 3, 1, 2, b)$.

The slice having no asterisks itself defines a complete n -tuple, which is included in the elemental set.

Being once specified the slice effects until either a new slice or the end of data block is encountered. Note that if no slice is specified in the data block, one, components of which are all asterisks, is assumed.

5.2.3 Simple data record

The *simple data record* defines one n -tuple in a simple format and has the following syntactic form:

$$t_1 , t_2 , \dots , t_n$$

where t_1, t_2, \dots, t_n are components of the n -tuple. Each component can be a number or symbol. Commas between components are optional and may be omitted.

5.2.4 Matrix data record

The *matrix data record* defines several 2-tuples (doublets) in a matrix format and has the following syntactic form:

```

      :   c1   c2   ...   cn   :=
r1  a11  a12  ...  a1n
r2  a21  a22  ...  a2n
      . . . . .
rm  am1  am2  ...  amn

```

where r_1, r_2, \dots, r_m are numbers and/or symbols corresponding to rows of the matrix; c_1, c_2, \dots, c_n are numbers and/or symbols corresponding to columns of the matrix, $a_{11}, a_{12}, \dots, a_{mn}$ are matrix elements, which can be either + or -. (In this data record the delimiter : preceding the column list and the delimiter := following the column list cannot be omitted.)

Each element a_{ij} of the matrix data block (where $1 \leq i \leq m, 1 \leq j \leq n$) corresponds to 2-tuple (r_i, c_j) . If a_{ij} is the plus sign (+), that 2-tuple (or a longer n -tuple, if a slice is used) is included in the elemental set. Otherwise, if a_{ij} is the minus sign (-), that 2-tuple is not included in the elemental set.

Since the matrix data record defines 2-tuples, either the elemental set must consist of 2-tuples or the slice currently used must be 2-dimensional.

5.2.5 Transposed matrix data record

The *transposed matrix data record* has the following syntactic form:

```

(tr) :   c1   c2   ...   cn   :=
      r1   a11  a12  ...  a1n
      r2   a21  a22  ...  a2n
      . . . . .
      rm   am1  am2  ...  amn

```

(In this case the delimiter : following the keyword (tr) is optional and may be omitted.)

This data record is completely analogous to the matrix data record (see above) with only exception that in this case each element a_{ij} of the matrix corresponds to 2-tuple (c_j, r_i) rather than (r_i, c_j) .

Being once specified the (tr) indicator affects all subsequent data records until either a slice or the end of data block is encountered.

5.3 Parameter data block

```

param name , record , ... , record ;
param name default value , record , ... , record ;
param : tabbing-data ;
param default value : tabbing-data ;

```

Where: *name* is a symbolic name of the parameter;
 value is an optional default value of the parameter;
 record, . . . , *record* are data records;
 tabbing-data specifies parameter data in the tabbing format.

Note: Commae preceding data records may be omitted.

Data records:

:= is a non-significant data record, which may be used freely to improve readability;

[*slice*] specifies a slice;

plain-data specifies parameter data in the plain format;

: *tabular-data*
 specifies parameter data in the tabular format;

(*tr*) : *tabular-data*
 specifies set data in the transposed tabular format. (In this case the colon following the keyword (*tr*) may be omitted.)

Examples

```
param T := 4;
param month := 1 'Jan' 2 'Feb' 3 'Mar' 4 'Apr' 5 'May';
param month := [1] Jan, [2] Feb, [3] Mar, [4] Apr, [5] May;
param day := [Sun] 0, [Mon] 1, [Tue] 2, [Wed] 3, [Thu] 4,
             [Fri] 5, [Sat] 6;
param init_stock := iron 7.32 nickel 35.8;
param init_stock [*] iron 7.32, nickel 35.8;
param cost [iron] .025 [nickel] .03;
param value := iron -.1, nickel .02;
param      : init_stock cost value :=
    iron      7.32   .025  -.1
    nickel   35.8   .03   .02 ;
param : raw : init_stock cost value :=
    iron      7.32   .025  -.1
    nickel   35.8   .03   .02 ;
param demand default 0 (tr)
      : FRA  DET  LAN  WIN  STL  FRE  LAF :=
bands 300   .   100  75   .   225  250
coils  500  750  400  250  .   850  500
plate  100   .   .    50  200  .   250 ;
```

```

param trans_cost :=
  [*,*,bands]:  FRA  DET  LAN  WIN  STL  FRE  LAF :=
                GARY   30   10   8   10   11   71   6
                CLEV   22   7   10   7   21   82  13
                PITT   19  11  12   10  25   83  15
  [*,*,coils]:  FRA  DET  LAN  WIN  STL  FRE  LAF :=
                GARY   39  14  11  14  16   82   8
                CLEV   27   9  12   9  26   95  17
                PITT   24  14  17  13  28   99  20
  [*,*,plate]:  FRA  DET  LAN  WIN  STL  FRE  LAF :=
                GARY   41  15  12  16  17   86   8
                CLEV   29   9  13   9  28   99  18
                PITT   26  14  17  13  31  104  20 ;

```

The *parameter data block* is used to specify complete data for a parameter (or parameters, if data are specified in the tabbing format).

Data blocks can be specified only for non-computable parameters, i.e. for parameters, which have no assign (`:=`) attribute in the corresponding parameter statements.

Data defined in the parameter data block are coded as a sequence of data records described below. Additionally the data block can be provided with the optional `default` attribute, which specifies a default numeric or symbolic value of the parameter (parameters). This default value is assigned to the parameter or its members, if no appropriate value is defined in the parameter data block. The `default` attribute cannot be used, if it is already specified in the corresponding parameter statement.

5.3.1 Assign data record

The *assign (`:=`) data record* is a non-significant element. It may be used for improving readability of data blocks.

5.3.2 Slice data record

The *slice data record* is a control record, which specifies a *slice* of the parameter array. It has the following syntactic form:

$$[s_1 , s_2 , \dots , s_n]$$

where s_1, s_2, \dots, s_n are components of the slice.

Each component of the slice can be a number or symbol or the asterisk (*). The number of components in the slice must be the same as the dimension of the parameter. For instance, if the parameter is a 4-dimensional array, the slice must have four components. The number of asterisks in the slice is called the *slice dimension*.

The effect of using slices is the following. If a m -dimensional slice (i.e. a slice having m asterisks) is specified in the data block, all subsequent data records

must specify subscripts of the parameter members as if the parameter were m -dimensional, not n -dimensional.

Whenever m subscripts are encountered, each asterisk in the slice is replaced by corresponding subscript that gives n subscripts, which define the actual parameter member. For example, if the slice $[a, *, 1, 2, *]$ is in effect, and subscripts 3 and b are encountered in a subsequent data record, the complete subscript list used to choose a parameter member is $[a, 3, 1, 2, b]$.

It is allowed to specify a slice having no asterisks. Such slice itself defines a complete subscript list, in which case the next data record should define only a single value of corresponding parameter member.

Being once specified the slice effects until either a new slice or the end of data block is encountered. Note that if no slice is specified in the data block, one, components of which are all asterisks, is assumed.

5.3.3 Plain data record

The *plain data record* defines a subscript list and a single value in the plain format. This record has the following syntactic form:

$$t_1, t_2, \dots, t_n, v$$

where t_1, t_2, \dots, t_n are subscripts, and v is a value. Each subscript as well as the value can be a number or symbol. Commae following subscripts are optional and may be omitted.

In case of 0-dimensional parameter or slice the plain data record has no subscripts and consists of a single value only.

5.3.4 Tabular data record

The *tabular data record* defines several values, where each value is provided with two subscripts. This record has the following syntactic form:

$$\begin{array}{cccccc}
 : & c_1 & c_2 & \dots & c_n & := \\
 r_1 & a_{11} & a_{12} & \dots & a_{1n} & \\
 r_2 & a_{21} & a_{22} & \dots & a_{2n} & \\
 & \cdot & \cdot & \cdot & \cdot & \cdot \\
 r_m & a_{m1} & a_{m2} & \dots & a_{mn} &
 \end{array}$$

where r_1, r_2, \dots, r_m are numbers and/or symbols corresponding to rows of the table; c_1, c_2, \dots, c_n are numbers and/or symbols corresponding to columns of the table, $a_{11}, a_{12}, \dots, a_{mn}$ are table elements. Each element can be a number or symbol or the single decimal point (\cdot). (In this data record the delimiter $:$ preceding the column list and the delimiter $:=$ following the column list cannot be omitted.)

Each element a_{ij} of the tabular data block ($1 \leq i \leq m, 1 \leq j \leq n$) defines two subscripts, where the first subscript is r_i , and the second one is c_j . These subscripts are used in conjunction with the current slice to form the complete subscript list that identifies a particular member of the parameter array. If a_{ij} is

a number or symbol, this value is assigned to the parameter member. However, if a_{ij} is the single decimal point, the member is assigned a default value specified either in the parameter data block or in the parameter statement, or, if no default value is specified, the member remains undefined.

Since the tabular data record provides two subscripts for each value, either the parameter or the slice currently used must be 2-dimensional.

5.3.5 Transposed tabular data record

The *transposed tabular data record* has the following syntactic form:

```
(tr) :  c1  c2  ...  cn  :=
        r1  a11 a12  ...  a1n
        r2  a21 a22  ...  a2n
        . . . . .
        rm  am1 am2  ...  amn
```

(In this case the delimiter `:` following the keyword `(tr)` is optional and may be omitted.)

This data record is completely analogous to the tabular data record (see above) with only exception that the first subscript defined by element a_{ij} is c_j while the second one is r_i .

Being once specified the `(tr)` indicator affects all subsequent data records until either a slice or the end of data block is encountered.

5.3.6 Tabbing data format

The parameter data block in the *tabbing format* has the following syntactic form:

```
param default value : s : p1 , p2 , ... , pk :=
    t11 , t12 , ... , t1n , a11 , a12 , ... , a1k
    t21 , t22 , ... , t2n , a21 , a22 , ... , a2k
    . . . . .
    tm1 , tm2 , ... , tmn , am1 , am2 , ... , amk ;
```

Notes:

1. The keyword `default` may be omitted along with a value following it.
2. Symbolic name `s` may be omitted along with the colon following it.
3. All comae are optional and may be omitted.

The data block in the tabbing format shown above is exactly equivalent to the following data blocks for $j = 1, 2, \dots, k$:

```
set s := (t11, t12, ..., t1n) (t21, t22, ..., t2n) ... (tm1, tm2, ..., tmn) ;
param pj default value :=
[t11, t12, ..., t1n] a1j [t21, t22, ..., t2n] a2j ... [tm1, tm2, ..., tmn] amj ;
```

A Using suffixes

Suffixes can be used to retrieve additional values associated with model variables, constraints, and objectives.

A *suffix* consists of a period (.) followed by a non-reserved keyword. For example, if x is a two-dimensional variable, $x[i, j].lb$ is a numeric value equal to the lower bound of elemental variable $x[i, j]$, which (value) can be used everywhere in expressions like a numeric parameter.

For model variables suffixes have the following meaning:

<code>.lb</code>	lower bound
<code>.ub</code>	upper bound
<code>.status</code>	status in the solution: 0 — undefined 1 — basic 2 — non-basic on lower bound 3 — non-basic on upper bound 4 — non-basic free (unbounded) variable 5 — non-basic fixed variable
<code>.val</code>	primal value in the solution
<code>.dual</code>	dual value (reduced cost) in the solution

For model constraints and objectives suffixes have the following meaning:

<code>.lb</code>	lower bound of the linear form
<code>.ub</code>	upper bound of the linear form
<code>.status</code>	status in the solution: 0 — undefined 1 — non-active 2 — active on lower bound 3 — active on upper bound 4 — active free (unbounded) row 5 — active equality constraint
<code>.val</code>	primal value of the linear form in the solution
<code>.dual</code>	dual value (reduced cost) of the linear form in the solution

Note that suffixes `.status`, `.val`, and `.dual` can be used only below the solve statement.

B Date and time functions

by Andrew Makhorin <mao@gnu.org>
and Heinrich Schuchardt <heinrich.schuchardt@gmx.de>

B.1 Obtaining current calendar time

To obtain the current calendar time in MathProg there exists the function `gmtime`. It has no arguments and returns the number of seconds elapsed since 00:00:00 on January 1, 1970, Coordinated Universal Time (UTC). For example:

```
param utc := gmtime();
```

MathProg has no function to convert UTC time returned by the function `gmtime` to *local* calendar times. Thus, if you need to determine the current local calendar time, you have to add to the UTC time returned the time offset from UTC expressed in seconds. For example, the time in Berlin during the winter is one hour ahead of UTC that corresponds to the time offset +1 hour = +3600 secs, so the current winter calendar time in Berlin may be determined as follows:

```
param now := gmtime() + 3600;
```

Similarly, the summer time in Chicago (Central Daylight Time) is five hours behind UTC, so the corresponding current local calendar time may be determined as follows:

```
param now := gmtime() - 5 * 3600;
```

Note that the value returned by `gmtime` is volatile, i.e. being called several times this function may return different values.

B.2 Converting character string to calendar time

The function `str2time(s, f)` converts a character string (timestamp) specified by its first argument *s*, which must be a symbolic expression, to the calendar time suitable for arithmetic calculations. The conversion is controlled by the specified format string *f* (the second argument), which also must be a symbolic expression.

The result of conversion returned by `str2time` has the same meaning as values returned by the function `gmtime` (see Subsection B.1, page 55). Note that `str2time` does **not** correct the calendar time returned for the local timezone, i.e. being applied to 00:00:00 on January 1, 1970 it always returns 0.

For example, the model statements:

```
param s, symbolic, := "07/14/98 13:47";  
param t := str2time(s, "%m/%d/%y %H:%M");  
display t;
```

produce the following printout:

```
t = 900424020
```

where the calendar time printed corresponds to 13:47:00 on July 14, 1998.

The format string passed to the function `str2time` consists of conversion specifiers and ordinary characters. Each conversion specifier begins with a percent (%) character followed by a letter.

The following conversion specifiers may be used in the format string:

- `%b` The abbreviated month name (case insensitive). At least three first letters of the month name must appear in the input string.
- `%d` The day of the month as a decimal number (range 1 to 31). Leading zero is permitted, but not required.
- `%h` The same as `%b`.
- `%H` The hour as a decimal number, using a 24-hour clock (range 0 to 23). Leading zero is permitted, but not required.
- `%m` The month as a decimal number (range 1 to 12). Leading zero is permitted, but not required.
- `%M` The minute as a decimal number (range 0 to 59). Leading zero is permitted, but not required.
- `%S` The second as a decimal number (range 0 to 60). Leading zero is permitted, but not required.
- `%y` The year without a century as a decimal number (range 0 to 99). Leading zero is permitted, but not required. Input values in the range 0 to 68 are considered as the years 2000 to 2068 while the values 69 to 99 as the years 1969 to 1999.
- `%z` The offset from GMT in ISO 8601 format.
- `%%` A literal % character.

All other (ordinary) characters in the format string must have a matching character in the input string to be converted. Exceptions are spaces in the input string which can match zero or more space characters in the format string.

If some date and/or time component(s) are missing in the format and, therefore, in the input string, the function `str2time` uses their default values corresponding to 00:00:00 on January 1, 1970, that is, the default value of the year is 1970, the default value of the month is January, etc.

The function `str2time` is applicable to all calendar times in the range 00:00:00 on January 1, 0001 to 23:59:59 on December 31, 4000 of the Gregorian calendar.

B.3 Converting calendar time to character string

The function `time2str(t, f)` converts the calendar time specified by its first argument t , which must be a numeric expression, to a character string (symbolic value). The conversion is controlled by the specified format string f (the second argument), which must be a symbolic expression.

The calendar time passed to `time2str` has the same meaning as values returned by the function `gmtime` (see Subsection B.1, page 55). Note that `time2str` does *not* correct the specified calendar time for the local timezone, i.e. the calendar time 0 always corresponds to 00:00:00 on January 1, 1970.

For example, the model statements:

```
param s, symbolic, := time2str(gmtime(), "%FT%TZ");
display s;
```

may produce the following printout:

```
s = '2008-12-04T00:23:45Z'
```

which is a timestamp in the ISO format.

The format string passed to the function `time2str` consists of conversion specifiers and ordinary characters. Each conversion specifier begins with a percent (%) character followed by a letter.

The following conversion specifiers may be used in the format string:

- `%a` The abbreviated (2-character) weekday name.
- `%A` The full weekday name.
- `%b` The abbreviated (3-character) month name.
- `%B` The full month name.
- `%C` The century of the year, that is the greatest integer not greater than the year divided by 100.
- `%d` The day of the month as a decimal number (range 01 to 31).
- `%D` The date using the format `%m/%d/%y`.
- `%e` The day of the month like with `%d`, but padded with blank rather than zero.
- `%F` The date using the format `%Y-%m-%d`.
- `%g` The year corresponding to the ISO week number, but without the century (range 00 to 99). This has the same format and value as `%y`, except that if the ISO week number (see `%V`) belongs to the previous or next year, that year is used instead.
- `%G` The year corresponding to the ISO week number. This has the same format and value as `%Y`, except that if the ISO week number (see `%V`) belongs to the previous or next year, that year is used instead.
- `%h` The same as `%b`.
- `%H` The hour as a decimal number, using a 24-hour clock (range 00 to 23).
- `%I` The hour as a decimal number, using a 12-hour clock (range 01 to 12).

%j	The day of the year as a decimal number (range 001 to 366).
%k	The hour as a decimal number, using a 24-hour clock like %H, but padded with blank rather than zero.
%l	The hour as a decimal number, using a 12-hour clock like %I, but padded with blank rather than zero.
%m	The month as a decimal number (range 01 to 12).
%M	The minute as a decimal number (range 00 to 59).
%p	Either AM or PM, according to the given time value. Midnight is treated as AM and noon as PM.
%P	Either am or pm, according to the given time value. Midnight is treated as am and noon as pm.
%R	The hour and minute in decimal numbers using the format %H:%M.
%S	The second as a decimal number (range 00 to 59).
%T	The time of day in decimal numbers using the format %H:%M:%S.
%u	The day of the week as a decimal number (range 1 to 7), Monday being 1.
%U	The week number of the current year as a decimal number (range 00 to 53), starting with the first Sunday as the first day of the first week. Days preceding the first Sunday in the year are considered to be in week 00.
%V	The ISO week number as a decimal number (range 01 to 53). ISO weeks start with Monday and end with Sunday. Week 01 of a year is the first week which has the majority of its days in that year; this is equivalent to the week containing January 4. Week 01 of a year can contain days from the previous year. The week before week 01 of a year is the last week (52 or 53) of the previous year even if it contains days from the new year. In other word, if 1 January is Monday, Tuesday, Wednesday or Thursday, it is in week 01; if 1 January is Friday, Saturday or Sunday, it is in week 52 or 53 of the previous year.
%w	The day of the week as a decimal number (range 0 to 6), Sunday being 0.
%W	The week number of the current year as a decimal number (range 00 to 53), starting with the first Monday as the first day of the first week. Days preceding the first Monday in the year are considered to be in week 00.
%y	The year without a century as a decimal number (range 00 to 99), that is the year modulo 100.

`%Y` The year as a decimal number, using the Gregorian calendar.

`%%` A literal `%` character.

All other (ordinary) characters in the format string are simply copied to the resultant string.

The first argument (calendar time) passed to the function `time2str` must be in the range from `-62135596800` to `+64092211199` that corresponds to the period from 00:00:00 on January 1, 0001 to 23:59:59 on December 31, 4000 of the Gregorian calendar.

C Table drivers

by Andrew Makhorin <mao@gnu.org>
and Heinrich Schuchardt <heinrich.schuchardt@gmx.de>

The *table driver* is a program module which provides transmitting data between MathProg model objects and data tables.

Currently the GLPK package has four table drivers:

- built-in CSV table driver;
- built-in xBASE table driver;
- ODBC table driver;
- MySQL table driver.

C.1 CSV table driver

The CSV table driver assumes that the data table is represented in the form of a plain text file in the CSV (comma-separated values) file format as described below.

To choose the CSV table driver its name in the table statement should be specified as "CSV", and the only argument should specify the name of a plain text file containing the table. For example:

```
table data IN "CSV" "data.csv": ... ;
```

The filename suffix may be arbitrary, however, it is recommended to use the suffix '.csv'.

On reading input tables the CSV table driver provides an implicit field named RECNO, which contains the current record number. This field can be specified in the input table statement as if there were the actual field having the name RECNO in the CSV file. For example:

```
table list IN "CSV" "list.csv": num <- [RECNO], ... ;
```

CSV format⁴

The CSV (comma-separated values) format is a plain text file format defined as follows.

1. Each record is located on a separate line, delimited by a line break. For example:

```
aaa,bbb,ccc\nxxx,yyy,zzz\n
```

where \n means the control character LF (0x0A).

⁴This material is based on the RFC document 4180.

2. The last record in the file may or may not have an ending line break. For example:

```
aaa,bbb,ccc\n
xxx,yyy,zzz
```

3. There should be a header line appearing as the first line of the file in the same format as normal record lines. This header should contain names corresponding to the fields in the file. The number of field names in the header line should be the same as the number of fields in the records of the file. For example:

```
name1,name2,name3\n
aaa,bbb,ccc\n
xxx,yyy,zzz\n
```

4. Within the header and each record there may be one or more fields separated by commas. Each line should contain the same number of fields throughout the file. Spaces are considered as part of a field and therefore not ignored. The last field in the record should not be followed by a comma. For example:

```
aaa,bbb,ccc\n
```

5. Fields may or may not be enclosed in double quotes. For example:

```
"aaa","bbb","ccc"\n
zzz,yyy,xxx\n
```

6. If a field is enclosed in double quotes, each double quote which is part of the field should be coded twice. For example:

```
"aaa","b""bb","ccc"\n
```

Example

```
FROM,TO,DISTANCE,COST
Seattle,New-York,2.5,0.12
Seattle,Chicago,1.7,0.08
Seattle,Topeka,1.8,0.09
San-Diego,New-York,2.5,0.15
San-Diego,Chicago,1.8,0.10
San-Diego,Topeka,1.4,0.07
```

C.2 xBASE table driver

The xBASE table driver assumes that the data table is stored in the .dbf file format.

To choose the xBASE table driver its name in the table statement should be specified as "xBASE", and the first argument should specify the name of a .dbf file containing the table. For the output table there should be the second

argument defining the table format in the form "FF...F", where F is either C(*n*), which specifies a character field of length *n*, or N(*n*[,*p*]), which specifies a numeric field of length *n* and precision *p* (by default *p* is 0).

The following is a simple example which illustrates creating and reading a .dbf file:

```
table tab1{i in 1..10} OUT "xBASE" "foo.dbf"
  "N(5)N(10,4)C(1)C(10)": 2*i+1 ~ B, Uniform(-20,+20) ~ A,
  "?" ~ F00, "[" & i & "]" ~ C;
set S, dimen 4;
table tab2 IN "xBASE" "foo.dbf": S <- [B, C, RECNO, A];
display S;
end;
```

C.3 ODBC table driver

The ODBC table driver allows connecting to SQL databases using an implementation of the ODBC interface based on the Call Level Interface (CLI).⁵

Debian GNU/Linux. Under Debian GNU/Linux the ODBC table driver uses the iODBC package,⁶ which should be installed before building the GLPK package. The installation can be effected with the following command:

```
sudo apt-get install libiodbc2-dev
```

Note that on configuring the GLPK package to enable using the iODBC library the option '`--enable-odbc`' should be passed to the configure script.

The individual databases must be entered for systemwide usage in `/etc/odbc.ini` and `/etc/odbcinst.ini`. Database connections to be used by a single user are specified by files in the home directory (`.odbc.ini` and `.odbcinst.ini`).

Microsoft Windows. Under Microsoft Windows the ODBC table driver uses the Microsoft ODBC library. To enable this feature the symbol:

```
#define ODBC_DLNAME "odbc32.dll"
```

should be defined in the GLPK configuration file '`config.h`'.

Data sources can be created via the Administrative Tools from the Control Panel.

To choose the ODBC table driver its name in the table statement should be specified as '`ODBC`' or '`iODBC`'.

The argument list is specified as follows.

⁵The corresponding software standard is defined in ISO/IEC 9075-3:2003.

⁶See <http://www.iodbc.org/>.

The first argument is the connection string passed to the ODBC library, for example:

```
'DSN=glpk;UID=user;PWD=password', or  
'DRIVER=MySQL;DATABASE=glpkdb;UID=user;PWD=password'.
```

Different parts of the string are separated by semicolons. Each part consists of a pair *fieldname* and *value* separated by the equal sign. Allowable fieldnames depend on the ODBC library. Typically the following fieldnames are allowed:

```
DATABASE  database;  
DRIVER    ODBC driver;  
DSN       name of a data source;  
FILEDSN   name of a file data source;  
PWD       user password;  
SERVER    database;  
UID       user name.
```

The second argument and all following are considered to be SQL statements. SQL statements may be spread over multiple arguments. If the last character of an argument is a semicolon this indicates the end of a SQL statement.

The arguments of a SQL statement are concatenated separated by space. The eventual trailing semicolon will be removed.

All but the last SQL statement will be executed directly.

For IN-table the last SQL statement can be a SELECT command starting with the capitalized letters 'SELECT '. If the string does not start with 'SELECT ' it is considered to be a table name and a SELECT statement is automatically generated.

For OUT-table the last SQL statement can contain one or multiple question marks. If it contains a question mark it is considered a template for the write routine. Otherwise the string is considered a table name and an INSERT template is automatically generated.

The writing routine uses the template with the question marks and replaces the first question mark by the first output parameter, the second question mark by the second output parameter and so forth. Then the SQL command is issued.

The following is an example of the output table statement:

```
table ta { 1 in LOCATIONS } OUT  
  'ODBC'  
  'DSN=glpkdb;UID=glpkuser;PWD=glpkpassword'  
  'DROP TABLE IF EXISTS result;'  
  'CREATE TABLE result ( ID INT, LOC VARCHAR(255), QUAN DOUBLE );'  
  'INSERT INTO result 'VALUES ( 4, ?, ? )' :  
  1 ~ LOC, quantity[1] ~ QUAN;
```

Alternatively it could be written as follows:

```
table ta { 1 in LOCATIONS } OUT  
  'ODBC'  
  'DSN=glpkdb;UID=glpkuser;PWD=glpkpassword'  
  'DROP TABLE IF EXISTS result;'
```

```
'CREATE TABLE result ( ID INT, LOC VARCHAR(255), QUAN DOUBLE );'
'result' :
1 ~ LOC, quantity[1] ~ QUAN, 4 ~ ID;
```

Using templates with '?' supports not only INSERT, but also UPDATE, DELETE, etc. For example:

```
table ta { 1 in LOCATIONS } OUT
'ODBC'
'DSN=glpkdb;UID=glpkuser;PWD=glpkpassword'
'UPDATE result SET DATE = ' & date & ' WHERE ID = 4;'
'UPDATE result SET QUAN = ? WHERE LOC = ? AND ID = 4' :
quantity[1], 1;
```

C.4 MySQL table driver

The MySQL table driver allows connecting to MySQL databases.

Debian GNU/Linux. Under Debian GNU/Linux the MySQL table driver uses the MySQL package,⁷ which should be installed before building the GLPK package. The installation can be effected with the following command:

```
sudo apt-get install libmysqlclient15-dev
```

Note that on configuring the GLPK package to enable using the MySQL library the option '--enable-mysql' should be passed to the configure script.

Microsoft Windows. Under Microsoft Windows the MySQL table driver also uses the MySQL library. To enable this feature the symbol:

```
#define MYSQL_DLNAME "libmysql.dll"
```

should be defined in the GLPK configuration file 'config.h'.

To choose the MySQL table driver its name in the table statement should be specified as 'MySQL'.

The argument list is specified as follows.

The first argument specifies how to connect the data base in the DSN style, for example:

```
'Database=glpk;UID=glpk;PWD=gnu'
```

Different parts of the string are separated by semicolons. Each part consists of a pair *fieldname* and *value* separated by the equal sign. The following fieldnames are allowed:

```
Server   server running the database (defaulting to localhost);
Database name of the database;
UID      user name;
```

⁷For download development files see <<http://dev.mysql.com/downloads/mysql/>>.

PWD user password;

Port port used by the server (defaulting to 3306).

The second argument and all following are considered to be SQL statements. SQL statements may be spread over multiple arguments. If the last character of an argument is a semicolon this indicates the end of a SQL statement.

The arguments of a SQL statement are concatenated separated by space. The eventual trailing semicolon will be removed.

All but the last SQL statement will be executed directly.

For IN-table the last SQL statement can be a SELECT command starting with the capitalized letters 'SELECT '. If the string does not start with 'SELECT ' it is considered to be a table name and a SELECT statement is automatically generated.

For OUT-table the last SQL statement can contain one or multiple question marks. If it contains a question mark it is considered a template for the write routine. Otherwise the string is considered a table name and an INSERT template is automatically generated.

The writing routine uses the template with the question marks and replaces the first question mark by the first output parameter, the second question mark by the second output parameter and so forth. Then the SQL command is issued.

The following is an example of the output table statement:

```
table ta { 1 in LOCATIONS } OUT
  'MySQL'
  'Database=glpkdb;UID=glpkuser;PWD=glpkpassword'
  'DROP TABLE IF EXISTS result; '
  'CREATE TABLE result ( ID INT, LOC VARCHAR(255), QUAN DOUBLE );'
  'INSERT INTO result VALUES ( 4, ?, ? )' :
  1 ~ LOC, quantity[1] ~ QUAN;
```

Alternatively it could be written as follows:

```
table ta { 1 in LOCATIONS } OUT
  'MySQL'
  'Database=glpkdb;UID=glpkuser;PWD=glpkpassword'
  'DROP TABLE IF EXISTS result; '
  'CREATE TABLE result ( ID INT, LOC VARCHAR(255), QUAN DOUBLE );'
  'result' :
  1 ~ LOC, quantity[1] ~ QUAN, 4 ~ ID;
```

Using templates with '?' supports not only INSERT, but also UPDATE, DELETE, etc. For example:

```
table ta { 1 in LOCATIONS } OUT
  'MySQL'
  'Database=glpkdb;UID=glpkuser;PWD=glpkpassword'
  'UPDATE result SET DATE = ' & date & ' WHERE ID = 4; '
  'UPDATE result SET QUAN = ? WHERE LOC = ? AND ID = 4' :
  quantity[1], 1;
```

D Solving models with glpsol

The GLPK package⁸ includes the program `glpsol`, which is a stand-alone LP/MIP solver. This program can be launched from the command line or from the shell to solve models written in the GNU MathProg modeling language.

In order to tell the solver that the input file contains a model description, you need to specify the option `--model` in the command line. For example:

```
glpsol --model foo.mod
```

Sometimes it is necessary to use the data section placed in a separate file, in which case you may use the following command:

```
glpsol --model foo.mod --data foo.dat
```

Note that if the model file also contains the data section, that section is ignored.

If the model description contains some display and/or printf statements, by default the output is sent to the terminal. In order to redirect the output to a file you may use the following command:

```
glpsol --model foo.mod --display foo.out
```

If you need to look at the problem, which has been generated by the model translator, you may use the option `--wlp` as follows:

```
glpsol --model foo.mod --wlp foo.lp
```

in which case the problem data is written to file `foo.lp` in CPLEX LP format suitable for visual analysis.

Sometimes it is needed merely to check the model description not solving the generated problem instance. In this case you may specify the option `--check`, for example:

```
glpsol --check --model foo.mod --wlp foo.lp
```

In order to write a numeric solution obtained by the solver you may use the following command:

```
glpsol --model foo.mod --output foo.sol
```

in which case the solution is written to file `foo.sol` in a plain text format.

The complete list of the `glpsol` options can be found in the reference manual included in the GLPK distribution.

⁸<http://www.gnu.org/software/glpk/>

E Example model description

E.1 Model description written in MathProg

Below here is a complete example of the model description written in the GNU MathProg modeling language.

```
# A TRANSPORTATION PROBLEM
#
# This problem finds a least cost shipping schedule that meets
# requirements at markets and supplies at factories.
#
# References:
#     Dantzig G B, "Linear Programming and Extensions."
#     Princeton University Press, Princeton, New Jersey, 1963,
#     Chapter 3-3.

set I;
/* canning plants */

set J;
/* markets */

param a{i in I};
/* capacity of plant i in cases */

param b{j in J};
/* demand at market j in cases */

param d{i in I, j in J};
/* distance in thousands of miles */

param f;
/* freight in dollars per case per thousand miles */

param c{i in I, j in J} := f * d[i,j] / 1000;
/* transport cost in thousands of dollars per case */

var x{i in I, j in J} >= 0;
/* shipment quantities in cases */

minimize cost: sum{i in I, j in J} c[i,j] * x[i,j];
/* total transportation costs in thousands of dollars */

s.t. supply{i in I}: sum{j in J} x[i,j] <= a[i];
/* observe supply limit at plant i */

s.t. demand{j in J}: sum{i in I} x[i,j] >= b[j];
/* satisfy demand at market j */
```

```

data;

set I := Seattle San-Diego;

set J := New-York Chicago Topeka;

param a := Seattle      350
          San-Diego     600;

param b := New-York     325
          Chicago       300
          Topeka        275;

param d :           New-York  Chicago  Topeka :=
          Seattle    2.5       1.7     1.8
          San-Diego  2.5       1.8     1.4 ;

param f := 90;

end;

```

E.2 Generated LP problem instance

Below here is the result of the translation of the example model produced by the solver `glpsol` and written in CPLEX LP format with the option `--wlp`.

```

\* Problem: transp *\

Minimize
  cost: + 0.225 x(Seattle,New~York) + 0.153 x(Seattle,Chicago)
        + 0.162 x(Seattle,Topeka) + 0.225 x(San~Diego,New~York)
        + 0.162 x(San~Diego,Chicago) + 0.126 x(San~Diego,Topeka)

Subject To
  supply(Seattle): + x(Seattle,New~York) + x(Seattle,Chicago)
                  + x(Seattle,Topeka) <= 350
  supply(San~Diego): + x(San~Diego,New~York) + x(San~Diego,Chicago)
                   + x(San~Diego,Topeka) <= 600
  demand(New~York): + x(Seattle,New~York) + x(San~Diego,New~York) >= 325
  demand(Chicago): + x(Seattle,Chicago) + x(San~Diego,Chicago) >= 300
  demand(Topeka): + x(Seattle,Topeka) + x(San~Diego,Topeka) >= 275

End

```

E.3 Optimal LP solution

Below here is the optimal solution of the generated LP problem instance found by the solver `glpsol` and written in plain text format with the option `--output`.

Problem: transp
 Rows: 6
 Columns: 6
 Non-zeros: 18
 Status: OPTIMAL
 Objective: cost = 153.675 (MINimum)

No.	Row name	St	Activity	Lower bound	Upper bound	Marginal
1	cost	B	153.675			
2	supply[Seattle]	B	300		350	
3	supply[San-Diego]	NU	600		600	< eps
4	demand[New-York]	NL	325	325		0.225
5	demand[Chicago]	NL	300	300		0.153
6	demand[Topeka]	NL	275	275		0.126

No.	Column name	St	Activity	Lower bound	Upper bound	Marginal
1	x[Seattle,New-York]	B	0	0		
2	x[Seattle,Chicago]	B	300	0		
3	x[Seattle,Topeka]	NL	0	0		0.036
4	x[San-Diego,New-York]	B	325	0		
5	x[San-Diego,Chicago]	NL	0	0		0.009
6	x[San-Diego,Topeka]	B	275	0		

End of output

Acknowledgment

The authors would like to thank the following people, who kindly read, commented, and corrected the draft of this document:

Juan Carlos Borrás <borras@cs.helsinki.fi>

Harley Mackenzie <hjm@bigpond.com>

Robbie Morrison <robbie@actrix.co.nz>