March 30, 1994

The ASCEND language syntax and semantics

Arthur W. Westerberg
Peter C. Piela[1]
Engineering Design Research Center
Carnegie Mellon University

We shall present an informal bottom-up description of the ASCEND language. Being informal, we shall often include descriptions of the intended semantics along with the syntax of the items. At times the inclusion of semantics will seem to anticipate later definitions. We do this because we would also like this section to be used as a reference for the ASCEND language after one generally understands it. Often one will need to clarify a point about a particular item and will not wish to have to search in several places to do so.

Syntax is the form or structure for the statements in ASCEND, where one worries about the exact words one uses, their ordering, the punctuation, etc. *Semantics* describe the meaning of a statement.

To distinguish between syntax and semantics, consider the statement

```
y IS_A fraction;
```

Rules on the syntax for this statement tell us we need a user supplied instance name (here, y) followed by the ASCEND operator IS_A followed by a model name (fraction). The statement terminates with a semicolon. The statement semantics says we are declaring the existence of an instance called y of the type fraction as a part within the current model definition.

[1] Currently at AspenTech, Cambridge, MA.

The building blocks

One often describes the syntax for a computer language by using a Bachus-Naur formal description for it. The following is an informal description motivated by such a formal description.

Letters	a through z, A through Z (the underscore is used as a letter, but it cannot be the first letter in a name)
Digits	0 through 9
Parentheses	( ) Rounded parentheses. Used to enclose arguments for functions, to embed terms in complex arithmetic expressions. { } Curly braces. Used to enclose units. For example when one states that a quantity is 1 {kg_mole/s}, we enclose the units with curly braces. [ ] Square brackets. Used to enclose sets or elements of sets.
Punctuation	colon (:), semicolon (;)
Binary math relational operators	plus (+), minus (-), time(*), to the power (^), divides (/)
Unary math relational operator	minus (-) (There is no unary + operator.)
Procedural equals	:= *Procedural equals* differs from the ordinary equals (=) in that it means the left-hand-side (LHS) variables is to be replaced by the value of the right-hand-side (RHS) expression immediately when this statement is encountered, either while compiling or when executing a procedure on request through the ASCEND user interface. The order the system encounters these statements matters, therefore, with a later one overwriting an earlier one if both have the same LHS variable. We use an ordinary equals (=) when defining an equation to state that the LHS expression is to equal the RHS expression at the solution for the model.
Math comparators	less than or equals (<), equals (=), greater than or equals (>). (These operators are used in *non-procedural* arithmetic statements only.)
Symbol, model or model part name	Any string of up to 30 letters, digits, and underscore (_), the first of which is a letter. If one uses a longer string, the system compares and retains only the first 30 characters.
Integer	An optional negative sign (-) followed by any number of consecutive digits up to the maximum that can be represented by the computer on which one is running ASCEND. Examples are 123 -5 10
Real	ASCEND represents reals almost exactly as any other mathematically oriented programming language does. The mantissa has an optional

negative sign followed by a string of digits and at most one decimal point. The exponent is the letter *e* followed by an integer. The number must not exceed the largest the computer is able to handle. There can be no blank characters in a real. The following are legitimate reals in ASCEND:

-1
1.2
1.3e-2
7.888888e34
.6

while the following are not:

1. 2 (contains a blank within it)
1.3e2.0 (exponent has a decimal in it)
+1.3 (contains illegal unary + sign)

String

A type of constant in ASCEND. The format for a string is that of a symbol enclosed between two single quotes. The following are legal strings in ASCEND:

'H2O'
'r1'

while the following are not:

"ethanol" (double quotes not allowed)
water (no single quotes)

If an ASCEND variable is of type string, once it is given a value, that value can never be altered. The system will prevent one from trying.

ASCEND operators

IS_A
ARE_THE_SAME
ARE_ALIKE
IS_REFINED_TO

Basic variable types

REAL and BOOLEAN. These types are built into the ASCEND system, and all variables in ASCEND are of these types or *refinements* of these types. Variables of type *real* and *boolean* are the only ones at present whose values may be altered during the course of solving a model. (We expect in the near future to add a form of integer whose value may be changed. This capability is needed to allow integer programming problems to be expressed.) Boolean variables may take on the values *true* or *false* only.

Basic constant types

INTEGER, BOOLEAN and STRING. These types are built into the system and all constants in the system are of this type or refinements of these basic types. Strings are any set of characters enclosed within single quotes. For example, '34xyz' is a string.

Built-in constants

ASCEND has the following built-in constants. If any of these is used, the units for the variables are the built in SI units for the dimensionality for the variable; thus these constants are used without any units being specified.

ZERO	equals 0
+INFINITY	the largest real the machine is capable of representing

-INFINITY	the most negative real the machine is capable of representing

An example of use is, assuming P is a pressure:

```
P := ZERO;
```

which sets P to zero (pascals).

Generic_real

The base-type for all *computable* variables in the current ASCEND system. It has wild card dimensionality. ASCEND has defined the following parts with associated initial values for it:

attribute	type	default
low_bound	real	-INFINITY
upper_bound	real	+INFINITY
fixed	boolean	FALSE
eligible	boolean	TRUE
solved	boolean	FALSE
defined	boolean	FALSE
initialized	boolean	FALSE
required	boolean	TRUE

low_bound and *upper_bound* are bounds for a variable which are monitored and maintained during solving. The flag *fixed* indicates if the variable is to be held fixed during solving. The flag *eligible* indicates the variable is structurally eligible to selected to be fixed. The flag*solved* indicates if the solver has set the value as the result of a converged solution. The flag *defined* indicates that the variable has been defined in an IS_A statement. If FALSE, it has been used in a statement but never defined with an IS_A statement. The flag *initialized* indicates that the value the variable has is from some process that set it (and is not just what was in the storage location by chance). *required* indicates that the user wishes this value to be computed.

A full definition of generic_real is in the base.lib.

Variables

A variable in ASCEND is an instance of a variable type. The base_type of a computable variable is generic_real; thus, all variables have the parts defined for generic_real.

Dimensionality

Variables have dimensionality such as length/time for velocity. Dimensionality is to be distinguished from the units such as ft/s. Dimensionality is built up from the following base dimensions:

symbol	dimension	typical unit
m	mass	kg
l	length	m
t	time	s
tmp	temperature	K
mole	moles	kg_mole
c	coulombs	coul
currency	currency	USdollar

The atom definitions in the library illustrate the use of dimensionality.

Dimensions may be any combination of these symbols along with rounded parentheses, (), and the operators *, ^ and /. Examples include

m/t or *m*l^2/t^2/tmp* {this latter means $((m*l^2)/t^2)/tmp$}. The second operand for the "to the power" operator ^ must be an integer constant (e.g., -2 or 3).

If the dimensionality for a type is undefined, then ASCEND gives it a wild card dimensionality. If ASCEND can later deduce its dimensionality from its use in a model definition, it will do so. For example, suppose *a* has wild card dimensionality, *b* has dimensionality of l/t. Then the statement:

a + b = 3 {ft/s};

requires that *a* have the same dimensionality as the other two terms, namely, *l/t*. ASCEND will assign this dimensionality to *a*.

Units

When a real number is used in a mathematical expression in ASCEND, it must have a set of units expressed with it. If it does not, ASCEND assumes the number is dimensionless, which may not be the intent of the modeler. An example is shown just above where the number 3 has the units ft/s associated with it. Units appear in curly braces {}.

Unit definitions are in the library called base.lib.

UNITS module

A modeler can add any set of units desired to the set defined in base.lib by creating a UNITS definition module. It has the same format as the UNITS module in the base.lib, namely:

```
UNITS
        any number of units definitions
END;
```

Any previously defined unit can be used in a subsequent definition. An example units definition statement is:

amp := coul/s

where coul and s have been previously defined.

Sets

Sets are lists of elements, all of type integer or all of type string, enclosed between square brackets []. The following are examples of sets:

```
['methane', 'ethane', 'propane']
[1..5, 7, 15]
[2..n_stages]
[1, 4, 2, 1, 16]
[]
```

The symbols 1..5 is an allowable shorthand for the integers 1, 2, 3, 4 and 5 while the symbol 2..n_stages (where n_stages must be of type integer) means all integers from 2 to n_stages. If n_stages is less than 2, then the third set is empty. The repeated occurence of 1 in the fourth set is ignored. The fifth set is the empty set.

We use the term *set* in the pure mathematical sense. The elements have no order. The only thing one can ask of a set is if an element is a member of it. Thus repeated elements used in defining it are ignored.

The elements of sets **cannot** themselved be sets in ASCEND; i.e., there can be no sets of set.

A set of integers may appear to be ordered to the modeler as the natural numbers have an order. One can find the successor to an integer (by adding one to it). It is the user imposing and using the ordering, not ASCEND. ASCEND simple sees these integers as elements in the set with NO ordering. Therefore, there are no operators in ASCEND such as successor or precursor member.

Arrays

An array is a list of instances indexed over a set. The instances are all of the same *base* type (as that is the only way they can be defined). Individual members of a list may be more refined than the other members (we shall illustrate that possibility later). The following are arrays in ASCEND.

```
stage[1..n_stages]
y[components]
column[areas][processes]
```

where *components*, *areas* and *processes* are sets. For example *components* could be the set of symbols ['ethylene','propylene'] , *areas* the set of symbols ['feed_prep','prod_purification'] while *processes* could be the set ['alcohol_manuf', 'poly_propropylene_manuf']. Note that the third example is a list of lists (the way that ASCEND permits a multiply subscripted list). There can be up to eight levels of subscripts for a list.

The following are elements in these arrays:

```
stage[1]
y['ethylene']
column['feed_prep'][alcohol_manuf']
```

provided that n_stages is 1 or larger.

Index variable

One can introduce a variable as an index which is to range over a set. Index variables are local to the statements in which they occur. An example of using an index variable is the following FOR statement (see later):

```
FOR i IN components CREATE
        VLE_equil[i]:  y[i] = K[i]*x[i];
END;
```

Label

One can label statements which define relationships (equalities and inequalities) in ASCEND. A label is a sequence of alphameric characters ending in a colon. An example of a labeled statement is

```
mass_balance: m_in = m_out;
```

If a relationship is defined within a FOR statement (see later), it must be indexed so each occurrence of the statement is distinguishable from the others. An example is:

```
FOR i IN components CREATE
        equil[i]: y[i] = K[i]*x[i];
END;
```

The ASCEND interactive user interface identifies relationships by their labels. If one has not provided such a label, the system generates one (a system generated label is not very descriptive as you might imagine).

List of names

Often in a statement one can include a list of names. A list is one or more names where multiple list entries are separated from each other by commas. Examples of a list of names are:

```
T1, inlet_T, outlet_T
y[components], y_in
stage[1..n_stages]
```

Single operand real ASCEND functions

exp	exponential (i.e., exp(x) = e^x)
ln	log to the base e
sin	sine
cos	cosine
tan	tangent
atn	arctan
erf	error function

The base type real operands for ASCEND functions are enclosed within rounded parentheses, (). An example of use is:

```
y = A*exp(-B/T);
```

The following function is available in procedures ONLY.

abs	absolute value

Multiple operand real functions

SUM(list_of_expressions) Add all expressions in the list

PROD(list_of_expressions) Multiply all expressions in the list.

For the SUM and PROD functions, the base type real items can be arbitrary arithmetic expressions. The resulting items must all be units compatible for the SUM operator.

An examples of the use is:

```
SUM(y[components]) = 1;
```

or, equivalently, one could write:

```
SUM(y[i], i IN components) = 1;
```

The following multiple operand functions are available in procedures only.

max, min	find the maximum (minimum) value of the items on the list

Multiple operand boolean functions

INTERSECTION(list_of_sets) Find the intersection (and) of the sets listed

UNION(list_of_sets) Find the union (or) of the sets listed

Set function UNION

A function taken over a list of sets. The result is a set whose members are the union of the sets in the list. The syntax is

UNION(list_of_sets)

Consider the following sets for the examples to follow.

```
A := [1, 2, 3, 5, 9];
B := [2, 4, 6, 8];
```

Then UNION(A, B) is equal to the set [1,2 3, 4, 5, 6, 8, 9] which equals [1..6, 6, 9].

Set function INTERSECTION

A function taken over a list of sets. The result is a set whose members are the intersection of the sets in the list. The syntax is

INTERSECTION(list_of_sets)

For the sets A and B defined just above, INTERSECTIION(A, B) is the set [2];

Set difference

One can subtract one set from another. The result is the first set less any members in the set union of the first and second set. The syntax is

first_set - second_set

For the sets A and B defined above, the set difference A - B is the set [1, 3, 5, 9] while the set difference
B - A is the set [4, 6, 8];

Set function CHOICE

The format for this instruction is

CHOICE(set)

The result of running the CHOICE function over a set is any arbitrary single element of that set.

The result of running CHOICE(A) is any member of the set A. The result is a member, not a set. To make the result into a set, it must be enclosed in square brackets. Thus [CHOICE(A)] is a set with a single element in it which could be any member of the set A.

To reduce a set by one element, one can use the following

```
A_less_one IS_A set OF integer;
A_less_one := A - [CHOICE(A)];
```

Statements

Statements

Statements in ASCEND terminate with a semicolon (;). Statements may extend over any number of lines. They may have blank lines in the middle of them. There may be several statements on a single line.

One constructs statements using the above building blocks.

Compound statements

Some statements in ASCEND can contain other statements as a part of them. The three examples are the FOR/CREATE, the FOR/DO and the IF statements.

In the following statement descriptions, we show keywords in capital letters. These words must appear in capital letters as shown in ASCEND statements. We show optional parts to a statement enclosed in double angle brackets (« ») and user supplied names in lower-case *italic* letters. (Remember that ASCEND treats the underscore (_) as a letter). The user may substitute any name desired for these names. We use names that describe the kind of name the user should use.

Begin model statement

This statement has the syntax

MODEL *model_name* « REFINES
name_of_previously_defined_model»;

It is used to initiate MODEL definitions. (The REFINES part is optional.)

Begin atom statement

The syntax for the begin atom statement is

ATOM atom_name REFINES variable_type
«DIMENSION dimension_expression» «DEFAULT
value»;

The options for variable type are those built into the system (real, integer, boolean) or the name of any other previously defined atom.

The dimension attribute is for variables whose base type is real. It is an optional field. If not defined for any atom whose base type is real, the dimensions will be left as undefined. Any variable which is later declared to be one of these types will be given *wild card* dimensionality (represented in the interactive display by an asterisk (*). The system will deduce the dimensionality from its use in the relationships in which it appears or in the declaring of default values for it, if it can.

The default field is also optional. If the atom has a declared dimensionality, then this value must be expressed with units which are compatible with this dimensionality.

End model/atom statement

This syntax for this statement is

END *model_name*;

or

END *atom_name*;

It is used to terminate model and atom definitions.

IS_A statement

This statement has the syntax

list_of_instance_names IS_A *model_name*;

The IS_A statement allows us to declare *instances* of a given *type* to exist within a model definition.

Instances and types

This is an opportune time to emphasize again the distinguish between the terms *instance* and *type*. A *type* in ASCEND is what we define when we declare an ASCEND model or atom. It is the definition of the attributes that something will have if the system creates one of them using this definition.

In ASCEND an *instance* is a *named part* that exists within a model definition. It does not physically exist for computing. (We can contrast this sense of an instance with that in most other object-oriented systems where it is an actual compiled object that exists for computations, etc. In ASCEND a computable object does not exist until someone asks for a model definition to be compiled. Thus we use *instance* in an in-between sense of its being a part of a model definition only.)

IS_REFINED_TO statement

This statement has the syntax

```
list_of_instances IS_REFINED_TO model_name;
```

We use this statement to change the type of the instances listed to being of a more refined type. The modeler has to have defined each member on the list of instance previously. The model_name has to be a type which is at an ancestor in the refinement hierarchy for all of the instances on the list.

An example of its use is as follows. First we define the the parts called fl1, fl2 and fl3 which are of type flash.

```
fl1, fl2, fl3 IS_A flash;
```

Assume that there exists in the previously define model definition an adiabatic_flash as a refinement of flash. Then we can make fl1 and fl3 into more refined types by stating:

```
fl1, fl3 IS_REFINED_TO adiabatic_flash;
```

ARE_THE_SAME statement

The format for this instruction is

```
list_of_instance_names ARE_THE_SAME;
```

All items on the list must have compatible types. For the example in Fig. 1, consider a model where we define the following parts:

```
a1 IS_A a;
b1 IS_A b;
c1 IS_A c;
d1 IS_A d;
e1 IS_A e;
```

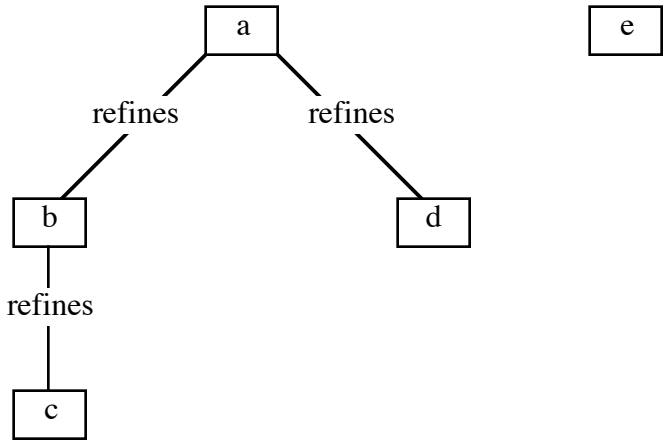

Figure 1. Diagram of the model hierarchy

Then the following ARE_THE_SAME statement is legal

a1, b1, c1 ARE_THE_SAME;

while the following are not

b1, d1 ARE_THE_SAME;
a1, c1, d1 ARE_THE_SAME;
b1, e1 ARE_THE_SAME;

When compiling a model, ASCEND will put all of the instances mentioned as being the same into an ARE_THE_SAME*clique*. Merging any other item with a member of the clique makes it the same as all the other items in the clique, i.e., it adds the newly mentioned items to the existing clique.

ASCEND merges all members of a clique by first checking that all members on the clique are type compatible. It then changes the type designation for all clique members to that of the most refined type for any member on the clique.

It next looks inside each of the instances, all of which are now of the same type, and puts all of the parts with the same name into their respective ARE_THE_SAME cliques. The process repeats by processing these cliques until all parts of all parts of all parts, etc., are their respective most refined type or discovered to be type incompatible.

There are now lots of cliques associated with the instances being merged. The type associated with each such clique is now either a model or an atom (i.e., a variable). If a model, only one member of the clique generates its equations. If a variable, it assigns all members to the same storage location.

ARE_ALIKE statement

The format for this statement is

list_of_instance_names ARE_ALIKE;

ARE_ALIKE is similar in implementation to ARE_THE_SAME. The compiler places all instances on the list onto an ARE_ALIKE clique. It checks that the members are type compatible and then it converts each into the most refined type of any instance in the clique. At that point it stops. It does not continue by placing the parts into cliques nor does it assign common storage locations, etc. It may be thought of as a partial merge. There are important consequences in modeling of such a partial merge.

One consequence of ARE_ALIKE is to prevent model misuse when configuring models. For example, suppose a modeler creates a pressure changing model. The modeler is not concerned about the type of the streams into and out of the device but does care that these streams are of the same final type. For example, the modeler wants both to be liquid streams if either is or both to be vapor streams if either is. By declaring both to be streams only but declaring the two streams to be alike, the modeler accomplishes this intent. Suppose the modeler merges the inlet stream with a liquid outlet stream from a reactor. The merge operation makes the inlet stream into a liquid stream. The outlet stream, being in an ARE_ALIKE clique with the inlet stream, also becomes a liquid stream. Any subsequent merge of the outlet stream with a vapor stream will lead to an error due to type incompatibility when ASCEND attempts to compile that merge. Without the ARE_ALIKE statement, the compiler would detect no such incompatibility.

Another purpose is the propagation of type through a model which we just saw taking place in the previous paragraph. Altering the type of the inlet stream through merging it with a liquid stream automatically made the outlet stream into a liquid stream.

If all the liquid streams within a distillation column are alike, then the modeler can make them all into streams with a particular set of components in them and with the same method used for physical property evaluation by merging only one of them with a liquid stream of this type.

UNIVERSAL qualifier

The begin model and the begin atom statements can have a qualifier preceding them which causes all instances of the type being defined to be merged. The syntax for these statements are

```
UNIVERSAL model_begin_statement
UNIVERSAL atom_begin_statement
```

Examples are

```
UNIVERSAL MODEL methane REFINES
  components_constants;

UNIVERSAL ATOM circle_constant REFINES
real DIMENSIONLESS;
```

In the former case, ASCEND will merge all instances of the physical property data for methane (such as it molecular weight, its normal boiling point, etc) into the same instance. In the later, all instances of circle_constant (3.1415926536) become the same instance.

Mathematical expression

The syntax for a mathematical expression is any legal combination of variable names and arithmetic operators in the normal notation. An expression may contain any number of matched rounded parentheses, ( ), to clarify meaning. The following is a legal arithmetic expression:

```
y^2+(sin(x)-tan(z))*q
```

Each term in a mathematical expression (terms are separated by + or - operators) must have the same dimensionality.

An expression may contain an index variable as a part of the calculation if that index variable is of type integer. (See the FOR/CREATE and FOR/DO statements below.) An example is:

```
term[i] = a[i]*x^(i-1);
```

Logical expression

An expression whose value is true or false is a logical expression. The mathematical comparitor operators (=, <, >) are common in such expressions. Such expressions may also contain boolean variables. If laminar is a boolean, then the following is an example:

```
(a>0) + ((b<3) * laminar)
```

The plus operator acts like an OR among the terms while the times operator acts like an AND. Think of true being equal to 1 and false being equal to 0 with the 1+1=0+1=1+0=1 and 1*1=1, 0*1=1*0=0*0=0. If *a* in fact equals -5, *b*=2.99 and *laminar* is true, this expression has the value

```
false OR (true AND true) = true
```

or in terms of ones and zeros

```
0 + (1 * 1) = 1
```

Mathematical relation

The syntax for a mathematical relation is

```
math_expression1  math_comparitor
        math_expression2;
```

where a math_comparator is one of the following: equals (=), less than or equals (<) or greater than or equals (>). Math relationships are equations or inequalities. The following are examples:

```
Ftot*y['methane'] = m['methane'];
y['ethanol'] > 0;
```

where this last one means that the mole fraction for ethanol is greater than or equal to 0. Equations must be dimensionally correct.

FOR/CREATE statement

The FOR/CREATE statement is a compound statement. Its format is:

```
FOR index_variable IN set CREATE
        list_of_statements;
END;
```

This statement must be in the nonprocedural part of the model definition only. Every statement in the list must have at least one

variable which is indexed by the index variable, else the same statement will be produced multiple times. An example is

```
FOR i IN components CREATE
        a.y[i], b[i]  ARE_THE_SAME;
        y[i] = K[i]*x[i];
END;
```

INITIALIZE statement

This statement introduces the procedure definitions in ASCEND. All statements following this statement are to define procedures in ASCEND while all before it are for the declarative part of ASCEND. The syntax for this statement is simply

```
INITIALIZE
```

with no punctuation. The next statement must be a procedure heading statement. If there are no procedure definitions, this statement is omitted.

Begin PROCEDURE statement

The syntax for a begin procedure statement is

```
PROCEDURE procedure_name;
```

End PROCEDURE statement

The syntax of a procedure ending statement is

```
END procedure_name;
```

LHS name

A LHS_name is one of the following

```
instance_name
array_name[set_name]
set_name
```

RHS expression

A RHS expression is one of the following

```
mathematical, boolean or set expression
```

Procedural assignment statement

The syntax is

```
        instance_name := mathematical_expression;
or
        array_name[set_name] :=
        mathematical_expression
or
        set_name := [list_of_symbols];
or
        set_name := set_expression;
```

Its meaning is that the value for the variable on the LHS is set to the value of the expression on the RHS.

FOR/DO statement

This statement is similar to the FOR/CREATE statement except it can only appear in a procedure definition. An example would be

```
FOR i IN [1..n_stages] DO
        T[i] := T[1] + (i-1)*DT;
END;
```

IF statement

This statement can only appear in a procedure definition. Its syntax is

```
IF (logical_expression)
THEN
        list_of_statements
ELSE
        list_of_statements
END;
```

or

```
IF (logical_expression)
THEN
        list_of_statements
END;
```

If the logical expression has a value of true, the statements ASCEND will execute the statements in the THEN part. If it is false, it executes the statements in the optional ELSE part.

RUN statement

This statement can appear only in a procedure. Its format is

```
RUN name_of_procedure;
```

The procedure can be within the current model or within any of its parts.

PROCEDURE

A procedure in ASCEND must appear following the INITIALIZE statement within a model. The first statement must be a begin procedure statement and the last an end procedure statement. In-between exist procedural assignment, FOR/DO and IF statements. The systems executes procedure statements in the order they are encountered.

Specifically disallowed are IS_A, IS_REFINED_TO, ARE_THE_SAME and ARE_ALIKE statements as these "declare" the structure of the model and belong only in the delarative section.

ASCEND model

An ASCEND model begins with a *begin model* statement and ends with an *end model* statement. Both must have the identical model_name. It has a declarative part and an optional procedural part headed by an INITIALIZE statement.