
AIMMS Language Reference - Set Declaration

This file contains only one chapter of the book. For a free download of the complete book in pdf format, please visit www.aimms.com or order your hard-copy at www.lulu.com/aimms.

Copyright © 1993–2011 by Paragon Decision Technology B.V. All rights reserved.

Paragon Decision Technology B.V.	Paragon Decision Technology Inc.	Paragon Decision Technology Pte.
Schipholweg 1	500 108th Avenue NE	Ltd.
2034 LS Haarlem	Ste. # 1085	80 Raffles Place
The Netherlands	Bellevue, WA 98004	UOB Plaza 1, Level 36-01
Tel.: +31 23 5511512	USA	Singapore 048624
Fax: +31 23 5511517	Tel.: +1 425 458 4024	Tel.: +65 9640 4182
	Fax: +1 425 458 4025	

Email: info@aimms.com
WWW: www.aimms.com

AIMMS is a registered trademark of Paragon Decision Technology B.V. IBM ILOG CPLEX and sc CPLEX is a registered trademark of IBM Corporation. GUROBI is a registered trademark of Gurobi Optimization, Inc. KNITRO is a registered trademark of Ziena Optimization, Inc. XPRESS-MP is a registered trademark of FICO Fair Isaac Corporation. MOSEK is a registered trademark of Mosek ApS. WINDOWS and EXCEL are registered trademarks of Microsoft Corporation. $\text{T}_{\text{E}}\text{X}$, $\text{L}_{\text{A}}\text{T}_{\text{E}}\text{X}$, and $\text{A}_{\text{M}}\text{S}_{\text{L}}\text{A}_{\text{T}}\text{E}_{\text{X}}$ are trademarks of the American Mathematical Society. LUCIDA is a registered trademark of Bigelow & Holmes Inc. ACROBAT is a registered trademark of Adobe Systems Inc. Other brands and their products are trademarks of their respective holders.

Information in this document is subject to change without notice and does not represent a commitment on the part of Paragon Decision Technology B.V. The software described in this document is furnished under a license agreement and may only be used and copied in accordance with the terms of the agreement. The documentation may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from Paragon Decision Technology B.V.

Paragon Decision Technology B.V. makes no representation or warranty with respect to the adequacy of this documentation or the programs which it describes for any particular purpose or with respect to its adequacy to produce any particular result. In no event shall Paragon Decision Technology B.V., its employees, its contractors or the authors of this documentation be liable for special, direct, indirect or consequential damages, losses, costs, charges, claims, demands, or claims for lost profits, fees or expenses of any nature or kind.

In addition to the foregoing, users should recognize that all complex software systems and their documentation contain errors and omissions. The authors, Paragon Decision Technology B.V. and its employees, and its contractors shall not be responsible under any circumstances for providing information or corrections to errors and omissions discovered at any time in this book or the software it describes, whether or not they are aware of the errors or omissions. The authors, Paragon Decision Technology B.V. and its employees, and its contractors do not recommend the use of the software described in this book for applications in which errors or omissions could threaten life, injury or significant loss.

This documentation was typeset by Paragon Decision Technology B.V. using $\text{L}_{\text{A}}\text{T}_{\text{E}}\text{X}$ and the LUCIDA font family.

Part II

Non-Procedural Language Components

Chapter 3

Set Declaration

This chapter covers all aspects associated with the declaration and use of sets in AIMMS models. The main topics are indexing with sets, simple sets with strings, simple sets with integers, relations, compound sets and indexed sets.

This chapter

3.1 Sets and indices

Sets and indices give your AIMMS model dimension and depth by providing a mechanism for grouping parameters, variables, and constraints. Sets and indices are also used as driving mechanism in arithmetic operations such as summation. The use of sets for indexing expressions helps to describe large models in a concise and understandable way.

General

Consider a set of *Cities* and an identifier called *Transport* defined between several pairs of cities (i, j) , representing the amount of product transported from supply city i to destination city j . Suppose that you are interested in the quantities arriving in each city. Rather than adding many individual terms, the following mathematical notation, using sets and indices, concisely describes the desired computation of these quantities.

Example

$$(\forall j \in Cities) \quad Arrival_j = \sum_{i \in Cities} Transport_{ij}.$$

This multidimensional index notation forms the foundation of the AIMMS modeling language, and can be used in all expressions. In this example, i and j are indices that refer to individual *Cities*.

A set in AIMMS

- has either *strings* or *integers* as elements,
- is either a *simple* set, a *relation*, or a *compound* set, and
- is either *indexed* or *not indexed*.

Several types of sets

Sets can either have strings as elements (such as the set *Cities* discussed above), or have integers as elements. An example of an integer set could be a set of *Trials* represented by the numbers $1, \dots, n$. The resulting integer set can then be used to refer to the results of each single experiment.

String versus integer

A *simple* set is a one-dimensional set, such as the set *Cities* mentioned above, while a *relation* or multidimensional set is the Cartesian product of a number of simple sets or a subset thereof. An example of a relation is the set of possible *Routes* between supply and destination cities, which can be represented as a subset of the Cartesian product $Cities \times Cities$.

Simple versus relation

Sets in AIMMS are the basis for creating multidimensional identifiers in your model. Through indices into sets you have access to individual values of these identifiers for each tuple of elements. In addition, the indexing notation in AIMMS is your basic mechanism for expressing iterative operations such as repeated addition, repeated multiplication, sequential search for a maximum or minimum, etc.

Indexing as basic mechanism

As you shall see, you can use both simple and relations for indexing. If you think, for example, that it is more convenient to express your model components in terms of an index r into the relation *Routes* rather than in terms of tuples (i, j) of cities, AIMMS allows you to do so. In AIMMS, *compound sets* are a special type of relations, namely those which have an associated index.

Compound sets

Both simple and compound sets may be indexed. An indexed set is a family of sets defined for every element in the index domain of the indexed set. An example of an indexed set is the set of transport destination cities defined for each supply city. On the other hand, the set *Cities* discussed above is not an indexed set.

Indexed sets

The contents of any simple or compound set can be sorted in AIMMS. Sorting can take place either automatically or manually. Automatic sorting is based on the value of some expression defined for all elements of the set. By using an index into a sorted subset, you can access any subselection of data in the specified order. Such a subselection may be of interest in your end-user interface or at a certain stage in your model.

Sorting of sets

3.2 SET declaration and attributes

Each set has an optional list of attributes which further specify its intended behavior in the model. The attributes of sets are given in Table 3.1. The attributes INDEX DOMAIN and TAGS are only relevant to compound sets and indexed sets, respectively.

Set attributes

Attribute	Value-type	See also page
INDEX DOMAIN	<i>index-domain</i>	38
SUBSET OF INDEX	<i>subset-domain</i>	
PARAMETER	<i>identifier-list</i>	
TAGS	<i>tags</i>	36
TEXT	<i>string</i>	19
COMMENT	<i>comment string</i>	19
PROPERTY	NoSave	
DEFINITION	<i>set-expression</i>	
ORDER BY	<i>expression-list</i>	

Table 3.1: SET attributes

3.2.1 Simple sets

A *simple* set in AIMMS is a finite collection of elements. These elements are either strings or integers. Strings are typically used to identify real-world objects such as products, locations, persons, etc.. Integers are typically used for algorithmic purposes. With every simple set you can associate indices through which you can refer (in succession) to all individual elements of that set in indexed statements and expressions.

Definition

An example of the most basic declaration for the set *Cities* from the previous example follows.

Most basic example

```
SET:
  identifier : Cities
  index     : i,j ;
```

This declares the identifier *Cities* as a simple set, and binds the identifiers *i* and *j* as indices to *Cities* throughout your model text.

Consider a set *SupplyCities* which is declared as follows:

More detailed example

```
SET:
  identifier : SupplyCities
  subset of : Cities
  parameter : LargestSupplyCity
  text      : The subset of cities that act as supply city
  definition : { i | Exists( j | Transport(i,j) ) }
  order by  : i ;
```

The “|” operator used in the definition is to be read as “such that” (it is explained in Chapter 5). Thus, *SupplyCities* is defined as the set of all cities from

which there is transport to at least one other city. All elements in the set are ordered lexicographically. The set has no index of its own, but does have an element parameter `LargestSupplyCity` that can hold any particular element with a specific property. For instance, the following assignment forms one way to specify the value of this element parameter:

```
LargestSupplyCity := ArgMax( i in SupplyCities, sum( j, Transport(i,j) ) );
```

Note that this assignment selects that particular element from the subset of `SupplyCities` for which the total amount of `Transport` leaving that element is the largest.

With the `SUBSET OF` attribute you can tell AIMMS that the set at hand is a subset of another set, called the *subset domain*. For simple sets, such a subset domain is denoted by a single set identifier. During the execution of the model AIMMS will assert that this subset relationship is satisfied at all times.

The SUBSET OF attribute

Each simple set that is not a subset of another set is called a *root set*. As will be explained later on, root sets have a special role in AIMMS with respect to data storage and ordering.

Root sets

An index takes the value of *all* elements of a set successively and in the order specified by its declaration. It is used in operations like summation and indexed assignment over the elements of a set. With the `INDEX` attribute you can associate identifiers as indices into the set at hand. The index attributes of all sets must be unique identifiers, i.e. every index can be declared only once.

The INDEX attribute

A parameter declared in the `PARAMETER` attribute of a set takes the value of a *specific* element of that set. Throughout the sequel we will refer to such a parameter as an *element parameter*. It is a very useful device for referring to set elements that have a special meaning in your model (as illustrated in the previous example). In a later chapter you will see that an element parameter can also be defined separately as a parameter which has a set as its range.

The PARAMETER attribute

With the `TEXT` attribute you can specify one line of descriptive text for the end-user. This description can be made visible in the graphical user interface when the data of an identifier is displayed in a page object. You can use the `COMMENT` attribute to provide a longer description of the identifier at hand. This description is intended for the modeler and cannot be made visible to an end-user. The `COMMENT` attribute is a multi-line string attribute.

The TEXT and COMMENT attributes

You can make AIMMS aware that specific words in your comment text are intended as identifier names by putting them in single quotes. This has the advantage that AIMMS will update your comment when you change the name of that identifier in the model editor, or, that AIMMS will warn you when a quoted name does not refer to an existing identifier.

*Quoting
identifier names
in COMMENT*

With the ORDER BY attribute you can indicate that you want the elements of a certain set to be ordered according to a single or multiple ordering criteria. Both simple and compound sets can be ordered.

*The ORDER BY
attribute*

A special word of caution is in place with respect to specifying an ordering principle for root sets. Root sets play a special role within AIMMS because all data defined over a root set or any of its subsets is stored in the original *data entry* order in which elements have been added to that root set. Thus, the data entry order defines the natural order of execution over a particular domain, and specifying the ORDER BY attribute of a root set may influence overall execution times of your model in a negative manner. Section 13.2.7 discusses these efficiency aspects in more detail, and provides alternative solutions.

*Ordering root
sets*

The value of the ORDER BY attribute can be a comma-separated list of one or more ordering criteria. The following ordering criteria (numeric, string or user-defined) can be specified.

*Ordering
criteria*

- If the value of the ORDER BY attribute is an indexed numerical expression defined over the elements of the set, AIMMS will order its elements in increasing order according to the numerical values of the expression.
- If the value of the ORDER BY attribute is either an index into the set, a set element-valued expression, or a string expression over the set, then its elements will be ordered lexicographically with respect to the strings associated with the expression. By preceding the expression with a minus sign, the elements will be ordered reverse lexicographically.
- If the value of the ORDER BY attribute is the keyword USER, the elements will be ordered according to the order in which they have been added to the subset, either by the user, the model, or by means of the Sort operator.

When applying a single ordering criterion, the resulting ordering may not be unique. For instance, when you order according to the size of transport taking place from a city, there may be multiple cities with equal transport. You may want these cities to be ordered too. In this case, you can enforce a more refined ordering principle by specifying multiple criteria. AIMMS applies all criteria in succession, and will order only those elements that could not be uniquely distinguished by previous criteria.

*Specifying
multiple criteria*

The following set declarations give examples of various types of automatic ordering. In the last declaration, the cities with equal transport are placed in a lexicographical order.

Example

```
SET:
  identifier : LexicographicSupplyCities
  subset of : SupplyCities
  order by : i ;
SET:
  identifier : ReverseLexicographicSupplyCities
  subset of : SupplyCities
  order by : - i ;
SET:
  identifier : SupplyCitiesByIncreasingTransport
  subset of : SupplyCities
  order by : sum( j, Transport(i,j) ) ;
SET:
  identifier : SupplyCitiesByDecreasingTransportThenLexicographic
  subset of : SupplyCities
  order by : - sum( j, Transport(i,j) ), i ;
```

In general, you can use the PROPERTY attribute to assign additional properties to an identifier in your model. The applicable properties depend on the identifier type. Sets, at the moment, only support a single property.

The PROPERTY attribute

- The property NoSave specifies that the contents of the set at hand will never be stored in a case file. This can be useful, for instance, for intermediate sets that are necessary during the model's computation, but are never important to an end-user.

The properties selected in the PROPERTY attribute of an identifier are on by default, while the nonselected properties are off by default. During execution of your model you can also dynamically change a property setting through the PROPERTY statement. The PROPERTY statement is discussed in Section 8.5.

Dynamic property selection

If an identifier can be uniquely defined throughout your model by a single expression, you can (and should) use the DEFINITION attribute to specify this global relationship. AIMMS stores the result of a DEFINITION and recomputes it only when necessary. For sets where a global DEFINITION is not possible, you can make assignments in procedures and functions. The value of the DEFINITION attribute must be a valid expression of the appropriate type, as exemplified in the declaration

The DEFINITION attribute

```
SET:
  identifier : SupplyCities
  subset of : Cities
  definition : { i | Exists( j | Transport(i,j) ) } ;
```

3.2.2 Integer sets

A special type of simple set is an integer set. Such a set is characterized by the fact that the value of the SUBSET OF attribute must be equal to the predefined set Integers or a subset thereof. Integer sets are most often used for algorithmic purposes.

Integer sets

Elements of integer sets can also be used as integer values in numerical expressions. In addition, the result of an integer-valued expression can be added as an element to an integer set. Elements of non-integer sets that represent numerical values cannot be used directly in numerical expressions. To obtain the numerical value of such non-integer elements, you can use the Val function (see Section 5.2.2).

Usage in expressions

In order to fill an integer set AIMMS provides the special operator “..” to specify an entire range of integer elements. This powerful feature is discussed in more detail in Section 5.1.1.

Construction

The following somewhat abstract example demonstrates some of the features of integer sets. Consider the following declarations.

Example

```
PARAMETER:
  identifier : LowInt
  range      : Integer ;
PARAMETER:
  identifier : HighInt
  range      : Integer ;
SET:
  identifier : EvenNumbers
  subset of  : Integers
  index      : i
  parameter  : LargestPolynomialValue
  order by   : - i ;
```

The following statements illustrate some of the possibilities to compute integer sets on the basis of integer expressions, or to use the elements of an integer set in expressions.

```
! Fill the integer set with the even numbers between
! LowInt and HighInt. The first term in the expression
! ensures that the first integer is even.

EvenNumbers := { (LowInt + mod(LowInt,2)) .. HighInt by 2 };

! Next the square of each element i of EvenNumbers is added
! to the set, if not already part of it (i.e. the union results)

for ( i | i <= HighInt ) do
  EvenNumbers += i^2;
```

```

endfor;

! Finally, compute that element of the set EvenNumbers, for
! which the polynomial expression assumes the maximum value.

LargestPolynomialValue := ArgMax( i, i^4 - 10*i^3 + 10*i^2 - 100*i );

```

By default, integer sets are ordered according to the numeric value of their elements. Like with ordinary simple sets, you can override this default ordering by using the `ORDER BY` attribute. When you use an index in specifying the order of an integer set, AIMMS will interpret it as a numeric expression.

*Ordering
integer sets*

3.2.3 Relations and compound sets

A *relation* or multidimensional set is the Cartesian product of a number of simple sets or a subset thereof. Relations are typically used as the domain space for multidimensional identifiers.

Relation

An element of a relation is called a *tuple* and is denoted by the usual mathematical notation, i.e. as a parenthesized list of comma-separated elements. Throughout, the word *index component* will be used to denote the index of a particular position inside a tuple.

*Tuples and
index
components*

To reference an element in a relation, you can use an *index tuple*, in which each tuple component contains an index corresponding to a simple set.

Index tuple

Like simple sets, the elements of a relation can be referenced using a single index, by associating a *compound index* with the relation. In AIMMS, such a special type of relation is called a *compound set*, and, from here on, we will reserve the word *relation* for multidimensional sets which

Compound set

- do not have an associated compound index, and
- are not a subset of a compound set.

Compound sets support all the attributes of a simple set. In addition, compound sets support the `TAGS` attribute, which enables you to access and use the individual components of an element tuple in expressions and statements.

*Compound set
attributes*

The `SUBSET OF` attribute is mandatory for both relations and compound sets, and must contain the *subset domain* of the set. This subset domain is denoted either as a parenthesized comma-separated list of simple set identifiers, or, if it is a subset of another relation or compound set, just the name of that set.

*The SUBSET OF
attribute*

The following example demonstrates some elementary declarations of relations and compound sets, given the two-dimensional parameters `Distance(i,j)` and `TransportCost(i,j)`. The following set declaration defines a relation.

Example

```
SET:
  identifier : HighCostConnections
  subset of  : (Cities, Cities)
  definition : { (i,j) | Distance(i,j) > 0 and TransportCost(i,j) > 100 } ;
```

The following two declarations define compound sets, either because they have an associated compound index, or are a subset of a compound set.

```
SET:
  identifier : ConnectedCities
  subset of  : (Cities, Cities)
  index     : cc
  definition : { (i,j) | Distance(i,j) > 0 } ;
SET:
  identifier : ExpensiveConnections
  subset of  : ConnectedCities
  definition : { cc | TransportCost(cc) > 100 } ;
```

The above example shows that you can specify indices and element parameters for compound sets just as you could for simple sets. They are called *compound indices* and *compound element parameters*. As long as the domains are compatible, AIMMS lets you unrestrictedly interchange compound indices with tuples of indices as illustrated in the definitions of `ConnectedCities` and `ExpensiveConnections`.

*Accessing
compound set
elements*

When you use compound indices or compound element parameters, you still have access to individual index components through the use of *tags*, which can be declared in the TAGS attribute of a compound set. Tags provide a placeholder for every individual component of a tuple in a compound set to be used in expressions and assignments. To reference a specific index component or subtuple, one can simply use the fairly standard dot notation followed by a single tag or a tuple of tags.

*The TAGS
attribute*

Consider an extension of the declaration of the set `ConnectedCities` with a TAGS attribute as follows.

Examples

```
SET:
  identifier : ConnectedCities
  subset of  : (Cities, Cities)
  index     : cc
  tags      : (orig, dest)
  comment   : The tag orig for the first index component stands
              for originating cities, while the second index
              component dest stands for destination cities. ;
```

The following compound set declarations demonstrate the use of these tags.

```
SET:
  identifier : LinksWithSupplyCities
  subset of  : ConnectedCities
  definition : { cc | cc.orig in SupplyCities } ;
SET:
  identifier : BidirectionalLinks
  subset of  : ConnectedCities
  definition : { cc | cc.(dest,orig) in ConnectedCities } ;
```

Note that the second example above reverses the tag order. As a result, only those connections that also have a reverse link are selected.

The tag names that you specify in the TAGS attribute are AIMMS identifiers, and hence must be unique within your application. The length of a tag tuple (i.e. the number of components) must coincide with the dimension of the compound set. Within the tree of a compound root set and all of its subsets, you can only specify a tuple of tag names for the root set. However, you are allowed to use these tags for all subsets in the entire tree.

Uniqueness

By default, compound sets are ordered componentwise from left to right, and consistent with the ordering of the underlying simple root sets. By using one or more ordering criteria, you can overrule this default ordering, as illustrated by the following declaration.

*Ordering
compound sets*

```
SET:
  identifier : LinksByDistanceThenDoublyLexicographic
  subset of  : ConnectedCities
  definition : ConnectedCities
  order by   : Distance(cc), cc.orig, cc.dest ;
```

Example

You may wonder why AIMMS makes such a clear distinction between relations and compound sets. The reason is that relations and compound sets lead to completely different storage and access characteristics.

*Relation versus
compound set*

- When you define or use a relation, for instance, in the domain restriction of a multidimensional parameter, AIMMS still uses multidimensional index tuples to refer to the elements in the relation.
- When you use a compound index in the declaration of a multidimensional parameter, AIMMS considers that index to refer to a simple set consisting of compound elements, and you need *tags* to refer to the individual components of the compound elements.

The use of compound sets will reduce the effective dimension of the multidimensional data in your model, possibly leading to a dramatically lower memory usage for very high-dimensional identifiers. However, you should also keep in mind that accessing the individual components of a compound element may become much more expensive in the presence of compound indices. Thus, whether or not to use compound sets is a trade-off which needs careful investigation.

Reduced memory versus increased access time

In addition, in deciding to use compound set, you should take into consideration that reading and writing of compound sets, and data declared over compound sets, from and to databases is not supported (see also Chapter 25). However, in such cases it is still possible to read in the underlying relation, and convert the data to a compound set representation by assigning the relation to the compound set.

Not all operations supported

3.2.4 Indexed sets

An *indexed set* represents a family of sets defined for all elements in another set, called the *index domain*. The elements of all members of the family must be from a single (sub)set. Although membership tables allow you to reach the same effect, indexed sets often make it possible to express certain operations very concisely and intuitively.

Definition

A set becomes an indexed set by specifying a value for the INDEX DOMAIN attribute. The value of this attribute must be a single index or a tuple of indices, optionally followed by a logical condition. The precise syntax of the INDEX DOMAIN attribute is discussed on page 42.

The INDEX DOMAIN attribute

The following declarations illustrate some indexed sets with a content that varies for all elements in their respective index domains.

Example

```
SET:
  identifier   : SupplyCitiesToDestination
  index domain : j
  subset of   : Cities
  definition  : { i | Transport(i,j) } ;
SET:
  identifier   : DestinationCitiesFromSupply
  index domain : i
  subset of   : Cities
  definition  : { j | Transport(i,j) } ;
SET:
  identifier   : IntermediateTransportCities
  index domain : (i,j)
  subset of   : Cities
  definition  : DestinationCitiesFromSupply(i) * SupplyCitiesToDestination(j)
  comment    :
    All intermediate cities via which an indirect transport
```

from city i to city j with one intermediate city takes place ;

The first two declarations both define a one-dimensional family of subsets of Cities, while the third declaration defines a two-dimensional family of subsets of Cities. Note that the $*$ operator is applied to sets, and therefore denotes intersection.

The subset domain of an indexed set family can be either a simple set identifier, a compound set identifier, or another family of indexed simple or compound sets of the same or lower dimension. The subset domain of an indexed set *cannot* be a relation.

Subset domains

Declarations of indexed sets do not allow you to specify either the INDEX or PARAMETER attribute. Consequently, if you want to use an indexed set for indexing, you must locally bind an index to it. For more details on the use of indices and index binding refer to Sections 3.3 and 9.1.

No default indices

3.3 INDEX declaration and attributes

Every index used in your model must be declared exactly once. You can declare indices *indirectly*, through the INDEX attribute of a simple or compound set, or *directly* using an INDEX declaration. Note that all previous examples show indirect declaration of indices.

Direct versus indirect declaration

When you choose to declare an index not as an attribute of a set declaration, you can use the INDEX declaration. The attributes of each single index declaration are given in Table 3.2.

INDEX declaration

Attribute	Value-type	See also page
RANGE	<i>set-identifier</i>	
TEXT	<i>string</i>	19
COMMENT	<i>comment string</i>	19

Table 3.2: INDEX attributes

You can assign a default binding with a specific set to directly declared indices by specifying the RANGE attribute. If you omit this RANGE attribute, the index has no default binding to a specific set and can only be used in the context of local or implicit index binding. The details of index binding are discussed in Section 9.1.

The RANGE attribute

The following declaration illustrates a direct INDEX declaration.

Example

```
INDEX:  
  identifier : c  
  range      : Customers ;
```