



Modelica[®] – A Unified Object-Oriented Language for Systems Modeling

Language Specification

Version 3.5-dev

December 16, 2020

Modelica Association

Abstract

This document defines the Modelica¹ language, version 3.5-dev, which is developed by the Modelica Association, a non-profit organization with seat in Linköping, Sweden. Modelica is a freely available, object-oriented language for modeling of large, complex, and heterogeneous systems. It is suited for multi-domain modeling, for example, mechatronic models in robotics, automotive and aerospace applications involving mechanical, electrical, hydraulic control and state machine subsystems, process oriented applications and generation and distribution of electric power. Models in Modelica are mathematically described by differential, algebraic and discrete equations. No particular variable needs to be solved for manually. A Modelica tool will have enough information to decide that automatically. Modelica is designed such that available, specialized algorithms can be utilized to enable efficient handling of large models having more than one hundred thousand equations. Modelica is suited and used for hardware-in-the-loop simulations and for embedded control systems. More information is available at <https://www.modelica.org>.

¹Modelica is a registered trademark of the Modelica Association.

Copyright © 1998-2020, Modelica Association (<https://www.modelica.org>)

All rights reserved. Reproduction or use of editorial or pictorial content is permitted, i.e., this document can be freely distributed especially electronically, provided the copyright notice and these conditions are retained. No patent liability is assumed with respect to the use of information contained herein. While every precaution has been taken in the preparation of this document no responsibility for errors or omissions is assumed.

The contributors to this and to previous versions of this document are listed in appendix E. All contributors worked voluntarily and without compensation.

Contents


Preface	5
1 Introduction	7
1.1 Overview of Modelica	7
1.2 Scope of the Specification	7
1.3 Some Definitions	8
1.4 Notation and Grammar	8
2 Lexical Structure	9
2.1 Character Set	9
2.2 Comments	9
2.3 Identifiers, Names, and Keywords	10
2.4 Literal Constants	11
2.5 Operator Symbols	12
3 Operators and Expressions	13
3.1 Expressions	13
3.2 Operator Precedence and Associativity	13
3.3 Evaluation Order	14
3.4 Arithmetic Operators	15
3.5 Equality, Relational, and Logical Operators	15
3.6 Miscellaneous Operators and Variables	16
3.7 Built-in Intrinsic Operators with Function Syntax	18
3.8 Variability of Expressions	31
4 Classes, Predefined Types, and Declarations	34
4.1 Access Control – Public and Protected Elements	34
4.2 Double Declaration not Allowed	35
4.3 Declaration Order and Usage before Declaration	35
4.4 Component Declarations	35
4.5 Class Declarations	41
4.6 Specialized Classes	44
4.7 Balanced Models	46
4.8 Predefined Types and Classes	52
5 Scoping, Name Lookup, and Flattening	58
5.1 Flattening Context	58
5.2 Enclosing Classes	58
5.3 Static Name Lookup	58
5.4 Instance Hierarchy Name Lookup of Inner Declarations	60
5.5 Simultaneous Inner/Outer Declarations	62
5.6 Flattening Process	63
6 Interface or Type Relationships	69
6.1 Interface Terminology	69
6.2 The Concepts of Type, Interface and Subtype	70
6.3 Interface or Type	71
6.4 Interface Compatibility or Subtyping	73

6.5	Plug-Compatibility or Restricted Subtyping	74
6.6	Function-Compatibility or Function-Subtyping for Functions	75
6.7	Type Compatible Expressions	76
7	Inheritance, Modification, and Redeclaration	78
7.1	Inheritance – Extends Clause	78
7.2	Modifications	81
7.3	Redeclaration	86
8	Equations	94
8.1	Equation Categories	94
8.2	Flattening and Lookup in Equations	94
8.3	Equations in Equation Sections	94
8.4	Synchronous Data-flow Principle and Single Assignment Rule	102
8.5	Events and Synchronization	102
8.6	Initialization, initial equation, and initial algorithm	104
9	Connectors and Connections	108
9.1	Connect-Equations and Connectors	108
9.2	Generation of Connection Equations	114
9.3	Restrictions of Connections and Connectors	116
9.4	Equation Operators for Overconstrained Connection-Based Equation Systems	119
10	Arrays	125
10.1	Array Declarations	125
10.2	Flexible Array Sizes	127
10.3	Built-in Array Functions	127
10.4	Vector, Matrix and Array Constructors	133
10.5	Array Indexing	136
10.6	Scalar, Vector, Matrix, and Array Operator Functions	138
10.7	Empty Arrays	142
11	Statements and Algorithm Sections	144
11.1	Algorithm Sections	144
11.2	Statements	145
12	Functions	152
12.1	Function Declaration	152
12.2	Function as a Specialized Class	153
12.3	Pure Modelica Functions	155
12.4	Function Call	157
12.5	Built-in Functions	164
12.6	Record Constructor Functions	164
12.7	Declaring Derivatives of Functions	168
12.8	Declaring Inverses of Functions	172
12.9	External Function Interface	174
13	Packages	188
13.1	Package as Specialized Class	188
13.2	Importing Definitions from a Package	188
13.3	The Modelica Library Path – MODELICAPATH	190
13.4	Mapping Package/Class Structures to a Hierarchical File System	191
13.5	External resources	192
14	Overloaded Operators	193
14.1	Overview of overloaded operators	193
14.2	Matching Function	194
14.3	Overloaded Constructors	194
14.4	Overloaded String Conversions	195

14.5	Overloaded Binary Operations	195
14.6	Overloaded Unary Operations	196
14.7	Example of Overloading for Complex Numbers	196
15	Stream Connectors	200
15.1	Definition of Stream Connectors	200
15.2	Stream Operator inStream and Connection Equations	201
15.3	Stream Operator actualStream	205
16	Synchronous Language Elements	206
16.1	Rationale for Clocked Semantics	207
16.2	Definitions	208
16.3	Clock Constructors	211
16.4	Clocked State Variables	214
16.5	Partitioning Operators	214
16.6	Clocked When-Clause	218
16.7	Clock Partitioning	218
16.8	Continuous-Time Equations in Clocked Partitions	221
16.9	Initialization of Clocked Partitions	225
16.10	Other Operators	225
16.11	Semantics	226
17	State Machines	229
17.1	Transitions	229
17.2	State Machine Graphics	231
17.3	State Machine Semantics	232
18	Annotations	241
18.1	Vendor-Specific Annotations	241
18.2	Annotations for Documentation	241
18.3	Annotations for Code Generation	245
18.4	Annotations for Simulations	247
18.5	Annotation for single use of class	248
18.6	Annotations for Graphical Objects	248
18.7	Annotations for the Graphical User Interface	258
18.8	Annotations for Version Handling	263
18.9	Annotations for Access Control to Protect Intellectual Property	269
18.10	Annotations for Functions	273
18.11	Annotation Choices for Modifications and Redclarations	273
18.12	Annotation for External Libraries and Include Files	273
19	Unit Expressions	274
19.1	The Syntax of Unit Expressions	274
19.2	Examples	275
20	The Modelica Standard Library	276
A	Glossary	277
B	Modelica Concrete Syntax	281
C	Modelica DAE Representation	289
D	Derivation of Stream Equations	292
E	Modelica Revision History	297
	Bibliography	326

Preface

Modelica is a freely available, object-oriented language for modeling of large, complex, and heterogeneous physical systems. From a user's point of view, models are described by schematics, also called object diagrams. Examples are shown below:



A schematic consists of connected components, like a resistor, or a hydraulic cylinder. A component has *connectors* (often also called *ports*) that describe the interaction possibilities, e.g., an electrical pin, a mechanical flange, or an input signal. By drawing connection lines between connectors a physical system or block diagram model is constructed. Internally a component is defined by another schematic, or on “bottom” level, by an equation-based description of the model in Modelica syntax.

The Modelica language is a textual description to define all parts of a model and to structure model components in libraries, called packages. An appropriate Modelica simulation environment is needed to graphically edit and browse a Modelica model (by interpreting the information defining a Modelica model) and to perform model simulations and other analysis. Information about such environments is available at www.modelica.org/tools. Basically, all Modelica language elements are mapped to differential, algebraic and discrete equations. There are no language elements to describe directly partial differential equations, although some types of discretized partial differential equations can be reasonably defined, e.g., based on the finite volume method and there are Modelica libraries to import results of finite-element programs.

This document defines the details of the Modelica language. It is not intended to learn the Modelica language with this text. There are better alternatives, such as the Modelica books referenced at www.modelica.org/publications. This specification is used by computer scientist to implement a Modelica translator and by modelers who want to understand the exact details of a particular language element.

The text directly under the chapter headings are non-normative introductions to the chapters.

The Modelica language has been developed since 1996. This document describes version 3.4 of the Modelica language. A complete summary is available in appendix E.1.

Chapter 1

Introduction

1.1 Overview of Modelica

Modelica is a language for modeling of physical systems, designed to support effective library development and model exchange. It is a modern language built on acausal modeling with mathematical equations and object-oriented constructs to facilitate reuse of modeling knowledge.

1.2 Scope of the Specification

The semantics of the Modelica language is specified by means of a set of rules for translating any class described in the Modelica language to a flat Modelica structure.

A class (of specialized class `model`, `class` or `block`) intended to be simulated on its own is called a *simulation model*.

The flat Modelica structure is also defined for other cases than simulation models; including functions (can be used to provide algorithmic contents), packages (used as a structuring mechanism), and partial models (used as base-models). This allows correctness to be verified for those classes, before using them to build the simulation model.

There are specific semantic restrictions for a simulation model to ensure that the model is complete; they allow its flat Modelica structure to be further transformed into a set of differential, algebraic and discrete equations (= flat hybrid DAE). Note that satisfying the semantic restrictions does not guarantee that the model can be initialized from the initial conditions and simulated.

Modelica was designed to facilitate symbolic transformations of models, especially by mapping basically every Modelica language construct to continuous or instantaneous equations in the flat Modelica structure. Many Modelica models, especially in the associated Modelica Standard Library, are higher index systems, and can only be reasonably simulated if symbolic index reduction is performed, i.e., equations are differentiated and appropriate variables are selected as states, so that the resulting system of equations can be transformed to state space form (at least locally numerically), i.e., a hybrid DAE of index zero. In order to allow this structural analysis, a tool may reject simulating a model if parameters cannot be evaluated during translation – due to calls of external functions or initial equations/initial algorithms for `fixed=false` parameters. Accepting such models is a quality of implementation issue. The Modelica specification does not define how to simulate a model. However, it defines a set of equations that the simulation result should satisfy as well as possible.

The key issues of the translation (or flattening) are:

- Expansion of inherited base classes
- Parameterization of base classes, local classes and components
- Generation of connection equations from connect-equations

The flat hybrid DAE form consists of:

- Declarations of variables with the appropriate basic types, prefixes and attributes, such as `parameter Real v=5`.
- Equations from equation sections.
- Function invocations where an invocation is treated as a set of equations which involves all input and all result variables (number of equations = number of basic result variables).
- Algorithm sections where every section is treated as a set of equations which involves the variables occurring in the algorithm section (number of equations = number of different assigned variables).
- When-clauses where every when-clause is treated as a set of conditionally evaluated equations, also called instantaneous equations, which are functions of the variables occurring in the clause (number of equations = number of different assigned variables).

Therefore, a flat hybrid DAE is seen as a set of equations where some of the equations are only conditionally evaluated (e.g. instantaneous equations are only evaluated when the corresponding when-condition becomes true). Initial setup of the model is specified using start-values and instantaneous equations that hold at the initial time only.

A Modelica class may also contain annotations, i.e. formal comments, which specify graphical representations of the class (icon and diagram), documentation text for the class, and version information.

1.3 Some Definitions

The semantic specification should be read together with the Modelica grammar. Non-normative text, i.e., examples and comments, are enclosed in [...] and set in italics. Additional terms are explained in the glossary in appendix A. Some important terms are defined below.

Definition 1.1. *Component.* An element defined by the production `component-clause` in the Modelica grammar (basically a variable or an instance of a class) □

Definition 1.2. *Element.* Class definitions, extends-clauses and component-clauses declared in a class (basically a class reference or a component in a declaration). □

Definition 1.3. *Flattening.* The translation of a model described in Modelica to the corresponding model described as a hybrid DAE, involving expansion of inherited base classes, parameterization of base classes, local classes and components, and generation of connection equations from connect-equations (basically, mapping the hierarchical structure of a model into a set of differential, algebraic and discrete equations together with the corresponding variable declarations and function definitions from the model). □

1.4 Notation and Grammar

The meta symbols (of the extended BNF-grammar) are defined in appendix B.1.

Boldface denotes keywords of the Modelica language. Keywords are reserved words and shall not be used as identifiers, with the exception of `initial` which is a keyword in section headings, and `der` which is a keyword for declaration functions, but it is also possible to call the functions `initial` and `der`.

See appendix B for a full lexical specification and grammar.

Chapter 2

Lexical Structure

This chapter describes several of the basic building blocks of Modelica such as characters and lexical units including identifiers and literals. Without question, the smallest building blocks in Modelica are single characters belonging to a character set. Characters are combined to form lexical units, also called tokens. These tokens are detected by the lexical analysis part of the Modelica translator. Examples of tokens are literal constants, identifiers, and operators. Comments are not really lexical units since they are eventually discarded. On the other hand, comments are detected by the lexical analyzer before being thrown away.

The information presented here is derived from the more formal specification in appendix B.

2.1 Character Set

The character set of the Modelica language is Unicode, but restricted to the Unicode characters corresponding to 7-bit ASCII characters in several places; for details see appendix B.1.

2.2 Comments

There are two kinds of comments in Modelica which are not lexical units in the language and therefore are treated as white-space by a Modelica translator. The white-space characters are space, tabulator, and line separators (carriage return and line feed); and white-space cannot occur inside tokens, e.g., `<=` must be written as two characters without space or comments between them. The following comment variants are available:

```
// comment & Characters from // to the end of the line are ignored.
/* comment */ & Characters between /* and */ are ignored, including line
  terminators.
```

[The comment syntax is identical to that of C++.]

Modelica comments do not nest, i.e., `/* */` cannot be embedded within `/* */`. The following is *invalid*:

```
/* Commented out — erroneous comment, invalid nesting of comments!
  /* This is an interesting model */
  model interesting
  ...
  end interesting;
*/
```

There is also a description-string, that is part of the Modelica language and therefore not ignored by the Modelica translator. Such a description-string may occur at the end of a declaration, equation, or statement or at the beginning of a class definition. For example:

```
model TempResistor "Temperature dependent resistor"
  ...
```

```

parameter Real R "Resistance for reference temp.";
...
end TempResistor;

```

2.3 Identifiers, Names, and Keywords

Identifiers are sequences of letters, digits, and other characters such as underscore, which are used for *naming* various items in the language. Certain combinations of letters are *keywords* represented as *reserved* words in the Modelica grammar and are therefore not available as identifiers.

2.3.1 Identifiers

Modelica *identifiers*, used for naming classes, variables, constants, and other items, are of two forms. The first form always starts with a letter or underscore (`_`), followed by any number of letters, digits, or underscores. Case is significant, i.e., the names `Inductor` and `inductor` are different. The second form (Q-IDENT) starts with a single quote, followed by a sequence of any printable ASCII character, where single-quote must be preceded by backslash, and terminated by a single quote, e.g. `'12H'`, `'13\H'`, `'+foo'`. Control characters in quoted identifiers have to use string escapes. The single quotes are part of the identifier, i.e., `'x'` and `x` are distinct identifiers. The redundant escapes (`'\?'` and `'\"'`) are the same as the corresponding non-escaped variants (`'?'` and `'\"'`), but are only for use in Modelica source code. A full BNF definition of the Modelica syntax and lexical units is available in appendix B.

```

IDENT    = NONDIGIT { DIGIT | NONDIGIT } | Q-IDENT
Q-IDENT  = "'" { Q-CHAR | S-ESCAPE } "'"
NONDIGIT = "_" | letters "a" ... "z" | letters "A" ... "Z"
DIGIT    = 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
Q-CHAR   = NONDIGIT | DIGIT | "!" | "#" | "$" | "%" | "&" | "(" | ")" | "*" |
          "+" | "," | "-" | "." | "/" | ":" | ";" | "<" | ">" | "=" | "?" | "@" | "["
          | "]" | "^" | "{" | "}" | "|" | "~" | " " | ""
S-ESCAPE = "\'" | "\"" | "\?" | "\\\" | "\a" | "\b" | "\f" | "\n" | "\r" | "\t"
          | "\v"

```

2.3.2 Names

A *name* is an identifier with a certain interpretation or meaning. For example, a name may denote an `Integer` variable, a `Real` variable, a function, a type, etc. A name may have different meanings in different parts of the code, i.e., different scopes. The interpretation of identifiers as names is described in more detail in chapter 5. The meaning of package names is described in more detail in chapter 13.

2.3.3 Modelica Keywords

The following Modelica *keywords* are reserved words and shall not be used as identifiers, except as listed in appendix B.1:

<code>algorithm</code>	<code>discrete</code>	<code>false</code>	<code>loop</code>	<code>pure</code>
<code>and</code>	<code>each</code>	<code>final</code>	<code>model</code>	<code>record</code>
<code>annotation</code>	<code>else</code>	<code>flow</code>	<code>not</code>	<code>redeclare</code>
	<code>elseif</code>	<code>for</code>	<code>operator</code>	<code>replaceable</code>
<code>block</code>	<code>elsewhen</code>	<code>function</code>	<code>or</code>	<code>return</code>
<code>break</code>	<code>encapsulated</code>	<code>if</code>	<code>outer</code>	<code>stream</code>
<code>class</code>	<code>end</code>	<code>import</code>	<code>output</code>	<code>then</code>
<code>connect</code>	<code>enumeration</code>	<code>impure</code>	<code>package</code>	<code>true</code>
<code>connector</code>	<code>equation</code>	<code>in</code>	<code>parameter</code>	<code>type</code>
<code>constant</code>	<code>expandable</code>	<code>initial</code>	<code>partial</code>	<code>when</code>
<code>constrainedby</code>	<code>extends</code>	<code>inner</code>	<code>protected</code>	<code>while</code>
<code>der</code>	<code>external</code>	<code>input</code>	<code>public</code>	<code>within</code>

2.4 Literal Constants

Literal constants are unnamed constants that have different forms depending on their type. Each of the predefined types in Modelica has a way of expressing unnamed constants of the corresponding type, which is presented in the ensuing subsections. Additionally, array literals and record literals can be expressed.

2.4.1 Floating Point Numbers

A floating point number is expressed as a decimal number in the form of a sequence of decimal digits followed by a decimal point, followed by decimal digits, followed by an exponent indicated by **E** or **e** followed by a sign and one or more decimal digits. The various parts can be omitted, see **UNSIGNED-REAL** in appendix B.1 for details and also the examples below. The minimal recommended range is that of IEEE double precision floating point numbers, for which the largest representable positive number is $1.7976931348623157 \times 10^{308}$ and the smallest positive number is $2.2250738585072014 \times 10^{-308}$. For example, the following are floating point number literal constants:

```
22.5, 3.141592653589793, 1.2E-35
```

The same floating point number can be represented by different literals. For example, all of the following literals denote the same number:

```
13., 13E0, 1.3e1, 0.13E2, .13E2
```

The last variant shows that the leading zero is optional (in that case decimal digits must be present). Note that `13` is not in this list, since it is not a floating point number, but can be converted to a floating point number.

2.4.2 Integer Literals

Literals of type **Integer** are sequences of decimal digits, e.g. as in the integer numbers `33`, `0`, `100`, `30030044`. The range of supported **Integer** literals shall be at least large enough to represent the largest positive **IntegerType** value, see section 4.8.2.

[Negative numbers are formed by unary minus followed by an integer literal.]

2.4.3 Boolean Literals

The two **Boolean** literal values are **true** and **false**.

2.4.4 Strings

String literals appear between double quotes as in `"between"`. Any character in the Modelica language character set (see appendix B.1 for allowed characters) apart from double quote (`"`) and backslash (`\`), including new-line, can be *directly* included in a string without using an escape code. Certain characters in string literals can be represented using escape codes, i.e., the character is preceded by a backslash (`\`) within the string. Those characters are:

<i>Character</i>	<i>Description</i>
<code>\'</code>	Single quote, may also appear without backslash in string constants
<code>\"</code>	Double quote
<code>\?</code>	Question-mark, may also appear without backslash in string constants
<code>\\</code>	Backslash itself
<code>\a</code>	Alert (bell, code 7, ctrl-G)
<code>\b</code>	Backspace (code 8, ctrl-H)
<code>\f</code>	Form feed (code 12, ctrl-L)
<code>\n</code>	Newline (code 10, ctrl-J), same as literal newline
<code>\r</code>	Carriage return (code 13, ctrl-M)
<code>\t</code>	Horizontal tab (code 9, ctrl-I)
<code>\v</code>	Vertical tab (code 11, ctrl-K)

For example, a string literal containing a tab, the words: *This is*, double quote, space, the word: *between*, double quote, space, the word: *us*, and new-line, would appear as follows:

```
"\tThis is\" between\" us\n"
```

Concatenation of string literals in certain situations (see the Modelica grammar) is denoted by the + operator in Modelica, e.g. "a" + "b" becomes "ab". This is useful for expressing long string literals that need to be written on several lines.

The "\n" character is used to conceptually indicate the end of a line within a Modelica string. Any Modelica program that needs to recognize line endings can check for a single "\n" character to do so on any platform. It is the responsibility of a Modelica implementation to make any necessary transformations to other representations when writing to or reading from a text file.

[For example, a "\n" is written and read as-is in a Unix or Linux implementation, but written as "\r\n" pair, and converted back to "\n" when read in a Windows implementation.]

[For long string comments, e.g., the info annotation to store the documentation of a model, it would be very inconvenient, if the string concatenation operator would have to be used for every line of documentation. It is assumed that a Modelica tool supports the non-printable newline character when browsing or editing a string literal. For example, the following statement defines one string that contains (non-printable) newline characters:

```
assert(noEvent(length > s_small),
"The distance between the origin of frame_a and the origin of frame_b
of a LineForceWithMass component became smaller as parameter s_small
(= a small number, defined in the
\"Advanced\" menu). The distance is
set to s_small, although it is smaller, to avoid a division by zero
when computing the direction of the line force.",
    level = AssertionLevel.warning);
```

]

2.5 Operator Symbols

The predefined operator symbols are formally defined on page 281 and summarized in the table of operators in section 3.2.

Chapter 3

Operators and Expressions

The lexical units are combined to form even larger building blocks such as expressions according to the rules given by the expression part of the Modelica grammar in appendix B.

This chapter describes the evaluation rules for expressions, the concept of expression variability, built-in mathematical operators and functions, and the built-in special Modelica operators with function syntax.

Expressions can contain variables and constants, which have types, predefined or user defined. The predefined built-in types of Modelica are `Real`, `Integer`, `Boolean`, `String`, and enumeration types which are presented in more detail in section 4.8.

3.1 Expressions

Modelica equations, assignments and declaration equations contain expressions.

Expressions can contain basic operations, `+`, `-`, `*`, `/`, `^`, etc. with normal precedence as defined in the Table in section 3.2 and the grammar in appendix B. The semantics of the operations is defined for both scalar and array arguments in section 10.6.

It is also possible to define functions and call them in a normal fashion. The function call syntax for both positional and named arguments is described in section 12.4.1 and for vectorized calls in section 12.4.4. The built-in array functions are given in section 10.1.1 and other built-in operators in section 3.7.

3.2 Operator Precedence and Associativity

Operator precedence determines the order of evaluation of operators in an expression. An operator with higher precedence is evaluated before an operator with lower precedence in the same expression.

The following table presents all the expression operators in order of precedence.

Table 3.1: Operators in order of precedence from highest to lowest, as derived from the Modelica grammar in appendix B. All operators are binary except the postfix operators and those shown as unary together with `expr`, the conditional operator, the array construction operator `{ }` and concatenation operator `[]`, and the array range constructor which is either binary or ternary. Operators with the same precedence occur at the same table row.

Operator group	Operator syntax	Examples
Postfix array index operator	<code>[]</code>	<code>arr[index]</code>
Postfix access operator	<code>.</code>	<code>a.b</code>
Postfix function call	<code>funcName(functionArguments)</code>	<code>sin(4.36)</code>
Array construction	<code>{expressions}</code>	<code>{2, 3}</code>
Horizontal concatenation	<code>[expressions]</code>	<code>[5, 6]</code>
Vertical concatenation	<code>[expressions; expressions ...]</code>	<code>[2, 3; 7, 8]</code>
Exponentiation	<code>^</code>	<code>2 ^ 3</code>
Multiplicative	<code>* /</code>	<code>2 * 3, 2 / 3</code>
Elementwise multiplicative	<code>.* ./</code>	<code>[1, 2; 3, 4] .* [2, 3; 5, 6]</code>
Additive	<code>+ -</code>	<code>1 + 2</code>
Additive unary	<code>+expr -expr</code>	<code>-0.5</code>
Array elementwise additive	<code>.+ .-</code>	<code>[1, 2; 3, 4] .+ [2, 3; 5, 6]</code>
Relational	<code>< <= > >= == <></code>	<code>a < b, a <= b, a > b, ...</code>
Unary negation	<code>not expr</code>	<code>not b1</code>
Logical and	<code>and</code>	<code>b1 and b2</code>
Logical or	<code>or</code>	<code>b1 or b2</code>
Array range	<code>expr : expr</code> <code>expr : expr : expr</code>	<code>1 : 5</code> <code>start : step : stop</code>
Conditional	<code>if expr then expr else expr</code>	<code>if b then 3 else x</code>
Named argument	<code>ident = expr</code>	<code>x = 2.26</code>

The conditional operator may also include `elseif`-clauses. Equality `=` and assignment `:=` are not expression operators since they are allowed only in equations and in assignment statements respectively. All binary expression operators are left associative, except exponentiation which is non-associative. The array range operator is non-associative.

[The unary minus and plus in Modelica is slightly different than in Mathematica¹ and in MATLAB², since the following expressions are illegal (whereas in Mathematica and in MATLAB these are valid expressions):

```

2*-2 // = -4 in Mathematica/MATLAB; is illegal in Modelica
--2 // = 2 in Mathematica/MATLAB; is illegal in Modelica
++2 // = 2 in Mathematica/MATLAB; is illegal in Modelica
2--2 // = 4 in Mathematica/MATLAB; is illegal in Modelica

```

Non-associative exponentiation and array range operator:

```

x^y^z // Not legal, use parenthesis to make it clear
a:b:c:d:e:f:g // Not legal, and scalar arguments gives no legal interpretation.

```

]

3.3 Evaluation Order

A tool is free to solve equations, reorder expressions and to not evaluate expressions if their values do not influence the result (e.g. short-circuit evaluation of `Boolean` expressions). If-statements and if-expressions guarantee that their clauses are only evaluated if the appropriate condition is true, but relational operators generating state or time events will during continuous integration have the value from the most recent event.

¹ *Mathematica* is a registered trademark of Wolfram Research Inc.

² *MATLAB* is a registered trademark of MathWorks Inc.

If a numeric operation overflows the result is undefined. For literals it is recommended to automatically convert the number to another type with greater precision.

3.3.1 Example: Guarding Expressions Against Incorrect Evaluation

[*Example: If one wants to guard an expression against incorrect evaluation, it should be guarded by an if:*

```

Boolean v[n];
Boolean b;
Integer I;
equation
  b=(I>=1 and I<=n) and v[I]; // Invalid
  b=if (I>=1 and I<=n) then v[I] else false; // Correct

```

To guard square against square root of negative number use `noEvent`:

```

der(h)=if h>0 then -c*sqrt(h) else 0; // Incorrect
der(h)=if noEvent(h>0) then -c*sqrt(h) else 0; // Correct

```

]

3.4 Arithmetic Operators

Modelica supports five binary arithmetic operators that operate on any numerical type:

<i>Operator</i>	<i>Description</i>
<code>^</code>	Exponentiation
<code>*</code>	Multiplication
<code>/</code>	Division
<code>+</code>	Addition
<code>-</code>	Subtraction

Some of these operators can also be applied to a combination of a scalar type and an array type, see section 10.6.

The syntax of these operators is defined by the following rules from the Modelica grammar:

```

arithmetic-expression :
  [ add-operator ] term { add-operator term }

add-operator :
  "+" | "-"

term :
  factor { mul-operator factor }

mul-operator :
  "*" | "/"

factor :
  primary [ "^" primary ]

```

3.5 Equality, Relational, and Logical Operators

Modelica supports the standard set of relational and logical operators, all of which produce the standard boolean values `true` or `false`:

<i>Operator</i>	<i>Description</i>
>	Greater than
>=	Greater than or equal
<	Less than
<=	Less than or equal to
==	Equality within expressions
<>	Inequality

A single equals sign = is never used in relational expressions, only in equations (chapter 8, section 10.6.1) and in function calls using named parameter passing (section 12.4.1).

The following logical operators are defined:

<i>Operator</i>	<i>Description</i>
<code>not</code>	Logical negation (unary operator)
<code>and</code>	Logical <i>and</i> (conjunction)
<code>or</code>	Logical <i>or</i> (disjunction)

The grammar rules define the syntax of the relational and logical operators.

```

logical-expression :
  logical-term { or logical-term }

logical-term :
  logical-factor { and logical-factor }

logical-factor :
  [ not ] relation

relation :
  arithmetic-expression [ relational-operator arithmetic-expression ]

relational-operator :
  "<" | "<=" | ">" | ">=" | "==" | "<>"
  
```

The following holds for relational operators:

- Relational operators <, <=, >, >=, ==, <>, are only defined for scalar operands of simple types. The result is `Boolean` and is true or false if the relation is fulfilled or not, respectively.
- For operands of type `String`, `str1 op str2` is for each relational operator, *op*, defined in terms of the C function `strcmp` as `strcmp(str1, str2) op 0`.
- For operands of type `Boolean`, `false < true`.
- For operands of enumeration types, the order is given by the order of declaration of the enumeration literals.
- In relations of the form `v1 == v2 or v1 <> v2`, `v1` or `v2` shall, unless used in a function, not be a subtype of `Real`.

[The reason for this rule is that relations with `Real` arguments are transformed to state events (see *Events*, section 8.5) and this transformation becomes unnecessarily complicated for the `==` and `<>` relational operators (e.g. two crossing functions instead of one crossing function needed, epsilon strategy needed even at event instants). Furthermore, testing on equality of `Real` variables is questionable on machines where the number length in registers is different to number length in main memory.]

- Relational operators can generate events, see section 3.8.3.

3.6 Miscellaneous Operators and Variables

Modelica also contains a few built-in operators which are not standard arithmetic, relational, or logical operators. These are described below, including `time`, which is a built-in variable, not an operator.

3.6.1 String Concatenation

Concatenation of strings (see the Modelica grammar) is denoted by the + operator in Modelica.

[*Example: "a" + "b" becomes "ab".*]

3.6.2 Array Constructor Operator

The array constructor operator { ... } is described in section 10.4.

3.6.3 Array Concatenation Operator

The array concatenation operator [...] is described in section 10.4.2.

3.6.4 Array Range Operator

The array range constructor operator : is described in section 10.4.3.

3.6.5 If-Expressions

An expression

```
if expression1 then expression2 else expression3
```

is one example of if-expression. First `expression1`, which must be `Boolean` expression, is evaluated. If `expression1` is true `expression2` is evaluated and is the value of the if-expression, else `expression3` is evaluated and is the value of the if-expression. The two expressions, `expression2` and `expression3`, must be type compatible expressions (section 6.7) giving the type of the if-expression. If-expressions with `elseif` are defined by replacing `elseif` by `else if`. For short-circuit evaluation see section 3.3.

[*elseif in expressions has been added to the Modelica language for symmetry with if-clauses.*]

[*Example:*

```
Integer i;
Integer sign_of_i1=if i<0 then -1 elseif i==0 then 0 else 1;
Integer sign_of_i2=if i<0 then -1 else if i==0 then 0 else 1;
```

]

3.6.6 Member Access Operator

It is possible to access members of a class instance using dot notation, i.e., the . operator.

[*Example: R1.R for accessing the resistance component R of resistor R1. Another use of dot notation: local classes which are members of a class can of course also be accessed using dot notation on the name of the class, not on instances of the class.*]

3.6.7 Built-in Variable time

All declared variables are functions of the independent variable `time`. The variable `time` is a built-in variable available in all models and blocks, which is treated as an input variable. It is implicitly defined as:

```
input Real time (final quantity = "Time",
                final unit = "s");
```

The value of the `start` attribute of `time` is set to the time instant at which the simulation is started.

[*Example:*

```
encapsulated model SineSource
  import Modelica.Math.sin;
  connector OutPort=output Real;
  OutPort y=sin(time); // Uses the built-in variable time.
end SineSource;
```

]

3.7 Built-in Intrinsic Operators with Function Syntax

Certain built-in operators of Modelica have the same syntax as a function call. However, they do not behave as a mathematical function, because the result depends not only on the input arguments but also on the status of the simulation.

There are also built-in functions that depend only on the input argument, but also may trigger events in addition to returning a value. Intrinsic means that they are defined at the Modelica language level, not in the Modelica library. The following built-in intrinsic operators/functions are available:

- Mathematical functions and conversion functions, see section 3.7.1 below.
- Derivative and special purpose operators with function syntax, see section 3.7.4 below.
- Event-related operators with function syntax, see section 3.7.5 below.
- Array operators/functions, see section 10.1.1.

Note that when the specification references a function having the name of a built-in function it references the built-in function, not a user-defined function having the same name, see also section 12.5. With exception of the built-in `String` operator, all operators in this section can only be called with positional arguments.

3.7.1 Numeric Functions and Conversion Functions

The mathematical functions and conversion operators are listed below do not generate events.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>abs(v)</code>	Absolute value (event-free)	Function 3.1
<code>sign(v)</code>	Sign of argument (event-free)	Function 3.2
<code>sqrt(v)</code>	Square root	Function 3.3
<code>Integer(e)</code>	Conversion from enumeration to <code>Integer</code>	Operator 3.1
<code>EnumTypeName(i)</code>	Conversion from <code>Integer</code> to enumeration	Operator 3.2
<code>String(...)</code>	Conversion to <code>String</code>	Operator 3.3

All of these except for the `String` conversion operator are vectorizable according to section 12.4.6.

Additional non-event generating mathematical functions are described in section 3.7.3, whereas the event-triggering mathematical functions are described in section 3.7.2.

Function 3.1 abs

`abs(v)`

Expands into `noEvent(if v >= 0 then v else -v)`. Argument *v* needs to be an `Integer` or `Real` expression.

Function 3.2 sign

`sign(v)`

Expands into `noEvent(if v > 0 then 1 else if v < 0 then -1 else 0)`. Argument *v* needs to be an `Integer` or `Real` expression.

Function 3.3 sqrt

`sqrt(v)`

Square root of *v* if $v \geq 0$, otherwise an error occurs. Argument *v* needs to be an `Integer` or `Real` expression.

Operator 3.1 Integer

`Integer(e)`

Ordinal number of the expression e of enumeration type that evaluates to the enumeration value $E.enumvalue$, where $Integer(E.e1) = 1$, $Integer(E.en) = n$, for an enumeration type $E = enumeration(e1, \dots, en)$. See also section 4.8.5.2.

Operator 3.2 <EnumTypeName>

`EnumTypeName (i)`

For any enumeration type `EnumTypeName`, returns the enumeration value `EnumTypeName.e` such that $Integer(EnumTypeName.e) = i$. Refer to the definition of `Integer` above.

It is an error to attempt to convert values of i that do not correspond to values of the enumeration type. See also section 4.8.5.3.

Operator 3.3 String

`String(b, <options>)`
`String(i, <options>)`
`String(r, significantDigits=d, <options>)`
`String(r, format=s)`
`String(e, <options>)`

Convert a scalar non-`String` expression to a `String` representation. The first argument may be a `Boolean` b , an `Integer` i , a `Real` r or an enumeration value e (section 4.8.5.2). The other arguments must use named arguments. For `Real` expressions the output shall be according to the Modelica grammar.

The optional `<options>` are:

- `Integer minimumLength = 0`: Minimum length of the resulting string. If necessary, the blank character is used to fill up unused space.
- `Boolean leftJustified = true`: If true, the converted result is left justified in the string; if false it is right justified in the string.
- `Integer significantDigits = 6`: Number of significant digits in the result string.

[*Examples of Real values formatted with 6 significant digits:* 12.3456, 0.0123456, 12345600, 1.23456E-10.]

The format string corresponding to `<options>` is:

- For `Real`:
`(if leftJustified then "-" else "") + String(minimumLength)`
`+ "." + String(significantDigits) + "g"`
- For `Integer`:
`(if leftJustified then "-" else "") + String(minimumLength) + "d"`

Form of the `format` string: According to ANSI-C the format string specifies one conversion specifier (excluding the leading `%`), shall not contain length modifiers, and shall not use `*` for width and/or precision. For all numeric values the format specifiers `'f'`, `'e'`, `'E'`, `'g'`, `'G'` are allowed. For integral values it is also allowed to use the `'d'`, `'i'`, `'o'`, `'x'`, `'X'`, `'u'`, and `'c'` format specifiers (for non-integral values a tool may round, truncate or use a different format if the integer conversion characters are used).

The `'x'`/`'X'` formats (hexa-decimal) and `c` (character) for `Integer` values give results that do not agree with the Modelica grammar.

3.7.2 Event Triggering Mathematical Functions

The operators listed below trigger events if used outside of a `when`-clause and outside of a clocked discrete-time partition (see section 16.8.1).

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>div(x, y)</code>	Division with truncation toward zero	Operator 3.4
<code>mod(x, y)</code>	Integer modulus	Operator 3.5
<code>rem(x, y)</code>	Integer remainder	Operator 3.6
<code>ceil(x)</code>	Smallest integer Real not less than x	Operator 3.7
<code>floor(x)</code>	Largest integer Real not greater than x	Operator 3.8
<code>integer(x)</code>	Largest Integer not greater than x	Operator 3.9

These expression for `div`, `ceil`, `floor`, and `integer` are event generating expression. The event generating expression for `mod(x,y)` is `floor(x/y)`, and for `rem(x,y)` it is `div(x,y) - i.e.` events are not generated when `mod` or `rem` changes continuously in an interval, but when they change discontinuously from one interval to the next.

[If this is not desired, the `noEvent` operator can be applied to them. E.g. `noEvent(integer(v))`.]

Operator 3.4 `div`

`div(x, y)`

Algebraic quotient x/y with any fractional part discarded (also known as truncation toward zero).

[This is defined for `/` in C99; in C89 the result for negative numbers is implementation-defined, so the standard function `div` must be used.]

Result and arguments shall have type **Real** or **Integer**. If either of the arguments is **Real** the result is **Real** otherwise **Integer**.

Operator 3.5 `mod`

`mod(x, y)`

Integer modulus of x/y , i.e. `mod(x, y) = x - floor(x / y) * y`. Result and arguments shall have type **Real** or **Integer**. If either of the arguments is **Real** the result is **Real** otherwise **Integer**.

[Note, outside of a when-clause state events are triggered when the return value changes discontinuously. Examples: `mod(3, 1.4) = 0.2`, `mod(-3, 1.4) = 1.2`, `mod(3, -1.4) = -1.2`.]

Operator 3.6 `rem`

`rem(x, y)`

Integer remainder of x/y , such that `div(x, y) * y + rem(x, y) = x`. Result and arguments shall have type **Real** or **Integer**. If either of the arguments is **Real** the result is **Real** otherwise **Integer**.

[Note, outside of a when-clause state events are triggered when the return value changes discontinuously. Examples: `rem(3, 1.4) = 0.2`, `rem(-3, 1.4) = -0.2`.]

Operator 3.7 `ceil`

`ceil(x)`

Smallest integer not less than x . Result and argument shall have type **Real**.

[Note, outside of a when-clause state events are triggered when the return value changes discontinuously.]

Operator 3.8 `floor`

`floor(x)`

Largest integer not greater than x . Result and argument shall have type **Real**.

[Note, outside of a when-clause state events are triggered when the return value changes discontinuously.]

Operator 3.9 `integer`

`integer(x)`

Largest integer not greater than x . The argument shall have type `Real`. The result has type `Integer`.

[*Note, outside of a when-clause state events are triggered when the return value changes discontinuously.*]

3.7.3 Elementary Mathematical Functions

The functions listed below are elementary mathematical functions. Tools are expected to utilize well known properties of these functions (derivatives, inverses, etc) for symbolic processing of expressions and equations.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>sin(x)</code>	Sine	
<code>cos(x)</code>	Cosine	
<code>tan(x)</code>	Tangent (x shall not be: $\dots, -\pi/2, \pi/2, 3\pi/2, \dots$)	
<code>asin(x)</code>	Inverse sine ($-1 \leq x \leq 1$)	
<code>acos(x)</code>	Inverse cosine ($-1 \leq x \leq 1$)	
<code>atan(x)</code>	Inverse tangent	
<code>atan2(y, x)</code>	Principal value of the arc tangent of y/x	Function 3.4
<code>sinh(x)</code>	Hyperbolic sine	
<code>cosh(x)</code>	Hyperbolic cosine	
<code>tanh(x)</code>	Hyperbolic tangent	
<code>exp(x)</code>	Exponential, base e	
<code>log(x)</code>	Natural (base e) logarithm ($x > 0$)	
<code>log10(x)</code>	Base 10 logarithm ($x > 0$)	

These functions are the only ones that can also be called using the deprecated "builtin" external language, see section 12.9.

[*End user oriented information about the elementary mathematical functions can be found for the corresponding functions in the Modelica.Math package.*]

Function 3.4 atan2

`atan2(y, x)`

Principal value of the arc tangent of y/x , using the signs of the two arguments to determine the quadrant of the result. The result φ is in the interval $[-\pi, \pi]$ and satisfies:

$$\begin{aligned} |(x, y)| \cos(\varphi) &= x \\ |(x, y)| \sin(\varphi) &= y \end{aligned}$$

3.7.4 Derivative and Special Purpose Operators with Function Syntax

The operators listed below include the derivative operator and special purpose operators with function syntax.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>der(expr)</code>	Time derivative	Operator 3.10
<code>delay(expr, ...)</code>	Time delay	Operator 3.11
<code>cardinality(c)</code>	Number of occurrences in connect-equations	Operator 3.12
<code>homotopy(actual, simplified)</code>	Homotpy initialization	Operator 3.13
<code>semiLinear(x, k⁺, k⁻)</code>	Sign-dependent slope	Operator 3.14
<code>inStream(v)</code>	Stream variable flow into component	Operator 3.15
<code>actualStream(v)</code>	Actual value of stream variable	Operator 3.16
<code>spatialDistribution(...)</code>	Variable-speed transport	Operator 3.17
<code>getInstanceName()</code>	Name of instance at call site	Operator 3.18

The special purpose operators with function syntax where the call below uses named arguments can be called with named arguments (with the specified names), or with positional arguments (the inputs of the functions are in the order given in the calls below).

Operator 3.10 `der`

```
der(expr)
```

The time derivative of *expr*. If the expression *expr* is a scalar it needs to be a subtype of `Real`. The expression and all its time-varying subexpressions must be continuous and semi-differentiable. If *expr* is an array, the operator is applied to all elements of the array. For non-scalar arguments the function is vectorized according to section 10.6.12.

[For `Real` parameters and constants the result is a zero scalar or array of the same size as the variable.]

Operator 3.11 `delay`

```
delay(expr, delayTime, delayMax)
delay(expr, delayTime)
```

Evaluates to $expr(\text{time} - \text{delayTime})$ for $\text{time} > \text{time.start} + \text{delayTime}$ and $expr(\text{time.start})$ for $\text{time} \leq \text{time.start} + \text{delayTime}$. The arguments, i.e., *expr*, *delayTime* and *delayMax*, need to be subtypes of `Real`. *delayMax* needs to be additionally a parameter expression. The following relation shall hold: $0 \leq \text{delayTime} \leq \text{delayMax}$, otherwise an error occurs. If *delayMax* is not supplied in the argument list, *delayTime* needs to be a parameter expression. For non-scalar arguments the function is vectorized according to section 10.6.12. For further details, see section 3.7.4.1.

Operator 3.12 `cardinality`

```
cardinality(c)
```

[This is a deprecated operator. It should no longer be used, since it will be removed in one of the next Modelica releases.]

Returns the number of (inside and outside) occurrences of connector instance *c* in a connect-equation as an `Integer` number. For further details, see section 3.7.4.3.

Operator 3.13 `homotopy`

```
homotopy(actual=actual, simplified=simplified)
```

The scalar expressions *actual* and *simplified* are subtypes of `Real`. A Modelica translator should map this operator into either of the two forms:

1. Returns *actual* (trivial implementation).
2. In order to solve algebraic systems of equations, the operator might during the solution process return a combination of the two arguments, ending at *actual*.

[Example: $actual \cdot \lambda + simplified \cdot (1 - \lambda)$, where λ is a homotopy parameter going from 0 to 1.]

The solution must fulfill the equations for `homotopy` returning *actual*.

For non-scalar arguments the function is vectorized according to section 12.4.6. For further details, see section 3.7.4.4.

Operator 3.14 `semiLinear`

```
semiLinear(x, k+, k-)
```

Returns: `smooth(0, if x >= 0 then k+ * x else k- * x)`. The result is of type `Real`. For non-scalar arguments the function is vectorized according to section 10.6.12. For further details, see section 3.7.4.5 (especially in the case when $x = 0$).

Operator 3.15 `inStream`

```
inStream(v)
```

`inStream(v)` is only allowed for stream variables *v* defined in stream connectors, and is the value of the stream variable *v* close to the connection point assuming that the flow is from the connection point into the component. This value is computed from the stream connection equations of the

flow variables and of the stream variables. The operator is vectorizable. For further details, see section 15.2.

Operator 3.16 `actualStream`

```
actualStream(v)
```

`actualStream(v)` returns the actual value of the stream variable v for any flow direction. The operator is vectorizable. For further details, see section 15.3.

Operator 3.17 `spatialDistribution`

```
spatialDistribution(
  in0=in0, in1=in1, x=x,
  positiveVelocity=...,
  initialPoints=...,
  initialValues=...)
```

`spatialDistribution` allows approximation of variable-speed transport of properties. For further details, see section 3.7.4.2.

Operator 3.18 `getInstanceName`

```
getInstanceName()
```

Returns a string with the name of the model/block that is simulated, appended with the fully qualified name of the instance in which this function is called. For further details, see section 3.7.4.6.

A few of these operators are described in more detail in the following.

3.7.4.1 `delay`

[`delay` allows a numerical sound implementation by interpolating in the (internal) integrator polynomials, as well as a more simple realization by interpolating linearly in a buffer containing past values of expression $expr$. Without further information, the complete time history of the delayed signals needs to be stored, because the delay time may change during simulation. To avoid excessive storage requirements and to enhance efficiency, the maximum allowed delay time has to be given via `delayMax`.

This gives an upper bound on the values of the delayed signals which have to be stored. For real-time simulation where fixed step size integrators are used, this information is sufficient to allocate the necessary storage for the internal buffer before the simulation starts. For variable step size integrators, the buffer size is dynamic during integration. In principle, `delay` could break algebraic loops. For simplicity, this is not supported because the minimum delay time has to be given as additional argument to be fixed at compile time. Furthermore, the maximum step size of the integrator is limited by this minimum delay time in order to avoid extrapolation in the delay buffer.]

3.7.4.2 `spatialDistribution`

[Many applications involve the modelling of variable-speed transport of properties. One option to model this infinite-dimensional system is to approximate it by an ODE, but this requires a large number of state variables and might introduce either numerical diffusion or numerical oscillations. Another option is to use a built-in operator that keeps track of the spatial distribution of $z(x,t)$, by suitable sampling, interpolation, and shifting of the stored distribution. In this case, the internal state of the operator is hidden from the ODE solver.]

`spatialDistribution` allows the infinite-dimensional problem below to be solved efficiently with good accuracy

$$\begin{aligned} \frac{\partial z(y,t)}{\partial t} + v(t) \frac{\partial z(y,t)}{\partial y} &= 0.0 \\ z(0.0,t) &= in_0(t) \text{ if } v \geq 0 \\ z(1.0,t) &= in_1(t) \text{ if } v < 0 \end{aligned}$$

where $z(y,t)$ is the transported quantity, y is the normalized spatial coordinate ($0.0 \leq y \leq 1.0$), t is the time, $v(t) = \text{der}(x)$ is the normalized transport velocity and the boundary conditions are set at either

$y = 0.0$ or $y = 1.0$, depending on the sign of the velocity. The calling syntax is:

```
(out0, out1) = spatialDistribution(in0, in1, x, positiveVelocity,
                                initialPoints = {0.0, 1.0},
                                initialValues = {0.0, 0.0});
```

where `in0`, `in1`, `out0`, `out1`, `x`, `v` are all subtypes of `Real`, `positiveVelocity` is a `Boolean`, `initialPoints` and `initialValues` are arrays of subtypes of `Real` of equal size, containing the y coordinates and the z values of a finite set of points describing the initial distribution of $z(y, t_0)$. The `out0` and `out1` are given by the solutions at $z(0.0, t)$ and $z(1.0, t)$; and `in0` and `in1` are the boundary conditions at $z(0.0, t)$ and $z(1.0, t)$ (at each point in time only one of `in0` and `in1` is used). Elements in the `initialPoints` array must be sorted in non-descending order. The operator can not be vectorized according to the vectorization rules described in section 12.4.6. The operator can be vectorized only with respect to the arguments `in0` and `in1` (which must have the same size), returning vectorized outputs `out0` and `out1` of the same size; the arguments `initialPoints` and `initialValues` are vectorized accordingly.

The solution, z , can be described in terms of characteristics:

$$z\left(y + \int_t^{t+\beta} v(\alpha) d\alpha, t + \beta\right) = z(y, t), \quad \text{for all } \beta \text{ as long as staying inside the domain}$$

This allows the direct computation of the solution based on interpolating the boundary conditions.

`spatialDistribution` can be described in terms of the pseudo-code given as a block:

```
block spatialDistribution
  input Real in0;
  input Real in1;
  input Real x;
  input Boolean positiveVelocity;
  parameter Real initialPoints(each min=0, each max=1)[:]= {0.0, 1.0};
  parameter Real initialValues[:]= {0.0, 0.0};
  output Real out0;
  output Real out1;
protected
  Real points[:];
  Real values[:];
  Real x0;
  Integer m;
algorithm
  /* The notation
   * x <and then> y
   * is used below as a shorthand for
   * if x then y else false
   * also known as "short-circuit evaluation of x and y".
   */
  if positiveVelocity then
    out1 := interpolate(points, values, 1 - (x - x0));
    out0 := values[1]; // Similar to in0 but avoiding algebraic loop.
  else
    out0 := interpolate(points, values, 0 - (x - x0));
    out1 := values[end]; // Similar to in1 but avoiding algebraic loop.
  end if;
  when <acceptedStep> then
    if x > x0 then
      m := size(points, 1);
      while m > 0 <and then> points[m] + (x - x0) >= 1 loop
        m := m - 1;
      end while;
      values := cat(1,
                   {in0},
                   values[1:m],
                   {interpolate(points, values, 1 - (x - x0))});
      points := cat(1, {0}, points[1:m] .+ (x-x0), {1});
```

```

elseif x < x0 then
  m := 1;
  while m < size(points, 1) <and then> points[m] + (x - x0) <= 0 loop
    m := m + 1;
  end while;
  values := cat(1,
                {interpolate(points, values, 0 - (x - x0))},
                values[m:end],
                {in1});
  points := cat(1, {0}, points[m:end] .+ (x - x0), {1});
end if;
x0 := x;
end when;
initial algorithm
  x0 := x;
  points := initialPoints;
  values := initialValues;
end spatialDistribution;
    
```

[Note that the implementation has an internal state and thus cannot be described as a function in Mod-*elica*; `initialPoints` and `initialValues` are declared as parameters to indicate that they are only used during initialization.

The infinite-dimensional problem stated above can then be formulated in the following way:

```

der(x) = v;
(out0, out1) = spatialDistribution(in0, in1, x, v >= 0,
                                  initialPoints, initialValues);
    
```

Events are generated at the exact instants when the velocity changes sign – if this is not needed, `noEvent` can be used to suppress event generation.

If the velocity is known to be always positive, then `out0` can be omitted, e.g.:

```

der(x) = v;
(, out1) = spatialDistribution(in0, 0, x, true, initialPoints, initialValues);
    
```

Technically relevant use cases for the use of `spatialDistribution` are modeling of electrical transmission lines, pipelines and pipeline networks for gas, water and district heating, sprinkler systems, impulse propagation in elongated bodies, conveyor belts, and hydraulic systems. Vectorization is needed for pipelines where more than one quantity is transported with velocity v in the example above.]

3.7.4.3 cardinality (deprecated)

[`cardinality` is deprecated for the following reasons and will be removed in a future release:

- Reflective operator may make early type checking more difficult.
- Almost always abused in strange ways
- Not used for Bond graphs even though it was originally introduced for that purpose.

]

[`cardinality` allows the definition of connection dependent equations in a model, for example:

```

connector Pin
  Real v;
  flow Real i;
end Pin;
model Resistor
  Pin p, n;
equation
  assert(cardinality(p) > 0 and cardinality(n) > 0,
         "Connectors p and n of Resistor must be connected");
  // Equations of resistor
  ...
    
```

```
end Resistor;
```

The cardinality is counted after removing conditional components, and shall not be applied to expandable connectors, elements in expandable connectors, or to arrays of connectors (but can be applied to the scalar elements of array of connectors). `cardinality` should only be used in the condition of `assert` and `if`-statements that do not contain `connect` and similar operators, see section 16.8.1).

3.7.4.4 homotopy

[During the initialization phase of a dynamic simulation problem, it often happens that large nonlinear systems of equations must be solved by means of an iterative solver. The convergence of such solvers critically depends on the choice of initial guesses for the unknown variables. The process can be made more robust by providing an alternative, simplified version of the model, such that convergence is possible even without accurate initial guess values, and then by continuously transforming the simplified model into the actual model. This transformation can be formulated using expressions of this kind:

$$\lambda \cdot \text{actual} + (1 - \lambda) \cdot \text{simplified}$$

in the formulation of the system equations, and is usually called a homotopy transformation. If the simplified expression is chosen carefully, the solution of the problem changes continuously with λ , so by taking small enough steps it is possible to eventually obtain the solution of the actual problem.

The operator can be called with ordered arguments or preferably with named arguments for improved readability.

It is recommended to perform (conceptually) one homotopy iteration over the whole model, and not several homotopy iterations over the respective non-linear algebraic equation systems. The reason is that the following structure can be present:

```
w = f1(x) // has homotopy
0 = f2(der(x), x, z, w)
```

Here, a non-linear equation system f_2 is present. `homotopy` is, however used on a variable that is an “input” to the non-linear algebraic equation system, and modifies the characteristics of the non-linear algebraic equation system. The only useful way is to perform the homotopy iteration over f_1 and f_2 together.

The suggested approach is “conceptual”, because more efficient implementations are possible, e.g. by determining the smallest iteration loop, that contains the equations of the first BLT block in which `homotopy` is present and all equations up to the last BLT block that describes a non-linear algebraic equation system.

A trivial implementation of `homotopy` is obtained by defining the following function in the global scope:

```
function homotopy
  input Real actual;
  input Real simplified;
  output Real y;
algorithm
  y := actual;
  annotation(Inline = true);
end homotopy;
```

[Example 1: In electrical systems it is often difficult to solve non-linear algebraic equations if switches are part of the algebraic loop. An idealized diode model might be implemented in the following way, by starting with a “flat” diode characteristic and then move with `homotopy` to the desired “steep” characteristic:

```
model IdealDiode
  ...
  parameter Real Goff = 1e-5;
protected
```

```

Real Goff_flat = max(0.01, Goff);
Real Goff2;
equation
  off = s < 0;
  Goff2 = homotopy(actual=Goff, simplified=Goff_flat);
  u = s*(if off then 1 else Ron2) + Vknee;
  i = s*(if off then Goff2 else 1 ) + Goff2*Vknee;
  ...
end IdealDiode;

```

[Example 2: In electrical systems it is often useful that all voltage sources start with zero voltage and all current sources with zero current, since steady state initialization with zero sources can be easily obtained. A typical voltage source would then be defined as:

```

model ConstantVoltageSource
  extends Modelica.Electrical.Analog.Interfaces.OnePort;
  parameter Modelica.Units.SI.Voltage V;
equation
  v = homotopy(actual=V, simplified=0.0);
end ConstantVoltageSource;

```

[Example 3: In fluid system modelling, the pressure/flowrate relationships are highly nonlinear due to the quadratic terms and due to the dependency on fluid properties. A simplified linear model, tuned on the nominal operating point, can be used to make the overall model less nonlinear and thus easier to solve without accurate start values. Named arguments are used here in order to further improve the readability.

```

model PressureLoss
  import Modelica.Units.SI;
  ...
  parameter SI.MassFlowRate m_flow_nominal "Nominal mass flow rate";
  parameter SI.Pressure dp_nominal "Nominal pressure drop";
  SI.Density rho "Upstream density";
  SI.DynamicViscosity lambda "Upstream viscosity";
equation
  ...
  m_flow = homotopy(actual = turbulentFlow_dp(dp, rho, lambda),
    simplified = dp/dp_nominal*m_flow_nominal);
  ...
end PressureLoss;

```

[Example 4: Note that `homotopy` shall not be used to combine unrelated expressions, since this can generate singular systems from combining two well-defined systems.

```

model DoNotUse
  Real x;
  parameter Real x0 = 0;
equation
  der(x) = 1-x;
initial equation
  0 = homotopy(der(x), x - x0);
end DoNotUse;

```

The initial equation is expanded into

$$0 = \lambda * \text{der}(x) + (1 - \lambda)(x - x_0)$$

and you can solve the two equations to give

$$x = \frac{\lambda + (\lambda - 1)x_0}{2\lambda - 1}$$

which has the correct value of x_0 at $\lambda = 0$ and of 1 at $\lambda = 1$, but unfortunately has a singularity at $\lambda = 0.5$.]

3.7.4.5 semiLinear

(See definition of `semiLinear` in section 3.7.4). In some situations, equations with `semiLinear` become underdetermined if the first argument (`x`) becomes zero, i.e., there are an infinite number of solutions. It is recommended that the following rules are used to transform the equations during the translation phase in order to select one meaningful solution in such cases:

- The equations

```
y = semiLinear(x, sa, s1);
y = semiLinear(x, s1, s2);
y = semiLinear(x, s2, s3);
...
y = semiLinear(x, sN, sb);
...
```

may be replaced by

```
s1 = if x >= 0 then sa else sb
s2 = s1;
s3 = s2;
...
sN = sN-1;
y = semiLinear(x, sa, sb);
```

- The equations

```
x = 0;
y = 0;
y = semiLinear(x, sa, sb);
```

may be replaced by

```
x = 0
y = 0;
sa = sb;
```

[For symbolic transformations, the following property is useful (this follows from the definition):

```
semiLinear(m_flow, port_h, h);
```

is identical to:

```
-semiLinear(-m_flow, h, port_h);
```

The `semiLinear` function is designed to handle reversing flow in fluid systems, such as

```
H_flow = semiLinear(m_flow, port.h, h);
```

i.e., the enthalpy flow rate `H_flow` is computed from the mass flow rate `m_flow` and the upstream specific enthalpy depending on the flow direction.]

3.7.4.6 getInstanceName

Returns a string with the name of the model/block that is simulated, appended with the fully qualified name of the instance in which this function is called.

[Example:

```
package MyLib
  model Vehicle
    Engine engine;
    ...
  end Vehicle;
```

```

model Engine
  Controller controller;
  ...
end Engine;
model Controller
equation
  Modelica.Utilities.Streams.print("Info from: " + getInstanceName());
end Controller;
end MyLib;
    
```

If `MyLib.Vehicle` is simulated, the call of `getInstanceName()` returns `"Vehicle.engine.controller"`.]

If this function is not called inside a model or block (e.g. the function is called in a function or in a constant of a package), the return value is not specified.

[The simulation result should not depend on the return value of this function.]

3.7.5 Event-Related Operators with Function Syntax

The operators listed below are event-related operators with function syntax. The operators `noEvent`, `pre`, `edge`, and `change`, are vectorizable according to section 12.4.6.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>initial()</code>	Predicate for the initialization phase	Operator 3.19
<code>terminal()</code>	Predicate for the end of a successful analysis	Operator 3.20
<code>noEvent(expr)</code>	Evaluate <i>expr</i> without triggering events	Operator 3.21
<code>smooth(p, expr)</code>	Treat <i>expr</i> as <i>p</i> times continuously differentiable	Operator 3.22
<code>sample(start, interval)</code>	Periodic triggering of events	Operator 3.23
<code>pre(y)</code>	Left limit $y(t^-)$ of variable $y(t)$	Operator 3.24
<code>edge(b)</code>	Expands into $(b \text{ and not pre}(b))$	Operator 3.25
<code>change(v)</code>	Expands into $(v <> \text{pre}(v))$	Operator 3.26
<code>reinit(x, expr)</code>	Reinitialize <i>x</i> with <i>expr</i>	Operator 3.27

Operator 3.19 `initial`

```
initial()
```

Returns `true` during the initialization phase and `false` otherwise.

[Hereby, `initial()` triggers a time event at the beginning of a simulation.]

Operator 3.20 `terminal`

```
terminal()
```

Returns `true` at the end of a successful analysis.

[Hereby, `terminal()` ensures an event at the end of successful simulation.]

Operator 3.21 `noEvent`

```
noEvent(expr)
```

Real elementary relations within *expr* are taken literally, i.e., no state or time event is triggered. No zero crossing functions shall be used to monitor any of the normally event-generating subexpressions inside *expr*. See also operator 3.22 `smooth` and section 8.5.

Operator 3.22 `smooth`

```
smooth(p, expr)
```

If $p \geq 0$ `smooth(p, expr)` returns *expr* and states that *expr* is *p* times continuously differentiable, i.e.: *expr* is continuous in all `Real` variables appearing in the expression and all partial derivatives with respect to all appearing real variables exist and are continuous up to order *p*. The argument *p* should be a scalar `Integer` parameter expression. The only allowed types for *expr* in `smooth`

are: `Real` expressions, arrays of allowed expressions, and records containing only components of allowed expressions.

`smooth` should be used instead of `noEvent` in order to avoid events for efficiency reasons. A tool is free to not generate events for expressions inside `smooth`. However, `smooth` does not guarantee that no events will be generated, and thus it can be necessary to use `noEvent` inside `smooth`.

[*Note that `smooth` does not guarantee a smooth output if any of the occurring variables change discontinuously.*]

[*Example:*

```

Real x, y, z;
parameter Real p;
equation
x = if time < 1 then 2 else time - 2;
z = smooth(0, if time < 0 then 0 else time);
y = smooth(1, noEvent(if x < 0 then 0 else sqrt(x) * x)); // noEvent is
  necessary.

```

]

Operator 3.23 `sample`

```
sample(start, interval)
```

Returns `true` and triggers time events at time instants $start + i \cdot interval$ for $i = 0, 1 \dots$, and is only true during the first event iteration at those times. At event iterations after the first one at each event and during continuous integration the operator always returns `false`. The starting time *start* and the sample interval *interval* must be parameter expressions and need to be a subtype of `Real` or `Integer`. The sample interval *interval* must be a positive number.

Operator 3.24 `pre`

```
pre(y)
```

Returns the *left limit* $y(t^-)$ of variable $y(t)$ at a time instant t . At an event instant, $y(t^-)$ is the value of y after the last event iteration at time instant t (see comment below). Any subscripts in the component expression y must be parameter expressions. `pre` can be applied if the following three conditions are fulfilled simultaneously: (a) variable y is either a subtype of a simple type or is a record component, (b) y is a discrete-time expression (c) the operator is *not* applied in a `function` class.

[*This can be applied to continuous-time variables in when-clauses, see section 3.8.3 for the definition of discrete-time expression.*]

The first value of `pre(y)` is determined in the initialization phase.

A new event is triggered if there is at least for one variable v such that `pre(v) <> v` after the active model equations are evaluated at an event instant. In this case the model is at once reevaluated. This evaluation sequence is called *event iteration*. The integration is restarted once `pre(v) == v` for all v appearing inside `pre(...)`.

[*If v and `pre(v)` are only used in when-clauses, the translator might mask event iteration for variable v since v cannot change during event iteration. It is a quality of implementation to find the minimal loops for event iteration, i.e., not all parts of the model need to be reevaluated.*]

The language allows mixed algebraic systems of equations where the unknown variables are of type `Real`, `Integer`, `Boolean`, or an enumeration. These systems of equations can be solved by a global fix point iteration scheme, similarly to the event iteration, by fixing the `Boolean`, `Integer`, and/or enumeration unknowns during one iteration. Again, it is a quality of implementation to solve these systems more efficiently, e.g., by applying the fix point iteration scheme to a subset of the model equations.

Operator 3.25 `edge`

```
edge(b)
```

Expands into $(b \text{ and not pre}(b))$ for Boolean variable b . The same restrictions as for `pre` apply (e.g. not to be used in `function` classes).

Operator 3.26 `change`

`change(v)`

Expands into $(v <> \text{pre}(v))$. The same restrictions as for `pre` apply.

Operator 3.27 `reinit`

`reinit(x, expr)`

In the body of a when clause, reinitializes x with $expr$ at an event instant. x is a scalar or array `Real` variable that is implicitly defined to have `StateSelect.always`.

[It is an error if the variable cannot be selected as a state.]

$expr$ needs to be type-compatible with x . `reinit` can only be applied once for the same variable – either as an individual variable or as part of an array of variables. It can only be applied in the body of a when clause in an equation section. See also section 8.3.6.

3.8 Variability of Expressions

The concept of variability of an expression indicates to what extent the expression can vary over time. See also section 4.4.4 regarding the concept of variability. There are four levels of variability of expressions, starting from the least variable:

- constant variability
- parameter variability
- discrete-time variability
- continuous-time variability

While many invalid models can be rejected based on the declared variabilities of variables alone (without the concept of expression variability), the following rules both help enforcing compliance of computed solutions to declared variability, and impose additional restrictions that simplify reasoning and reporting of errors:

- For an assignment $v := \text{expr}$ or binding equation $v = \text{expr}$, v must be declared to be at least as variable as expr .
- When determining whether an equation can contribute to solving for a variable v (for instance, when applying the perfect matching rule, see section 8.4), the equation can only be considered contributing if the resulting solution would be at most as variable as v .
- The right-hand side expression in a binding equation (that is, expr) of a parameter component and of the base type attributes (such as `start`) needs to be a parameter or constant expression.
- If v is a discrete-time component then expr needs to be a discrete-time expression.

3.8.1 Constant Expressions

Constant expressions are:

- `Real`, `Integer`, `Boolean`, `String`, and `enumeration` literals.
- Variables declared as `constant`.
- Except for the special built-in operators `initial`, `terminal`, `der`, `edge`, `change`, `sample`, and `pre`, a function or operator with constant subexpressions as argument (and no parameters defined in the function) is a constant expression.
- Some function calls are constant expressions regardless of the arguments:
 - `ndims(A)`

Components declared as `constant` shall have an associated declaration equation with a constant expression, if the constant is directly in the simulation model, or used in the simulation model. The value of a constant can be modified after it has been given a value, unless the constant is declared `final` or modified with a `final` modifier. A constant without an associated declaration equation can be given one by using a modifier.

3.8.2 Parameter Expressions

Parameter expressions are:

- Constant expressions.
- Variables declared as `parameter`.
- Input variables in functions behave as though they were parameter expressions.
- Except for the special built-in operators `initial`, `terminal`, `der`, `edge`, `change`, `sample`, and `pre`, a function or operator with parameter subexpressions is a parameter expression.
- Some function calls are parameter expressions even if the arguments are not:
 - `cardinality(c)`, see restrictions for use in section 3.7.4.3.
 - `end` in `A[... end ...]` if `A` is variable declared in a non-`function` class.
 - `size(A)` (including `size(A, j)` where `j` is parameter expression) if `A` is variable declared in a non-`function` class.
 - `Connections.isRoot(A.R)`
 - `Connections.rooted(A.R)`

3.8.3 Discrete-Time Expressions

Discrete-time expressions are:

- Parameter expressions.
- Discrete-time variables, i.e., `Integer`, `Boolean`, `String` variables and `enumeration` variables, as well as `Real` variables assigned in `when`-clauses.
- Function calls where all input arguments of the function are discrete-time expressions.
- Expressions where all the subexpressions are discrete-time expressions.
- Expressions in the body of a `when`-clause, `initial equation`, or `initial algorithm`.
- Unless inside `noEvent`: Ordered relations (`>`, `<`, `>=`, `<=`) and the event generating functions `ceil`, `floor`, `div`, and `integer`, if at least one argument is non-discrete time expression and subtype of `Real`.

[These will generate events, see section 8.5. Note that `rem` and `mod` generate events but are not discrete-time expressions. In other words, relations inside `noEvent`, such as `noEvent(x>1)`, are not discrete-time expressions.]
- The functions `pre`, `edge`, and `change` result in discrete-time expressions.
- Expressions in functions behave as though they were discrete-time expressions.

For an equation `expr1 = expr2` where neither expression is of base type `Real`, both expressions must be discrete-time expressions. For record equations the equation is split into basic types before applying this test.

[This restriction guarantees that `noEvent` cannot be applied to `Boolean`, `Integer`, `String`, or `enumeration` equations outside of a `when`-clause, because then one of the two expressions is not discrete-time.]

Inside an if-expression, `if`-clause, `while`-statement or `for`-clause, that is controlled by a non-discrete-time (that is continuous-time, but not discrete-time) switching expression and not in the body of a

when-clause, it is not legal to have assignments to discrete-time variables, equations between discrete-time expressions, or real elementary relations/functions that should generate events.

[The restriction above is necessary in order to guarantee that all equations for discrete-time variable are discrete-time expressions, and to ensure that crossing functions do not become active between events.]

[Example: The (underdetermined) model **Test** below illustrates two kinds of consequences due to variability constraints. First, it contains variability errors for declaration equations and assignments. Second, it illustrates the impact of variability on the matching of equations to variables, which can lead to violation of the perfect matching rule.

```

model Constants
  parameter Real p1 = 1;
  constant Real c1 = p1 + 2; // error, not a constant expression
  parameter Real p2 = p1 + 2; // fine
end Constants;
model Test
  Constants c1(p1=3); // fine
  Constants c2(p2=7); // fine, declaration equation can be modified
  Real x;
  Boolean b1 = noEvent(x > 1); // error, since b1 is a discrete-time variable
                                // and noEvent(x > 1) is not discrete-time.

  Boolean b2;
  Integer i1;
  Integer i2;
algorithm
  i1 := x; // error, assignment to variable of lesser variability.
equation
  b2 = noEvent(x > 1); // no variability error, but equation cannot be matched.
  i2 = x; // no variability error, and can be matched to x.
end Test;

```

]

3.8.4 Continuous-Time Expressions

All expressions are continuous-time expressions including constant, parameter and discrete expressions. The term *non-discrete-time expression* refers to expressions that are not constant, parameter or discrete expressions.

Chapter 4

Classes, Predefined Types, and Declarations

The fundamental structuring unit of modeling in Modelica is the class. Classes provide the structure for objects, also known as instances. Classes can contain equations which provide the basis for the executable code that is used for computation in Modelica. Conventional algorithmic code can also be part of classes. All data objects in Modelica are instantiated from classes, including the basic data types – `Real`, `Integer`, `String`, `Boolean` – and enumeration types, which are built-in classes or class schemata.

Declarations are the syntactic constructs needed to introduce classes and objects (i.e., components).

4.1 Access Control – Public and Protected Elements

Members of a Modelica class can have two levels of visibility: `public` or `protected`. The default is `public` if nothing else is specified

A protected element, P, in classes and components shall not be accessed via dot notation (e.g., `A.P`, `a.P`, `a[1].P`, `a.b.P`, `.A.P`; but there is no restriction on using `P` or `P.x` for a protected element P). They shall not be modified or redeclared except for modifiers applied to protected elements in a base-class modification (not inside any component or class) and the modifier on the declaration of the protected element.

[*Example:*

```
package A
  model B
    protected
      parameter Real x;
    end B;
  protected
    model C end C;
  public
    model D
      C c; // Legal use of protected class C from enclosing scope
      extends A.B(x=2); // Legal modifier for x in derived class
                        // also x.start=2 and x(start=2) are legal.
      Real y=x; // Legal use of x in derived class
    end D;
    model E
      A.B a(x=2); // Illegal modifier, also x.start=2 and x(start=2) are illegal
      A.C c; // Illegal use of protected class C
      model F=A.C; // Illegal use of protected class C
    end E;
  end A;
```

]

All elements defined under the heading `protected` are regarded as protected. All other elements (i.e., defined under the heading `public`, without headings or in a separate file) are public (i.e. not protected). Regarding inheritance of protected and public elements, see section 7.1.2.

4.2 Double Declaration not Allowed

The name of a declared element shall not have the same name as any other element in its partially flattened enclosing class. However, the internal flattening of a class can in some cases be interpreted as having two elements with the same name; these cases are described in section 5.5, and section 7.3.

[*Example:*

```

record R
  Real x;
end R;
model M // wrong Modelica model
  R R; // not correct, since component name and type specifier are identical
equation
  R.x = 0;
end M;

```

]

4.3 Declaration Order and Usage before Declaration

Variables and classes can be used before they are declared.

[*In fact, declaration order is only significant for:*

- *Functions with more than one input variable called with positional arguments, section 12.4.1.*
- *Functions with more than one output variable, section 12.4.3.*
- *Records that are used as arguments to external functions, section 12.9.1.3.*
- *Enumeration literal order within enumeration types, section 4.8.5.*

]

4.4 Component Declarations

Component declarations are described in this section.

4.4.1 Syntax and Examples of Component Declarations

The formal syntax of a component declaration clause is given by the following syntactic rules:

```

component-clause :
  type-prefix type-specifier [ array-subscripts ] component-list

type-prefix :
  [ flow | stream ]
  [ discrete | parameter | constant ] [ input | output ]

type-specifier :
  name

component-list :
  component-declaration { "," component-declaration }

component-declaration :
  declaration [ condition-attribute ] comment

```

```

condition-attribute:
  if expression

declaration :
  IDENT [ array-subscripts ] [ modification ]
  
```

[The declaration of a component states the type, access, variability, data flow, and other properties of the component. A component-clause i.e., the whole declaration, contains type prefixes followed by a type-specifier with optional array-subscripts followed by a component-list.

There is no semantic difference between variables declared in a single declaration or in multiple declarations. For example, regard the following single declaration (component-clause) of two matrix variables:

```
Real [2,2] A, B;
```

That declaration has the same meaning as the following two declarations together:

```
Real [2,2] A;
Real [2,2] B;
```

The array dimension descriptors may instead be placed after the variable name, giving the two declarations below, with the same meaning as in the previous example:

```
Real A [2,2];
Real B [2,2];
```

The following declaration is different, meaning that the variable *a* is a scalar but *B* is a matrix as above:

```
Real a, B [2,2];
```

]

4.4.2 Component Declaration Static Semantics

If the type-specifier of the component declaration denotes a built-in type (`RealType`, `IntegerType`, etc.), the flattened or instantiated component has the same type.

If the type-specifier of the component does not denote a built-in type, the name of the type is looked up (section 5.3). The found type is flattened with a new environment and the partially flattened enclosing class of the component. It is an error if the type is partial in a simulation model, or if a simulation model itself is partial. The new environment is the result of merging

- the modification of enclosing class element-modification with the same name as the component
- the modification of the component declaration

in that order.

Array dimensions shall be scalar non-negative parameter expressions of type `Integer`, a reference to a type (which must be an enumeration type or `Boolean`, see section 4.8.5), or the colon operator denoting that the array dimension is left unspecified (see section 10.1). All variants can also be part of short class definitions.

[Example of variables with array dimensions.

```

model ArrayVariants
  type T=Real[:]; // Unspecified size for type
  parameter T x=ones(4);
  parameter T y[3]=ones(3, 4);
  parameter Real a[2]=ones(2); // Specified using Integer
  parameter Real b[2, 0]=ones(2, 0); // Size 0 is allowed
  parameter Real c[:]=ones(0); // Unspecified size for variable
  parameter Integer n=0;
  Real x[n*2]=cat(1,ones(n),zeros(n)); // Parameter expressions are allowed
  Boolean notV[Boolean]={true,false}; // Indexing with type
end ArrayVariants;
  
```

]

The rules for components in functions are described in section 12.2.

Conditional declarations of components are described in section 4.4.5.

4.4.2.1 Declaration Equations

An environment that defines the value of a component of built-in type is said to define a declaration equation associated with the declared component. For declarations of vectors and matrices, declaration equations are associated with each element.

Only components of the restricted classes `type`, `record`, `operator record`, and `connector`, or components of classes inheriting from `ExternalObject` may have declaration equations. See also the corresponding rule for algorithms, section 11.2.1.2.

4.4.2.2 Prefix Rules

Variables declared with the `flow` or the `stream` type prefix shall be a subtype of `Real`.

Type prefixes (that is, `flow`, `stream`, `discrete`, `parameter`, `constant`, `input`, `output`) shall only be applied for type, record and connector components – see also record specialized class, section 4.6.

An exception is `input` for components whose type is of the special class function type (these can only be used for function formal parameters and has special semantics, see section 12.4.2), and the `input` prefix is not applied to the elements of the component and is allowed even if the elements have `input` or `output` prefix.

In addition, instances of classes extending from `ExternalObject` may have type prefixes `parameter` and `constant`, and in functions also type prefixes `input` and `output`, see section 12.9.7.

The type prefixes `flow`, `stream`, `input` and `output` of a structured component (except as described above) are also applied to the elements of the component (this is done after verifying that the type prefixes occurring on elements of the component are correct; e.g. the `flow` prefix can be used on a record component and all the record elements will generate zero-sum equations, even if elements of a record shall not be declared with the `flow` prefix). When any of the type prefixes `flow`, `stream`, `input` and `output` are applied for a structured component, no element of the component may have any of these type prefixes. The corresponding rules for the type prefixes `discrete`, `parameter` and `constant` are described in section 4.4.4.1 for structured components.

[*Example: `input` can only be used, if none of the elements has a `flow`, `stream`, `input` or `output` type prefix.*]

The prefixes `input` and `output` have a slightly different semantic meaning depending on the context where they are used:

- In *functions*, these prefixes define the computational causality of the function body, i.e., given the variables declared as `input`, the variables declared as `output` are computed in the function body, see section 12.4.
- In *simulation models* and *blocks* (i.e., on the top level of a model or block that shall be simulated), these prefixes define the interaction with the environment where the simulation model or block is used. Especially, the `input` prefix defines that values for such a variable have to be provided from the simulation environment and the `output` prefix defines that the values of the corresponding variable can be directly utilized in the simulation environment, see the notion of *globally balanced* in section 4.7.
- In component *models* and *blocks*, the `input` prefix defines that a binding equation has to be provided for the corresponding variable when the component is utilized in order to guarantee a locally balanced model (i.e., the number of local equations is identical to the local number of unknowns), see section 4.7.

[*Example:*

```

block FirstOrder
  input Real u;
```

```

    ...
  end FirstOrder;
  model UseFirstOrder
    FirstOrder firstOrder(u=time); // binding equation for u
    ...
  end UseFirstOrder;

```

]

The `output` prefix does not have a particular effect in a model or block component and is ignored.

- In *connectors*, prefixes `input` and `output` define that the corresponding connectors can only be connected according to block diagram semantics, see section 9.1 (e.g., a connector with an `output` variable can only be connected to a connector where the corresponding variable is declared as `input`). There is the restriction that connectors which have at least one variable declared as `input` must be externally connected, see section 4.7 (in order to get a locally balanced model, where the number of local unknowns is identical to the number of unknown equations). Together with the block diagram semantics rule this means, that such connectors must be connected *exactly once externally*.
- In *records*, prefixes `input` and `output` are not allowed, since otherwise a record could not be, e.g., passed as input argument to a function.

4.4.3 Acyclic Bindings of Constants and Parameters

The unexpanded binding equations for parameters and constants in the translated model must be acyclic after flattening; except that cycles are allowed if the cycles disappear when evaluating parameters having annotation `Evaluate = true` that are not part of the cycle. Thus it is not possible to introduce equations for parameters by cyclic dependencies.

[There is no exception for parameters with `fixed = false`, despite the fact that such parameters are generally allowed to be initialized from systems of dependent equations. However, a parameter with `fixed = false` can use an initial equation instead of a binding equation, allowing for cyclic dependencies.]

[Example:

```

constant Real p = 2 * q;
constant Real q = sin(p); // Illegal since p = 2 * q, q = sin(p) are cyclical

model ABCD
  parameter Real A[n, n];
  parameter Integer n = size(A, 1);
end ABCD;

final ABCD a;
// Illegal since cyclic dependencies between size(a.A,1) and a.n

ABCD b(redeclare Real A[2, 2] = [1, 2; 3, 4]);
// Legal since size of A is no longer dependent on n.

ABCD c(n = 2); // Legal since n is no longer dependent on the size of A.

parameter Real r = 2 * sin(r); // Illegal, since r = 2 * sin(r) is cyclic

partial model PartialLumpedVolume
  parameter Boolean use_T_start = true "= true, use T_start, otherwise h_start"
  annotation(Dialog(tab = "Initialization"), Evaluate = true);
  parameter Medium.Temperature T_start=if use_T_start then system.T_start else
  Medium.temperature_phX(p_start, h_start, X_start)
  annotation(Dialog(tab = "Initialization", enable = use_T_start));
  parameter Medium.SpecificEnthalpy h_start=if use_T_start then
  Medium.specificEnthalpy_pTX(p_start, T_start, X_start) else Medium.
  h_default
  annotation(Dialog(tab = "Initialization", enable = not use_T_start));

```

```

end PartialLumpedVolume;
// Cycle for T_start and h_start, but ok since disappears
// when evaluating use_T_start

// Illegal since the unexpanded bindings have cycles for both x and y
// (even if they would disappear if bindings were expanded).
model HasCycles
  parameter Integer n = 10;
  final constant Real A[3, 3] = [0, 0, 0; 1, 0, 0; 2, 3, 0];
  parameter Real y[3] = A * y + ones(3);
  parameter Real x[n] = cat(1, {3.4}, x[1:(n-1)]);
end HasCycles;

```

]

4.4.4 Component Variability Prefixes `discrete`, `parameter`, `constant`

The prefixes `discrete`, `parameter`, `constant` of a component declaration are called variability prefixes and define in which situation the variable values of a component are initialized (see section 8.5 and section 8.6) and when they are changed in transient analysis (= solution of initial value problem of the hybrid DAE):

- A variable `vc` declared with `constant` prefix remains constant during transient analysis, with a value that is unaffected by the initialization problem.
- A variable `vc` declared with the `parameter` prefix remains constant during transient analysis, with a value determined by the initialization problem.
- A *discrete-time* variable `vd` has a vanishing time derivative between events. Note that this is not the same as saying that `der(vd)=0` almost everywhere, as the derivative is not even defined at the events, and it is not legal to apply `der` to discrete-time variables as they are not continuous. During transient analysis the variable can only change its value at event instants (see section 8.5).
- A *continuous-time* variable `vn` may have a non-vanishing time derivative (`der(vn) <> 0` possible) and may also change its value discontinuously at any time during transient analysis (see section 8.5). If there are any discontinuities the variable is not differentiable.

If a `Real` variable is declared with the prefix `discrete` it must in a simulation model be assigned in a when-clause, either by an assignment or an equation. The variable assigned in a when-clause shall not be defined in a sub-component of `model` or `block` specialized class. (This is to keep the property of balanced models.)

A `Real` variable assigned in a when-clause is a discrete-time variable, even though it was not declared with the prefix `discrete`. A `Real` variable not assigned in any when-clause and without any type prefix is a continuous-time variable.

The default variability for `Integer`, `String`, `Boolean`, or `enumeration` variables is discrete-time, and it is not possible to declare continuous-time `Integer`, `String`, `Boolean`, or `enumeration` variables.

[The restriction that discrete-valued variables (of type `Boolean`, etc) cannot be declared with continuous-time variability is one of the foundations of the expression variability rules that will ensure that any discrete-valued expression has at most discrete-time variability, see section 3.8.]

The variability of expressions and restrictions on variability for definition equations is given in section 3.8.

[A discrete-time variable is a piecewise constant signal which changes its values only at event instants during simulation. Such types of variables are needed in order that special algorithms, such as the algorithm of Pantelides for index reduction, can be applied (it must be known that the time derivative of these variables is identical to zero). Furthermore, memory requirements can be reduced in the simulation environment, if it is known that a component can only change at event instants.]

A `parameter` variable is constant during simulation. This prefix gives the library designer the possibility to express that the physical equations in a library are only valid if some of the used components are constant during simulation. The same also holds for discrete-time and constant variables. Additionally, the `parameter` prefix allows a convenient graphical user interface in an experiment environment, to

support quick changes of the most important constants of a compiled model. In combination with an *if*-clause, a **parameter** prefix allows removing parts of a model before the symbolic processing of a model takes place in order to avoid variable causalities in the model (similar to **#ifdef** in C). Class parameters can be sometimes used as an alternative.

Example:

```

model Inertia
  parameter Boolean state = true;
  ...
equation
  J*a = t1 - t2;
  if state then // code which is removed during symbolic
    der(v) = a; // processing, if state=false
    der(r) = v;
  end if;
end Inertia;

```

A constant variable is similar to a parameter with the difference that constants cannot be changed after translation and usually not changed after they have been given a value. It can be used to represent mathematical constants, e.g.

```

final constant Real PI=4*atan(1);

```

There are no continuous-time Boolean, Integer or String variables. In the rare cases they are needed they can be faked by using Real variables, e.g.:

```

Boolean off1, off1a;
Real off2;
equation
  off1 = s1 < 0;
  off1a = noEvent(s1 < 0); // error, since off1a is discrete
  off2 = if noEvent(s2 < 0) then 1 else 0; // possible
  u1 = if off1 then s1 else 0; // state events
  u2 = if noEvent(off2 > 0.5) then s2 else 0; // no state events

```

Since *off1* is a discrete-time variable, state events are generated such that *off1* is only changed at event instants. Variable *off2* may change its value during continuous integration. Therefore, *u1* is guaranteed to be continuous during continuous integration whereas no such guarantee exists for *u2*.]

4.4.4.1 Variability of Structured Entities

For elements of structured entities with variability prefixes the most restrictive of the variability prefix and the variability of the component wins (using the default variability for the component if there is no variability prefix on the component).

[Example:

```

record A
  constant Real pi=3.14;
  Real y;
  Integer i;
end A;

parameter A a;
// a.pi is a constant
// a.y and a.i are parameters

A b;
// b.pi is a constant
// b.y is a continuous-time variable
// b.i is a discrete-time variable

```

]

4.4.5 Conditional Component Declaration

A component declaration can have a condition-attribute: `if expression`.

[Example:

```

parameter Integer level(min=1)=1;
Motor motor;
Level1 component1(J=J) if level==1 "Conditional component";
Level2 component2 if level==2 "Conditional component";
Level3 component3(J=component1.J) if level<2 "Conditional component";
// Illegal modifier on component3 since component1.J is conditional
// Even if we can see that component1 always exist if component3 exist
equation
connect(component1..., ...) "Connection to conditional component 1";
connect(component2.n, motor.n) "Connection to conditional component 2";
connect(component3.n, motor.n) "Connection to conditional component 3";
component1.u=0; // Illegal

```

]

The *expression* must be a `Boolean` scalar expression, and must be a parameter expression.

[A *parameter expression* is required since it shall be evaluated at compile time.]

A redeclaration of a component shall not include a condition attribute; and the condition attribute is kept from the original declaration (see section 6.4).

If the `Boolean` expression is false the component (including its modifier) is removed from the flattened DAE, and connections to/from the component are removed. A component declared with a condition-attribute can only be modified and/or used in connections.

[Adding the component and then removing it ensures that the component is valid.

If a `connect` equation defines the connection of a non-conditional component `c1` with a conditional component `c2` and `c2` is de-activated, then `c1` must still be a declared element.]

If the condition is true for a public connector containing flow variables the connector must be connected from the outside.

[The reason for this restriction is that the default flow equation is probably incorrect (since it could otherwise be an unconditional connector) and the model cannot check that connector is connected.]

4.5 Class Declarations

Essentially everything in Modelica is a class, from the predefined classes `Integer` and `Real`, to large packages such as the Modelica standard library.

[Example: A rather typical structure of a Modelica class is shown below. A class with a name, containing a number of declarations followed by a number of equations in an equation section.

```

class ClassName
  Declaration1
  Declaration2
  ...
equation
  equation1
  equation2
  ...
end ClassName;

```

]

The following is the formal syntax of class definitions, including the special variants described in later sections.

```

class-definition :
  [ encapsulated ] class-prefixes
  class-specifier

class-prefixes :
  [ partial ]
  ( class | model | [ operator ] record | block | [ expandable ] connector |
    type |
  package | [ ( pure | impure ) ] [ operator ] function | operator )

class-specifier :
  long-class-specifier | short-class-specifier | der-class-specifier

long-class-specifier :
  IDENT description-string composition end IDENT
  | extends IDENT [ class-modification ] description-string
  composition end IDENT

short-class-specifier :
  IDENT "=" base-prefix name [ array-subscripts ]
  [ class-modification ] comment
  | IDENT "=" enumeration "(" ( [enum-list] | ":" ) ")" comment

der-class-specifier :
  IDENT "=" der "(" name "," IDENT { "," IDENT } ")" comment

base-prefix :
  [ input | output ]

enum-list : enumeration-literal { "," enumeration-literal }

enumeration-literal : IDENT comment

composition :
  element-list
  { public element-list |
    protected element-list |
    equation-section |
    algorithm-section
  }
  [ external [ language-specification ]
  [ external-function-call ] [ annotation ] ";" ]
  [ annotation ";" ]
  
```

4.5.1 Short Class Definitions

A class definition of the form

```
class IDENT1 = IDENT2 class-modification;
```

is identical, except that IDENT2 may be replaceable and for the lexical scope of modifiers, where the short class definition does not introduce an additional lexical scope for modifiers, to the longer form

```
class IDENT1
  extends IDENT2 class-modification;
end IDENT1;
```

An exception to the above is that if the short class definition is declared as **encapsulated**, then the modifiers follow the rules for encapsulated classes and cannot be looked up in the enclosing scope.

[Example: Demonstrating the difference in scopes:

```
model Resistor
  parameter Real R;
```

```

...
end Resistor;
model A
  parameter Real R;
  replaceable model Load=Resistor(R=R) constrainedby TwoPin;
  // Correct, sets the R in Resistor to R from model A.
  replaceable model LoadError
    extends Resistor(R=R);
    // Gives the singular equation R=R, since the right-hand side R
    // is searched for in LoadError and found in its base-class Resistor.
  end LoadError constrainedby TwoPin;
  encapsulated model Load2=.Resistor(R=2); // Ok
  encapsulated model LoadR=.Resistor(R=R); // Illegal
  Load a,b,c;
  ConstantSource ...;
  ...
end A;

```

]

A short class definition of the form

```
type TN = T[N] (optional modifier);
```

where N represents arbitrary array dimensions, conceptually yields an array class

```

'array' TN
  T[n] _ (optional modifiers);
'end' TN;

```

Such an array class has exactly one anonymous component (_); see also section 4.5.2. When a component of such an array class type is flattened, the resulting flattened component type is an array type with the same dimensions as _ and with the optional modifier applied.

[Example: The types of f1 and f2 are identical:

```

type Force = Real [3] (unit={"Nm", "Nm", "Nm"});
Force f1;
Real f2 [3] (unit={"Nm", "Nm", "Nm"});

```

]

If a short class definition inherits from a partial class the new class definition will be partial, regardless of whether it is declared with the keyword partial or not.

[Example:

```

replaceable model Load=TwoPin;
Load R; // Error unless Load is redeclared since TwoPin is a partial class.

```

]

If a short class definition does not specify any specialized class the new class definition will inherit the specialized class (this rule applies iteratively and also for redeclare).

A base-prefix applied in the short-class definition does not influence its type, but is applied to components declared of this type or types derived from it; see also section 4.5.2.

[Example:

```

type InArgument = input Real;
type OutArgument = output Real [3];

function foo
  InArgument u; // Same as: input Real u
  OutArgument y; // Same as: output Real [3] y
algorithm
  y:=fill(u,3);

```

```
end foo;

Real x[:]=foo(time);
```

]

4.5.2 Restriction on combining base-classes and other elements

It is not legal to combine other components or base-classes with an extends from an array class, a class with non-empty base-prefix, a simple type (`Real`, `Boolean`, `Integer`, `String` and enumeration types), or any class transitively extending from an array class, a class with non-empty base-prefix, or a simple type (`Real`, `Boolean`, `Integer`, `String` and enumeration types).

[Example:

```
model Integrator
  input Real u;
  output Real y = x;
  Real x;
equation
  der(x) = u;
end Integrator;

model Integrators = Integrator[3]; // Legal

model IllegalModel
  extends Integrators;
  Real x; // Illegal combination of component and array class
end IllegalModel;

connector IllegalConnector
  extends Real;
  Real y; // Illegal combination of component and simple type
end IllegalConnector;
```

]

4.5.3 Local Class Definitions – Nested Classes

The local class should be statically flattenable with the partially flattened enclosing class of the local class apart from local class components that are `partial` or `outer`. The environment is the modification of any enclosing class element modification with the same name as the local class, or an empty environment.

The unflattened local class together with its environment becomes an element of the flattened enclosing class.

[Example: The following example demonstrates parameterization of a local class:

```
model C1
  type Voltage = Real(nominal=1);
  Voltage v1, v2;
end C1;

model C2
  extends C1(Voltage(nominal=1000));
end C2;
```

Flattening of class C2 yields a local class Voltage with nominal-modifier 1000. The variables v1 and v2 are instances of this local class and thus have a nominal value of 1000.]

4.6 Specialized Classes

Specialized kinds of classes (earlier known as *restricted classes*) `record`, `type`, `model`, `block`, `package`, `function` and `connector` have the properties of a general class, apart from restrictions. Moreover, they

have additional properties called *enhancements*. The definitions of the specialized classes are given below (additional restrictions on inheritance are in section 7.1.3):

- **record** – Only public sections are allowed in the definition or in any of its components (i.e., **equation**, **algorithm**, **initial equation**, **initial algorithm** and **protected** sections are not allowed). The elements of a record shall not have prefixes **input**, **output**, **inner**, **outer**, **stream**, or **flow**. Enhanced with implicitly available record constructor function, see section 12.6. The components directly declared in a record may only be of specialized class record or type.
- **type** – May only be predefined types, enumerations, array of type, or classes extending from type.
- **model** – The normal modeling class in Modelica.
- **block** – Same as **model** with the restriction that each connector component of a **block** must have prefixes **input** and/or **output** for all connector variables.

*[The purpose is to model input/output blocks of block diagrams. Due to the restrictions on **input** and **output** prefixes, connections between blocks are only possible according to block diagram semantics.]*

- **function** – See section 12.2 for restrictions and enhancements of functions. Enhanced to allow the function to contain an external function interface.

[Non-function specialized classes do not have this property.]

- **connector** – Only public sections are allowed in the definition or in any of its components (i.e., **equation**, **algorithm**, **initial equation**, **initial algorithm** and **protected** sections are not allowed).

Enhanced to allow **connect** to components of connector classes. The elements of a connector shall not have prefixes **inner**, or **outer**. May only contain components of specialized class **connector**, **record** and **type**.

- **package** – May only contain declarations of classes and constants. Enhanced to allow **import** of elements of packages. (See also chapter 13 on packages.)
- **operator record** – Similar to **record**; but operator overloading is possible, and due to this the typing rules are different, see chapter 6. It is not legal to extend from an **operator record** (or **connector** inheriting from **operator record**), except if the new class is an **operator record** or **connector** that is declared as a short class definition, whose modifier is either empty or only modify the default attributes for the component elements directly inside the **operator record**. An **operator record** can only extend from an **operator record**. It is not legal to extend from any of its enclosing scopes. (See chapter 14).
- **operator** – Similar to **package**; but may only contain declarations of functions. May only be placed directly in an **operator record**. (See also chapter 14).
- **operator function** – Shorthand for an **operator** with exactly one function; same restriction as **function** class and in addition may only be placed directly in an **operator record**.

[A function declaration

```
operator function foo ... end foo;
```

is conceptually treated as

```
operator foo function foo1
...
end foo1; end foo;
```

]

Additionally only components which are of specialized classes **record**, **type**, **operator record**, and connector classes based on any of those can be used as component references in normal expressions and in the left hand side of assignments, subject to normal type compatibility rules. Additionally components of connectors may be arguments of connect-equations, and any component can be used as argument to

the `ndims` and `size`-functions, or for accessing elements of that component (possibly in combination with array indexing).

[Example: Use of operator:

```

operator record Complex
  Real re;
  Real im;
  ...
  encapsulated operator function '*'
    import Complex;
    input Complex c1;
    input Complex c2;
    output Complex result
  algorithm
    result := Complex(re=c1.re*c2.re - c1.im*c2.im,
                     im=c1.re*c2.im + c1.im*c2.re);
  end '*';
end Complex;
record MyComplex
  extends Complex; // Error; extending from enclosing scope.
  Real k;
end MyComplex;
operator record ComplexVoltage = Complex(re(unit="V"),im(unit="V")); // allowed

```

]

4.7 Balanced Models

[In this section restrictions for `model` and `block` classes are present, in order that missing or too many equations can be detected and localized by a Modelica translator before using the respective `model` or `block` class. A non-trivial case is demonstrated in the following example:

```

partial model BaseCorrelation
  input Real x;
  Real y;
end BaseCorrelation;

model SpecialCorrelation // correct in Modelica 2.2 and 3.0
  extends BaseCorrelation(x=2);
equation
  y=2/x;
end SpecialCorrelation;

model UseCorrelation // correct according to Modelica 2.2
// not valid according to Modelica 3.0
  replaceable model Correlation=BaseCorrelation;
  Correlation correlation;
equation
  correlation.y=time;
end UseCorrelation;

model Broken // after redeclaration, there is 1 equation too much in Modelica
2.2
  UseCorrelation example(redeclare Correlation=SpecialCorrelation);
end Broken;

```

In this case one can argue that both `UseCorrelation` (adding an acausal equation) and `SpecialCorrelation` (adding a default to an input) are correct. Still, when combined they lead to a model with too many equations, and it is not possible to determine which model is incorrect without strict rules – as the ones defined here.

In Modelica 2.2, model `Broken` will work with some models. However, by just redeclaring it to model

SpecialCorrelation, an error will occur and it will be very difficult in a larger model to figure out the source of this error.

In Modelica 3.0, model **UseCorrelation** is no longer allowed and the translator will give an error. In fact, it is guaranteed that a redeclaration cannot lead to an unbalanced model any more.]

The restrictions below apply after flattening – i.e. inherited components are included – possibly modified. The corresponding restrictions on connectors and connections are in section 9.3.

Definition 4.1. Local number of unknowns. The local number of unknowns of a **model** or **block** class is the sum based on the components:

- For each declared component of specialized class **type** (**Real**, **Integer**, **String**, **Boolean**, enumeration and arrays of those, etc.) or **record**, or **operator record** not declared as **outer**, it is the number of unknown variables inside it (i.e., excluding parameters and constants and counting the elements after expanding all records, operator record, and arrays to a set of scalars of primitive types).
- Each declared component of specialized class **type** or **record** declared as **outer** is ignored.
[I.e., all variables inside the component are treated as known.]
- For each declared component of specialized class **connector** component, it is the number of unknown variables inside it (i.e., excluding parameters and constants and counting the elements after expanding all records and arrays to a set of scalars of primitive types).
- For each declared component of specialized class **block** or **model**, it is the sum of the number of inputs and flow variables in the (top level) public connector components of these components (and counting the elements after expanding all records and arrays to a set of scalars of primitive types).

□

Definition 4.2. Local equation size. The local equation size of a **model** or **block** class is the sum of the following numbers:

- The number of equations defined locally (i.e. not in any **model** or **block** component), including binding equations, and equations generated from connect-equations.
[This includes the proper count for when-clauses (see section 8.3.5), and algorithms (see section 11.1), and is also used for the flat Hybrid DAE formulation (see appendix C).]
- The number of input and flow variables present in each (top-level) public connector component.
[This represents the number of connection equations that will be provided when the class is used.]
- The number of (top level) public input variables that neither are connectors nor have binding equations.
[I.e., top-level inputs are treated as known variables. This represents the number of binding equations that will be provided when the class is used.]

□

[To clarify top-level inputs without binding equation (for non-inherited inputs binding equation is identical to declaration equation, but binding equations also include the case where another model extends **M** and has a modifier on **u** giving the value):

```

model M
  input Real u;
  input Real u2=2;
end M;
  
```

Here **u** and **u2** are top-level inputs and not connectors. The variable **u2** has a binding equation, but **u** does not have a binding equation. In the equation count, it is assumed that an equation for **u** is supplied when using the model.]

Definition 4.3. Locally balanced. A **model** or **block** class is locally balanced if the *local number of unknowns* is identical to the *local equation size* for all legal values of constants and parameters. □

[Here, legal values must respect final bindings and min/max-restrictions. A tool shall verify the locally balanced property for the actual values of parameters and constants in the simulation model. It is a quality of implementation for a tool to verify this property in general, due to arrays of (locally) undefined sizes, conditional declarations, for-loops etc.]

Definition 4.4. Globally balanced. Similarly as locally balanced, but including all unknowns and equations from all components. The global number of unknowns is computed by expanding all unknowns (i.e. excluding parameters and constants) into a set of scalars of primitive types. This should match the global equation size defined as:

- The number of equations defined (included in any `model` or `block` component), including equations generated from connect-equations.
- The number of input and flow variables present in each (top-level) public connector component.
- The number of (top level) public input variables that neither are connectors nor have binding equations.

[I.e., top-level inputs are treated as known variables.]

□

The following restrictions hold:

- In a non-partial `model` or `block`, all non-connector inputs of `model` or `block` components must have binding equations.

[E.g. if the model contains a component, `firstOrder` (of specialized class `model`) and `firstOrder` has `input Real u` then there must be a binding equation for `firstOrder.u`.]

- A component declared with the `inner` or `outer` prefix shall not be of a class having top-level public connectors containing inputs.
- In a declaration of a component of a record, connector, or simple type, modifiers can be applied to any element, and these are also considered for the equation count.

[Example:

```
Flange support(phi=phi, tau=torque1+torque2) if use_support;
```

If `use_support=true`, there are two additional equations for `support.phi` and `support.tau` via the modifier.]

- In other cases (declaration of a component of a `model` or `block` class, modifiers on extends, and modifier on short-class-definitions): Modifiers for components shall only contain redeclarations of replaceable elements and binding equations. The binding equations in modifiers for components may in these cases only be for parameters, constants, inputs and variables having a default binding equation.
- All non-partial `model` and `block` classes must be locally balanced.

[This means that the local number of unknowns equals the local equation size.]

Based on these restrictions, the following strong guarantee can be given:

- All simulation models and blocks are globally balanced.

[Therefore the number of unknowns equal to the number of equations of a simulation model or block, provided that every used non-partial `model` or `block` class is locally balanced.]

[Example: Example 1:

```
connector Pin
  Real v;
  flow Real i;
end Pin;

model Capacitor
  parameter Real C;
```

```

Pin p, n;
Real u;
equation
  0 = p.i + n.i;
  u = p.v - n.v;
  C*der(u) = p.i;
end Capacitor;

```

Model `Capacitor` is a locally balanced model according to the following analysis:

Locally unknown variables: `p.i`, `p.v`, `n.i`, `n.v`, `u`

Local equations:

$$\begin{aligned}
 0 &= p.i + n.i; \\
 u &= p.v - n.v; \\
 C \cdot \text{der}(u) &= p.i;
 \end{aligned}$$

and 2 equations corresponding to the 2 flow variables `p.i` and `n.i`.

These are 5 equations in 5 unknowns (locally balanced model). A more detailed analysis would reveal that this is structurally non-singular, i.e. that the hybrid DAE will not contain a singularity independent of actual values.

If the equation `u = p.v - n.v` would be missing in the `Capacitor` model, there would be 4 equations in 5 unknowns and the model would be locally unbalanced and thus simulation models in which this model is used would be usually structurally singular and thus not solvable.

If the equation `u = p.v - n.v` would be replaced by the equation `u = 0` and the equation `C*der(u) = p.i` would be replaced by the equation `C*der(u) = 0`, there would be 5 equations in 5 unknowns (locally balanced), but the equations would be singular, regardless of how the equations corresponding to the flow variables are constructed because the information that `u` is constant is given twice in a slightly different form.]

[Example: Example 2:

```

connector Pin
  Real v;
  flow Real i;
end Pin;

partial model TwoPin
  Pin p,n;
end TwoPin;

model Capacitor
  parameter Real C;
  extends TwoPin;
  Real u;
equation
  0 = p.i + n.i;
  u = p.v - n.v;
  C*der(u) = p.i;
end Capacitor;

model Circuit
  extends TwoPin;
  replaceable TwoPin t;
  Capacitor c(C=12);
equation
  connect(p, t.p);
  connect(t.n, c.p);
  connect(c.n, n);
end Circuit;

```

Since \mathbf{t} is partial we cannot check whether this is a globally balanced model, but we can check that `Circuit` is locally balanced.

Counting on model `Circuit` results in the following balance sheet:

Locally unknown variables (8): `p.i`, `p.v`, `n.i`, `n.v`, and 2 flow variables for `t` (`t.p.i`, `t.n.i`), and 2 flow variables for `c` (`c.p.i`, `c.n.i`).

Local equations:

$$\begin{aligned}
 p.v &= t.p.v; \\
 0 &= p.i - t.p.i; \\
 c.p.v &= t.n.v; \\
 0 &= c.p.i + t.n.i; \\
 n.v &= c.n.v; \\
 0 &= n.i - c.n.i;
 \end{aligned}$$

and 2 equation corresponding to the flow variables `p.i`, `n.i`.

In total we have 8 scalar unknowns and 8 scalar equations, i.e., a locally balanced model (and this feature holds for any models used for the replaceable component `t`).

Some more analysis reveals that this local set of equations and unknowns is structurally non-singular. However, this does not provide any guarantees for the global set of equations, and specific combinations of models that are locally non-singular may lead to a globally singular model.]

[Example: Example 3:

```

import Modelica.Units.SI;

partial model BaseProperties
  "Interface of medium model for all type of media"
  parameter Boolean preferredMediumStates=false;
  constant Integer nXi "Number of independent mass fractions";
  InputAbsolutePressure      p;
  InputSpecificEnthalpy      h;
  InputMassFraction          Xi[nXi];
  SI.Temperature              T;
  SI.Density                  d;
  SI.SpecificInternalEnergy  u;

  connector InputAbsolutePressure = input SI.AbsolutePressure;
  connector InputSpecificEnthalpy = input SI.SpecificEnthalpy;
  connector InputMassFraction = input SI.MassFraction;
end BaseProperties;
    
```

The use of connector here is a special design pattern. The variables `p`, `h`, `Xi` are marked as input to get correct equation count. Since they are connectors they should neither be given binding equations in derived classes nor when using the model. The design pattern is to give textual equations for them (as below); using connect-equations for these connectors would be possible (and would work) but is not part of the design.

This partial model defines that `T`, `d`, `u` can be computed from the medium model, provided `p`, `h`, `Xi` are given. Every medium with one or multiple substances and one or multiple phases, including incompressible media, has the property that `T`, `d`, `u` can be computed from `p`, `h`, `Xi`. A particular medium may have different “independent variables” from which all other intrinsic thermodynamic variables can be recursively computed. For example, a simple air model could be defined as:

```

model SimpleAir "Medium model of simple air. Independent variables: p,T"
  extends BaseProperties(nXi = 0,
    p(stateSelect = if preferredMediumStates then StateSelect.prefer
      else StateSelect.default),
    T(stateSelect = if preferredMediumStates then StateSelect.prefer
      else StateSelect.default));
    
```

```

constant SI.SpecificHeatCapacity R = 287;
constant SI.SpecificHeatCapacity cp = 1005.45;
constant SI.Temperature T0 = 298.15
equation
  d = p/(R*T);
  h = cp*(T-T0);
  u = h - p/d;
end SimpleAir;
    
```

The local number of unknowns in model `SimpleAir` (after flattening) is:

- 3 (T , d , u : variables defined in `BaseProperties` and inherited in `SimpleAir`), plus
- $2+nXi$ (p , h , Xi : variables inside connectors defined in `BaseProperties` and inherited in `SimpleAir`)

resulting in $5 + nXi$ unknowns. The local equation size is:

- 3 (equations defined in `SimpleAir`), plus
- $2 + nXi$ (input variables in the connectors inherited from `BaseProperties`)

Therefore, the model is locally balanced.

The generic medium model `BaseProperties` is used as a **replaceable model** in different components, like a dynamic volume or a fixed boundary condition:

```

import Modelica.Units.SI;

connector FluidPort
  replaceable model Medium = BaseProperties;
  SI.AbsolutePressure p;
  flow SI.MassFlowRate m_flow;
  SI.SpecificEnthalpy h;
  flow SI.EnthalpyFlowRate H_flow;
  SI.MassFraction Xi [Medium.nXi] "Independent mixture mass fractions";
  flow SI.MassFlowRate mXi_flow[Medium.nXi] "Independent subst. mass flow rates";
end FluidPort;

model DynamicVolume
  parameter SI.Volume V;
  replaceable model Medium = BaseProperties;
  FluidPort port(redeclare model Medium = Medium);
  Medium medium(preferredMediumStates=true); // No modifier for p,h,Xi
  SI.InternalEnergy U;
  SI.Mass M;
  SI.Mass MXi[medium.nXi];
equation
  U = medium.u*M;
  M = medium.d*V;
  MXi = medium.Xi*M;
  der(U) = port.H_flow; // Energy balance
  der(M) = port.m_flow; // Mass balance
  der(MXi) = port.mXi_flow; // Substance mass balance
// Equations binding to medium (inputs)
  medium.p = port.p;
  medium.h = port.h;
  medium.Xi = port.Xi;
end DynamicVolume;
    
```

The local number of unknowns of `DynamicVolume` is:

- $4 + 2 \cdot nXi$ (inside the port connector), plus
- $2 + nXi$ (variables U , M and MXi), plus

- $2 + nXi$ (the input variables in the connectors of the medium model)

resulting in $8 + 4 \cdot nXi$ unknowns; the local equation size is

- $6 + 3 \cdot nXi$ from the equation section, plus
- $2 + nXi$ flow variables in the port connector.

Therefore, `DynamicVolume` is a locally balanced model.

Note, when the `DynamicVolume` is used and the `Medium` model is redeclared to `SimpleAir`, then a tool will try to select `p`, `T` as states, since these variables have `StateSelect.prefer` in the `SimpleAir` model (this means that the default states `U`, `M` are derived quantities). If this state selection is performed, all intrinsic medium variables are computed from `medium.p` and `medium.T`, although `p` and `h` are the input arguments to the medium model. This demonstrates that in Modelica input/output does not define the computational causality. Instead, it defines that equations have to be provided here for `p`, `h`, `Xi`, in order that the equation count is correct. The actual computational causality can be different as it is demonstrated with the `SimpleAir` model.

```

model FixedBoundary_pTX
  parameter SI.AbsolutePressure p "Predefined boundary pressure";
  parameter SI.Temperature T "Predefined boundary temperature";
  parameter SI.MassFraction Xi[medium.nXi]
    "Predefined boundary mass fraction";
  replaceable model Medium = BaseProperties;
  FluidPort port(redeclare model Medium = Medium);
  Medium medium;
equation
  port.p = p;
  port.H_flow = semiLinear(port.m_flow, port.h , medium.h);
  port.MXi_flow = semiLinear(port.m_flow, port.Xi, medium.Xi);
  // Equations binding to medium (note: T is not an input).
  medium.p = p;
  medium.T = T;
  medium.Xi = Xi;
end FixedBoundary_pTX;
    
```

The number of local variables in `FixedBoundary_pTX` is:

- $4 + 2 \cdot nXi$ (inside the port connector), plus
- $2 + nXi$ (the input variables in the connectors of the medium model)

resulting in $6 + 3 \cdot nXi$ unknowns, while the local equation size is

- $4 + 2 \cdot nXi$ from the equation section, plus
- $2 + nXi$ flow variables in the port connector.

Therefore, `FixedBoundary_pTX` is a locally balanced model. The predefined boundary variables `p` and `Xi` are provided via equations to the input arguments `medium.p` and `medium.Xi`, in addition there is an equation for `T` in the same way – even though `T` is not an input. Depending on the flow direction, either the specific enthalpy in the port (`port.h`) or `h` is used to compute the enthalpy flow rate `H_flow`. `h` is provided as binding equation to the medium. With the equation `medium.T = T`, the specific enthalpy `h` of the reservoir is indirectly computed via the medium equations. Again, this demonstrates, that an `input` just defines the number of equations have to be provided, but that it not necessarily defines the computational causality.]

4.8 Predefined Types and Classes

The attributes of the predefined variable types (`Real`, `Integer`, `Boolean`, `String`) and `enumeration` types are described below with Modelica syntax although they are predefined. Attributes cannot be accessed using dot notation, and are not constrained by equations and algorithm sections. E.g. in `Real x(unit="kg") = y;` only the values of `x` and `y` are declared to be equal, but not their `unit` attributes, nor any other attribute of `x` and `y`. It is not possible to combine extends from the predefined types,

enumeration types, or this `Clock` type with other components. The names `Real`, `Integer`, `Boolean` and `String` are reserved such that it is illegal to declare an element with these names.

[Hence, it is possible to define a normal class called `Clock` in a package and extend from it.]

[It also follows that the only way to declare a subtype of e.g. `Real` is to use the `extends` mechanism.]

The definitions use `RealType`, `IntegerType`, `BooleanType`, `StringType`, `EnumType` as mnemonics corresponding to machine representations.

4.8.1 Real Type

The following is the predefined `Real` type:

```

type Real // Note: Defined with Modelica syntax although predefined
  RealType value; // Accessed without dot-notation
  parameter StringType quantity = "";
  parameter StringType unit = "" "Unit used in equations";
  parameter StringType displayUnit = "" "Default display unit";
  parameter RealType min=-Inf, max=+Inf; // Inf denotes a large value
  parameter RealType start = 0; // Initial value
  parameter BooleanType fixed = true, // default for parameter/constant;
                                = false; // default for other variables

  parameter RealType nominal; // Nominal value
  parameter BooleanType unbounded=false; // For error control
  parameter StateSelect stateSelect = StateSelect.default;
equation
  assert(value >= min and value <= max, "Variable value out of limit");
end Real;

```

The `nominal` attribute is meant to be used for scaling purposes and to define tolerances in relative terms, see section 4.8.6.

[For external functions in C89, `RealType` maps to `double`. In the mapping proposed in Annex F of the C99 standard, `RealType/double` matches the IEC 60559:1989 (ANSI/IEEE 754-1985) `double` format.]

4.8.2 Integer Type

The following is the predefined `Integer` type:

```

type Integer // Note: Defined with Modelica syntax although predefined
  IntegerType value; // Accessed without dot-notation
  parameter StringType quantity = "";
  parameter IntegerType min=-Inf, max=+Inf;
  parameter IntegerType start = 0; // Initial value
  parameter BooleanType fixed = true, // default for parameter/constant;
                                = false; // default for other variables

equation
  assert(value >= min and value <= max, "Variable value out of limit");
end Integer;

```

The minimal recommended number range for `IntegerType` is from -2147483648 to +2147483647, corresponding to a two's-complement 32-bit integer implementation.

4.8.3 Boolean Type

The following is the predefined `Boolean` type:

```

type Boolean // Note: Defined with Modelica syntax although predefined
  BooleanType value; // Accessed without dot-notation
  parameter StringType quantity = "";
  parameter BooleanType start = false; // Initial value
  parameter BooleanType fixed = true, // default for parameter/constant;
                                = false, // default for other variables

end Boolean;

```

4.8.4 String Type

The following is the predefined `String` type:

```

type String // Note: Defined with Modelica syntax although predefined
  StringType value; // Accessed without dot-notation
  parameter StringType quantity = "";
  parameter StringType start = ""; // Initial value
  parameter BooleanType fixed = true, // default for parameter/constant;
                                = false, // default for other variables
end String;
    
```

4.8.5 Enumeration Types

A declaration of the form

```

type E = enumeration([enum-list]);
    
```

defines an enumeration type `E` and the associated enumeration literals of the `enum-list`. The enumeration literals shall be distinct within the enumeration type. The names of the enumeration literals are defined inside the scope of `E`. Each enumeration literal in the `enum-list` has type `E`.

[Example:

```

type Size = enumeration(small, medium, large, xlarge);
Size t_shirt_size = Size.medium;
    
```

]

An optional comment string can be specified with each enumeration literal.

[Example:

```

type Size2 = enumeration(small "1st", medium "2nd", large "3rd", xlarge "4th");
    
```

]

An enumeration type is a simple type and the attributes are defined in section 4.8.5.1. The `Boolean` type name or an enumeration type name can be used to specify the dimension range for a dimension in an array declaration and to specify the range in a for-loop range expression; see section 11.2.2.2. An element of an enumeration type can be accessed in an expression.

[Uses of elements of enumeration type in expressions include indexing into an array.]

[Example:

```

type DigitalCurrentChoices = enumeration(zero, one);
// Similar to Real, Integer
    
```

Setting attributes:

```

type DigitalCurrent = DigitalCurrentChoices(quantity="Current",
                                           start = DigitalCurrentChoices.one, fixed = true)
;
DigitalCurrent c(start = DigitalCurrent.one, fixed = true);
DigitalCurrentChoices c(start = DigitalCurrentChoices.one, fixed = true);
    
```

Using enumeration types as expressions:

```

Real x[DigitalCurrentChoices];

// Example using the type name to represent the range
for e in DigitalCurrentChoices loop
  x[e] := 0.;
end for;

for e loop // Equivalent example using short form
    
```

```

    x[e] := 0.;
end for;

// Equivalent example using the colon range constructor
for e in DigitalCurrentChoices.zero : DigitalCurrentChoices.one loop
    x[e] := 0.;
end for;

model Mixing1 "Mixing of multi-substance flows, alternative 1"
    replaceable type E=enumeration(:)"Substances in Fluid";
    input Real c1[E], c2[E], mdot1, mdot2;
    output Real c3[E], mdot3;
equation
    0 = mdot1 + mdot2 + mdot3;
    for e in E loop
        0 = mdot1*c1[e] + mdot2*c2[e]+ mdot3*c3[e];
    end for;
    /* Array operations on enumerations are NOT (yet) possible:
       zeros(n) = mdot1*c1 + mdot2*c2 + mdot3*c3 // error
    */
end Mixing1;

model Mixing2 "Mixing of multi-substance flows, alternative 2"
    replaceable type E=enumeration(:)"Substances in Fluid";
    input Real c1[E], c2[E], mdot1, mdot2;
    output Real c3[E], mdot3;
protected
    // No efficiency loss, since cc1, cc2, cc3
    // may be removed during translation
    Real cc1[:]=c1, cc2[:]=c2, cc3[:]=c3;
    final parameter Integer n = size(cc1,1);
equation
    0 = mdot1 + mdot2 + mdot3;
    zeros(n) = mdot1*cc1 + mdot2*cc2 + mdot3*cc3
end Mixing2;

```

]

4.8.5.1 Attributes of Enumeration Types

For each enumeration:

```
type E=enumeration(e1, e2, ..., en);
```

a new simple type is conceptually defined as

```

type E // Note: Defined with Modelica syntax although predefined
    EnumType value; // Accessed without dot-notation
    parameter StringType quantity = "";
    parameter EnumType min=e1, max=en;
    parameter EnumType start = e1; // Initial value
    parameter BooleanType fixed = true, // default for parameter/constant;
                                = false; // default for other variables

    constant EnumType e1=...;
    ...
    constant EnumType en=...;
equation
    assert(value >= min and value <= max, "Variable value out of limit");
end E;

```

[Since the attributes and enumeration literals are on the same level, it is not possible to use the enumeration attribute names (quantity, min, max, start, fixed) as enumeration literals.]

4.8.5.2 Type Conversion of Enumeration Values to String or Integer

The type conversion function `Integer(<expression of enumeration type>)` returns the ordinal number of the enumeration value `E.enumvalue`, to which the expression is evaluated, where `Integer(E.e1) = 1`, `Integer(E.en) = n`, for an enumeration type `E = enumeration(e1, ..., en)`.

`String(E.enumvalue)` gives the `String` representation of the enumeration value.

[*Example: `String(E.Small)` gives "Small".*]

See also section 3.7.1.

4.8.5.3 Type Conversion of Integer to Enumeration Values

Whenever an enumeration type is defined, a type conversion function with the same name and in the same scope as the enumeration type is implicitly defined. This function can be used in an expression to convert an integer value to the corresponding (as described in section 4.8.5.2) enumeration value.

For an enumeration type named `EnumTypeName`, the expression `EnumTypeName(<Integer expression>)` returns the enumeration value `EnumTypeName.e` such that `Integer(EnumTypeName.e)` is equal to the original integer expression.

Attempting to convert an integer argument that does not correspond to a value of the enumeration type is an error.

[*Example:*

```
type Colors = enumeration ( RED, GREEN, BLUE, CYAN, MAGENTA, YELLOW );
```

Converting from Integer to Colors:

```
c = Colors(i);  
c = Colors(10); // An error
```

]

4.8.5.4 Unspecified enumeration

An enumeration type defined using `enumeration(:)` is unspecified and can be used as a replaceable enumeration type that can be freely redeclared to any enumeration type. There can be no enumeration variables declared using `enumeration(:)` in a simulation model.

4.8.6 Attributes `start`, `fixed`, `nominal`, and `unbounded`

The attributes `start` and `fixed` define the initial conditions for a variable. `fixed = false` means an initial guess, i.e., value may be changed by static analyzer. `fixed = true` means a required value. The resulting consistent set of values for *all* model variables is used as initial values for the analysis to be performed.

The attribute `nominal` gives the nominal value for the variable. The user need not set it even though the standard does not define a default value. The lack of default allows the tool to propagate the nominal attribute based on equations, and if there is no value to propagate the tool should use a non-zero value, it may use additional information (e.g. `min` attribute) to find a suitable value, and as last resort use 1. If `unbounded = true` it indicates that the state may grow without bound, and the error in absolute terms shall be controlled.

[*The nominal value can be used by an analysis tool to determine appropriate tolerances or epsilons, or may be used for scaling. For example, the tolerance for an integrator could be computed as `tol * (abs(nominal) + (if x.unbounded then 0 else abs(x)))`. A default value is not provided in order that in cases such as `a = b`, where `b` has a nominal value but not `a`, the nominal value can be propagated to the other variable.*]

4.8.7 Other Predefined Types

4.8.7.1 StateSelect

The predefined `StateSelect` enumeration type is the type of the `stateSelect` attribute of the `Real` type. It is used to explicitly control state selection.

```

type StateSelect = enumeration(
  never "Do not use as state at all.",
  avoid "Use as state, if it cannot be avoided (but only if variable appears
        differentiated and no other potential state with attribute
        default, prefer, or always can be selected).",
  default "Use as state if appropriate, but only if variable appears
          differentiated.",
  prefer "Prefer it as state over those having the default value
         (also variables can be selected, which do not appear
         differentiated). ",
  always "Do use it as a state."
);

```

4.8.7.2 ExternalObject

See section 12.9.7 for information about the predefined type `ExternalObject`.

4.8.7.3 AssertionLevel

The predefined `AssertionLevel` enumeration type is used together with `assert`, section 8.3.7.

```

type AssertionLevel=enumeration(warning, error);

```

4.8.7.4 Connections

The package `Connections` is used for over-constrained connection graphs, section 8.3.9.

4.8.7.5 Graphical Annotation Types

A number of “predefined” record types and enumeration types for graphical annotations are described in chapter 18. These types are not predefined in the usual sense since they cannot be referenced in ordinary Modelica code, only within annotations.

4.8.7.6 Clock Types

See section 16.2.1 and section 16.3.

Chapter 5

Scoping, Name Lookup, and Flattening

This chapter describes the scope rules, and most of the name lookup and flattening of Modelica.

5.1 Flattening Context

Flattening is made in a context which consists of a modification environment (section 7.2.2) and an ordered set of enclosing classes.

5.2 Enclosing Classes

The classes lexically enclosing an element form an ordered set of enclosing classes. A class defined inside another class definition (the enclosing class) precedes its enclosing class definition in this set.

Enclosing all class definitions is an unnamed enclosing class that contains all top-level class definitions, and not-yet read classes defined externally as described in section 13.4. The order of top-level class definitions in the unnamed enclosing class is undefined.

During flattening, the enclosing class of an element being flattened is a partially flattened class.

[For example, this means that a declaration can refer to a name inherited through an extends-clause.]

[Example:

```
class C1 ... end C1;
class C2 ... end C2;
class C3
  Real x=3;
  C1 y;
  class C4
    Real z;
  end C4;
end C3;
```

The unnamed enclosing class of class definition C3 contains C1, C2, and C3 in arbitrary order. When flattening class definition C3, the set of enclosing classes of the declaration of *x* is the partially flattened class C3 followed by the unnamed enclosing class with C1, C2, and C3. The set of enclosing classes of *z* is C4, C3 and the unnamed enclosing class in that order.]

5.3 Static Name Lookup

Names are looked up at class flattening to find names of base classes, component types, etc. Implicitly defined names of record constructor functions and enumeration type conversion functions are ignored

during type name lookup. Names of record classes and enumeration types are ignored during function name lookup.

[The reason to ignore the implicitly defined names is that a record and the implicitly created record constructor function, see section 12.6, and an enumeration type and the implicitly created conversion function (section 4.8.5.3), have the same name.]

5.3.1 Simple Name Lookup

When an element, equation, or section is flattened, any simple name (not composed using dot notation) is first looked up sequentially among iteration variables (if any; see below), and then looked up sequentially in each member of the ordered set of *instance scopes* (see section 5.6.1.1) corresponding to *lexically enclosing classes* until a match is found or an enclosing class is encapsulated. In the latter case the lookup stops except for the predefined types, functions and operators defined in this specification. For these cases the lookup continues in the global scope, where they are defined.

The iteration variables are the implicitly declared iteration variable(s) if inside the body of a for-loop, section 8.3.2 and section 11.2.2, or the body of a reduction expression, section 10.3.4.

Reference to variables successfully looked up in an enclosing class is only allowed for variables declared as **constant**. The values of modifiers are thus resolved in the *instance* scope of which the modifier appears; if the use is in a modifier on a short class definition, see section 4.5.1.

This lookup in each *instance* scope is performed as follows:

- Among declared named elements (**class-definition** and **component-declaration**) of the class (including elements inherited from base-classes).
- Among the import names of qualified import-clauses in the *instance* scope. The import name of `import A.B.C;` is C and the import name of `import D=A.B.C;` is D.
- Among the public members of packages imported via unqualified import-clauses in the *instance* scope. It is an error if this step produces matches from several unqualified imports.

Import statements defined in inherited classes are ignored for the lookup, i.e. import-clauses are not inherited.

5.3.2 Composite Name Lookup

For a composite name of the form A.B or A.B.C, etc. lookup is performed as follows:

- The first identifier (A) is looked up as defined above.
- If the first identifier denotes a component, the rest of the name (e.g., B or B.C) is looked up among the declared named component elements of the component.
- If not found, and if the first identifier denotes a scalar component, or component[j] where component is an array of components and the indices j can be evaluated at translation time and component[j] is a scalar; and if the composite name is used as a function call, the lookup is also performed among the declared named class elements of the scalar component, and must find a non-operator function. All identifiers of the rest of the name (e.g., B and B.C) must be classes.
- If the identifier denotes a class, that class is temporarily flattened (as if instantiating a component without modifiers of this class, see section 7.2.2) and using the enclosing classes of the denoted class. The rest of the name (e.g., B or B.C) is looked up among the declared named elements of the temporary flattened class. If the class does not satisfy the requirements for a package, the lookup is restricted to encapsulated elements only. The class we look inside shall not be partial in a simulation model.

[The temporary class flattening performed for composite names follow the same rules as class flattening of the base class in an extends-clause, local classes and the type in a component clause, except that the environment is empty. See also MoistAir2 example in section 7.3 for further explanations regarding looking inside partial packages.]

[Example: Components and classes are part of the same name-space and thus a component cannot have the same name as its class or the first part of the class-name as that would prevent lookup of the class name.]

```

model A
  M M;    // Illegal, component 'M' prevents finding class 'M'
  P.Q P;  // Illegal, component 'P' prevents finding package 'P'
  .R R;   // Legal, see next section
  S.Q Q;  // Legal

  Y a;    // Illegal, component 'Y' (below) prevents finding class 'Y'
  Y.X b;  // Illegal, component 'Y' (below) prevents finding package 'Y'
  .Y c;   // Legal, see next section
  Real Y;
end A;

```

]

5.3.3 Global Name Lookup

For a name starting with dot, e.g.: `.A` (or `.A.B`, `.A.B.C` etc.) lookup is performed as follows:

- The first identifier (`A`) is looked up in the global scope. This is possible even if the class is encapsulated and import-clauses are not used for this. If there does not exist a class `A` in global scope this is an error.
- If the name is simple then the class `A` is the result of lookup.
- If the name is a composite name then the class `A` is temporarily flattened with an empty environment (i.e. no modifiers, see section 7.2.2) and using the enclosing classes of the denoted class. The rest of the name (e.g., `B` or `B.C`) is looked up among the declared named elements of the temporary flattened class. If the class does not satisfy the requirements for a package, the lookup is restricted to encapsulated elements only. The class we look inside shall not be partial.

[The package-restriction ensures that global name lookup of component references can only find global constants.]

5.3.4 Lookup of Imported Names

See section 13.2.1.

5.4 Instance Hierarchy Name Lookup of Inner Declarations

An element declared with the prefix `outer` references an element instance with the same name but using the prefix `inner` which is nearest in the enclosing instance hierarchy of the `outer` element declaration.

Outer component declarations shall not have modifications (including binding equations). Outer class declarations should be defined using short-class definitions without modifications. However, see also section 5.5.

An `outer` element reference in a simulation model requires that one corresponding `inner` element declaration exist or can be created in a unique way:

- If there are two (or more) `outer` declarations with the same name, both lacking matching `inner` declarations, and the `outer` declarations are not of the same class it is in error.
- If there is one (or more) `outer` declarations of a partial class it is an error.
- In other cases, i.e. if a unique non-partial class is used for all `outer` declarations of the same name lacking a matching inner declaration, then an `inner` declaration of that class is automatically added at the top of the model and diagnostics is given.
- The annotations defined in section 18.7 does not affect this process, other than that:
 - `missingInnerMessage` can be used for the diagnostic (and possibly error messages)

An **outer** element component may be of a partial class (but the referenced **inner** component must be of a non-partial class).

[*inner/outer components may be used to model simple fields, where some physical quantities, such as gravity vector, environment temperature or environment pressure, are accessible from all components in a specific model hierarchy. Inner components are accessible throughout the model, if they are not “shadowed” by a corresponding inner declaration in a more deeply nested level of the model hierarchy.*]

[*Example: Simple Example:*

```
class A
  outer Real T0;
  ...
end A;
class B
  inner Real T0=1;
  A a1, a2; // B.T0, B.a1.T0 and B.a2.T0 will have the same value
  A a3(T0=4); // illegal as T0 is an outer variable.
  ...
end B;
```

More complicated example:

```
class A
  outer Real TI;
  class B
    Real TI;
    class C
      Real TI;
      class D
        outer Real TI; //
      end D;
      D d;
    end C;
    C c;
  end B;
  B b;
end A;

class E
  inner Real TI;
  class F
    inner Real TI;
    class G
      Real TI;
      class H
        A a;
      end H;
      H h;
    end G;
    G g;
  end F;
  F f;
end E;

class I
  inner Real TI;
  E e;
  // e.f.g.h.a.TI, e.f.g.h.a.b.c.d.TI, and e.f.TI is the same variable
  // But e.f.TI, e.TI and TI are different variables
  A a; // a.TI, a.b.c.d.TI, and TI is the same variable
end I;
```

]

The **inner** component shall be a subtype of the corresponding **outer** component.

[If the two types are not identical, the type of the **inner** component defines the instance and the **outer** component references just part of the **inner** component.]

[Example:

```
class A
  inner Real TI;
  class B
    outer Integer TI; // error , since A.TI is no subtype of A.B.TI
  end B;
end A;
```

]

5.4.1 Example of Field Functions using Inner/Outer

[Inner declarations can be used to define field functions, such as position dependent gravity fields, e.g.:

```
partial function A
  input Real u;
  output Real y;
end A;

function B // B is a subtype of A
  extends A;
algorithm
  ...
end B;

class D
  outer function fc = A;
  ...
equation
  y = fc(u);
end D;

class C
  inner function fc = B; // define function to be actually used
  D d; // The equation is now treated as y = B(u)
end C;
```

]

5.5 Simultaneous Inner/Outer Declarations

An element declared with both the prefixes **inner** and **outer** conceptually introduces two declarations with the same name: one that follows the above rules for **inner** and another that follows the rules for **outer**.

[Local references for elements with both the prefix **inner** and **outer** references the **outer** element. That in turn references the corresponding element in an enclosing scope with the prefix **inner**.]

Modifications of elements declared with both the prefixes **inner** and **outer** may have modifications, those modifications are only applied to the **inner** declaration.

[Example:

```
class A
  outer parameter Real p=2; // error , since modification
end A;
```

Intent of the following example: Propagate enabled through the hierarchy, and also be able to disable subsystems locally.

```

model ConditionalIntegrator "Simple differential equation if isEnabled"
  outer Boolean isEnabled;
  Real x(start=1);
equation
  der(x)=if isEnabled then -x else 0;
end ConditionalIntegrator;

model SubSystem "subsystem that 'enable' its conditional integrators"
  Boolean enableMe = time<=1;
  // Set inner isEnabled to outer isEnabled and enableMe
  inner outer Boolean isEnabled = isEnabled and enableMe;
  ConditionalIntegrator conditionalIntegrator;
  ConditionalIntegrator conditionalIntegrator2;
end SubSystem;

model System
  SubSystem subSystem;
  inner Boolean isEnabled = time>=0.5;
  // subSystem.conditionallIntegrator.isEnabled will be
  // 'isEnabled and subSystem.enableMe'
end System;

```

]

5.6 Flattening Process

In order to guarantee that elements can be used before they are declared and that elements do not depend on the order of their declaration (section 4.3) in the enclosing class, the flattening proceeds in the following two major steps:

1. Instantiation process
2. Generation of the flat equation system

The result is an equation system of all equations/algorithms, initial equations/algorithms and instances of referenced functions. Modifications of constants, parameters and variables are included in the form of equations.

The constants, parameters and variables are defined by globally unique identifiers and all references are resolved to the identifier of the referenced variable. No other transformations are performed.

5.6.1 Instantiation

The instantiation is performed in two steps. First a class tree is created and then from that an instance tree for a particular model is built up. This forms the basis for derivation of the flat equation system.

An implementation may delay and/or omit building parts of these trees, which means that the different steps can be interleaved. If an error occurs in a part of the tree that is not used for the model to be instantiated the corresponding diagnostics can be omitted (or be given). However, errors that should only be reported in a simulation model must be omitted there, since they are not part of the simulation model.

5.6.1.1 The Class Tree

All necessary libraries including the model which is to be instantiated are loaded from e.g. file system and form a so called *class tree*. This tree represents the syntactic information from the class definitions. It contains also all modifications at their original locations in syntactic form. The builtin classes are put into the unnamed root of the class tree.

[The class tree is built up directly during parsing of the Modelica texts. For each class a local tree is created which is then merged into the one big tree, according to the location of the class in the class hierarchy. This tree can be seen as the abstract syntax tree (AST) of the loaded libraries.]

5.6.1.2 The Instance Tree

The output of the instantiation process is an *instance tree*. The instance tree consists of nodes representing the elements of a class definition from the class tree. For a component the subtree of a particular node is created using the information from the class of the component clause and a new modification environment as result of merging the current modification environment with the modifications from the current element declaration (see section 7.2.3).

The instance tree has the following properties:

- It contains the instantiated elements of the class definitions, with redeclarations taken into account and merged modifications applied.
- Each instance knows its source class definition from the class tree and its modification environment.
- Each modification knows its instance scope.

The instance tree is used for lookup during instantiation. To be prepared for that, it has to be based on the structure of the class tree with respect to the class definitions. The builtin classes are instantiated and put in the unnamed root prior to the instantiation of the user classes, to be able to find them.

[The existence of the two separate trees (instance tree and class tree) is conceptual. Whether they really exist or are merged into only one tree or the needed information is held completely differently is an implementation detail. It is also a matter of implementation to have only these classes instantiated which are needed to instantiate the class of interest.]

A node in the instance tree is the instance scope for the modifiers and elements syntactically defined in the class it is instantiated from. The instance scope is the starting point for name lookup.

[If the name is not found the lookup is continued in the instance scope corresponding to the lexically enclosing class. Extends clauses are treated as unnamed nodes in the instance tree – when searching for an element in an instance scope the search also recursively examines the elements of the extends clauses. Except that inherited import-clauses are ignored.]

5.6.1.3 The Instantiation Procedure.

The instantiation is a recursive procedure with the following inputs:

- the class to be instantiated (current class)
- the modification environment with all applicable redeclarations and merged modifications (initially empty)
- a reference to the node of the instance tree, which the new instance should go into (parent instance)

The instantiation starts with the class to be instantiated, an empty modification environment, and an unnamed root node as parent node.

During instantiation all lookup is performed using the instance tree, starting from the instance scope of the current element. References in modifications and equations are resolved later (during generation of flat equation system) using the same lookup.

5.6.1.4 Steps of Instantiation

The element itself A partially instantiated class or component is an element that is ready to be instantiated; a partially instantiated element (i.e. class or component) is comprised of a reference to the original element (from the class tree) and the modifiers for that element (including a possible redeclaration).

The possible redeclaration of the element itself takes effect.

The class of a partially instantiated component is found in the instance tree (using the redeclaration if any), modifiers merged to that class forming a new partially instantiated class that is instantiated as below.

The local contents of the element For local classes and components in the current class, instance nodes are created and inserted into the current instance. Modifiers (including class redeclarations) are merged and associated with the instance and the element is partially instantiated.

[*The partially instantiated elements are used later for lookup during the generation of the flat equation system and are instantiated fully, if necessary, using the stored modification environment.*]

Equations, algorithms, and annotations of the class and the component declaration are copied to the instance without merging.

[*The annotations can be relevant for simulations, e.g. annotations for code generation (section 18.3.), simulation experiments (section 18.4.1) or functions (section 12.7, section 12.8 and section 12.9).*]

Extends clauses are not looked up, but empty extends clause nodes are created and inserted into the current instance (to be able to preserve the declaration order of components).

The inherited contents of the element Classes of extends clauses of the current class are looked up in the instance tree, modifiers (including redeclarations) are merged, the contents of these classes are partially instantiated using the new modification environment, and are inserted into an extends clause node, which is an unnamed node in the current instance that only contains the inherited contents from that base-class.

The classes of extends-clauses are looked up before and after handling extends-clauses; and it is an error if those lookups generate different results.

At the end, the current instance is checked whether their children (including children of extends-clauses) with the same name are identical and only the first one of them is kept. It is an error if they are not identical.

[*Only keeping the first among the children with the same name is important for function arguments where the order matters.*]

Recursive instantiation of components Components (local and inherited) are recursively instantiated.

[*Example: As an example, consider:*

```

model M
  model B
    A a;
    replaceable model A = C;
    type E = Boolean;
  end B;
  B b(redeclare model A = D (p=1));
  partial model C
    E e;
  end C;

  model D
    extends C;
    parameter E p;
    type E = Integer;
  end D;

  type E = Real;
end M;

```

To recursively instantiate M allowing the generation of flat equation system we have the following steps (not including checks):

1. Instantiate M: which partially instantiates B, b, C, D, and E.
2. Instantiate M.b:
 - 2.1. First find the class B in M (the partially instantiated elements have correct name allowing lookup)

- 2.2. instantiate the partially instantiated $M.B$ with the modifier `redeclare model A=D(p=1)`
- 2.3. partially instantiate $M.b.a$ (no modifier), and $M.b.A$ (with modifier `=D(p=1)`)
3. Instantiate $M.b.a$:
 - 3.1. First find the class A in $M.b$ (the partially instantiated elements have correct name allowing lookup)
 - 3.2. Instantiate the partially instantiated $M.b.A$ with the modifier `=D(p=1)`.
 - 3.2.1. Find the base-class `=D` from the modifier. This performs lookup for D in M , and finds the partially instantiated class D
 - 3.2.2. Instantiate the base-class $M.D$ with modifier `p=1`, and insert as unnamed node in $M.b.A$.
 - 3.2.2.1. Partially instantiate the component p with modifier `=1`
 - 3.2.2.2. Find the base-class C in $M.D$. Since there is no local element called C the search is then continued in M and finds the partially instantiated class $M.C$
 - 3.2.2.3. Instantiate the base-class $M.C$ as below
4. Instantiate the base-class $M.C$ inserting the result into unnamed node in $M.b.a$
 - 4.1. Partially instantiate e
 - 4.2. Instantiate e which requires finding E . First looking for E in the un-named node for `extends M.C`, and, since there is no local element E the search is then continued in M (which lexically encloses $M.C$) and finds E class inheriting from `Real`. The e is then instantiated using class E inheriting from `Real`.
5. Instantiate $M.b.a.p$
 - 5.1. First the class E in $M.b.a$ finding E class inheriting from `Integer`.
 - 5.2. Instantiate the $M.b.a.p$ using the class E inheriting from `Integer` with modifier `=1`
 - 5.3. Instantiate the base-class `Integer` with modifier `=1`, and insert as unnamed node in $M.b.a.p$.

An implementation can use different heuristics to be more efficient by re-using instantiated elements as long as the resulting flat equation system is identical.

Note that if D was consistently replaced by A in the example above the result would be identical (but harder to read due to two different classes called A .)]

5.6.2 Generation of the flat equation system

During this process, all references by name in conditional declarations, modifications, dimension definitions, annotations, equations and algorithms are resolved to the real instance to which they are referring to, and the names are replaced by the global unique identifier of the instance.

[This identifier is normally constructed from the names of the instances along a path in the instance tree (and omitting the unnamed nodes of `extends` clauses), separated by dots. Either the referenced instance belongs to the model to be simulated the path starts at the model itself, or if not, it starts at the unnamed root of the instance tree, e.g. in case of a constant in a package.]

[To resolve the names, a name lookup using the instance tree is performed, starting at the instance scope (unless the name is fully qualified) of the modification, algorithm or equation. If it is not found locally the search is continued at the instance of the lexically enclosing class of the scope (this is normally not equal to the parent of the current instance), and then continued with their parents as described in section 5.3. If the found component is an outer declaration, the search is continued using the direct parents in the instance tree (see section 5.4). If the lookup has to look into a class which is not instantiated yet (or only partially instantiated), it is instantiated in place.]

The flat equation system consists of a list of variables with dimensions, flattened equations and algorithms, and a list of called functions which are flattened separately. A flattened function consists of algorithm or external-clause and top-level variables (variables directly declared in the function or one of

its base-classes) – which recursively can contain other variables; the list of non-top level variables is not needed.

The instance tree is recursively walked through as follows for elements of the class (if necessary a partially instantiated component is first instantiated):

- At each visited component instance, the name is inserted into the variables list. Then the conditional declaration expression is evaluated if applicable.
 - The variable list is updated with the actual instance
 - The variability information and all other properties from the declaration are attached to this variable.
 - Dimension information from the declaration and all enclosing instances are resolved and attached to the variable to define their complete dimension.
 - If it is of record or simple type (`Boolean`, `Integer`, enumeration, `Real`, `String`, `Clock`, `ExternalObject`):
 - * In the modifications of *value* attribute references are resolved using the instance scope of the modification. An equation is formed from a reference to the name of the instance and the resolved modification value of the instance, and included into the equation system. Except if the value for an element of a record is overridden by the value for an entire record; section 7.2.3.
 - If it is of simple type (`Boolean`, `Integer`, enumeration, `Real`, `String`, `Clock`, `ExternalObject`):
 - * In the modifications of *non-value* attributes, e.g. `start`, `fixed` etc. references are resolved using the instance scope of the modification. An equation is formed from a reference to the name of the instance appended by a dot and the attribute name and the resolved modification value of the instance, and included into the equation system.
 - If it is of a non-simple type the instance is recursively handled.
- If there are equation or algorithm sections in the class definition of the instance, references are resolved using the instance scope of the instance and are included in the equation system. Some references – in particular to non simple, non record objects like connectors in connect-equations and states in `transition` equations are not resolved yet and handled afterwards.
- Instances of local classes are ignored.
- The unnamed nodes corresponding to extends-clauses are recursively handled.
- If there are function calls encountered during this process, the call is filled up with default arguments as defined in section 12.4.1. These are built from the modifications of input arguments which are resolved using their instance scope. The called function itself is looked up in the instance tree. All used functions are flattened and put into the list of functions.
- Conditional components with false condition are removed afterwards and they are not part of the simulation model.

[Thus e.g. parameters don't need values in them. However, type-error can be detected.]
- Each reference is checked, whether it is a valid reference, e.g. the referenced object belongs to or is an instance, where all existing conditional declaration expressions evaluate to true or it is a constant in a package.

[Conditional components can be used in connect-equations, and if the component is conditionally disabled the connect-statement is removed.]

This leads to a flattened equation system, except for connect and `transition` equations. These have to be transformed as described in chapter 9 and chapter 17. This may lead to further changes in the instance tree (e.g. from expandable connectors (section 9.1.3)) and additional equations in the flattened equation system (e.g. connect-equations (section 9.2), generated equations for state machine semantics (section 17.3.4)).

*[After flattening, the resulting equation system is self contained and covers all information needed to transform it to a simulatable model, but the class and instance trees are still needed: in the transformation process, there might be the need to instantiate further functions, e.g. from **derivative** annotation or from **inverse** annotation etc., on demand.]*

Chapter 6

Interface or Type Relationships

A class or component, e.g. denoted **A**, can in some cases be used at a location designed for another class or component, e.g. denoted **B**. In Modelica this is the case for replaceable classes (see section 7.3) and for `inner/outer` elements (see section 5.4). Replaceable classes are the primary mechanism to create very flexible models. In this chapter, the precise rules are defined when **A** can be used at a location designed for **B**. The restrictions are defined in terms of compatibility rules (section 6.4 and section 6.5) between “interfaces” (section 6.2); this can also be viewed as sub-typing (section 6.2).

6.1 Interface Terminology

In this chapter, two kinds of terminology is used for identical concepts to get better understanding (e.g. by both engineers and computer scientists). A short summary of the terms is given in the following table. The details are defined in the rest of this chapter.

Definition 6.1. Type or interface. The “essential” part of the public declaration sections of a class that is needed to decide whether **A** can be used instead of **B**.

[E.g. a declaration `Real x` is part of the type (also called interface), but `import A` is not.] □

Definition 6.2. Class type or inheritance interface. The “essential” part of the public and protected declaration sections of a class that is needed to decide whether **A** can be used instead of **B**. The class type, also called inheritance interface, is needed when inheritance takes place, since then the protected declarations have to be taken into account. □

Definition 6.3. Subtype or compatible interface. **A** is a subtype of **B**, or equivalently, the interface of **A** is compatible to the interface of **B**, if the “essential” part of the public declaration sections of **B** is also available in **A**.

[E.g., if **B** has a declaration `Real x`, this declaration must also be present in **A**. If **A** has a declaration `Real y`, this declaration must not be present in **B**.] □

Definition 6.4. Restricted subtype or plug compatible interface. **A** is a restricted subtype of **B**, or equivalently, the interface of **A** is plug compatible to the interface of **B**, if **A** is a subtype of **B** and if connector components in **A** that are not in **B**, are default connectable.

[E.g. it is not allowed that these connectors have variables with the `input` prefix, because then they must be connected.]

A model or block **A** cannot be used instead of **B**, if the particular situation does not allow to make a connection to these additional connectors. In such a case the stricter *plug compatible* is required for a redeclaration. □

Definition 6.5. Function subtype or function compatible interface. **A** is a function subtype of **B**, or equivalently, the interface of **A** is function compatible to the interface of **B**, if **A** is a subtype of **B** and if the additional arguments of function **A** that are not in function **B** are defined in such a way, that **A** can be called at places where **B** is called.

[E.g. an additional argument must have a default value.] □

6.2 The Concepts of Type, Interface and Subtype

A *type* can conceptually be viewed as a *set of values*. When we say that the variable `x` has the type `Real`, we mean that the value of `x` belongs to the set of values represented by the type `Real` i.e., roughly the set of floating point numbers representable by `Real`, for the moment ignoring the fact that `Real` is also viewed as a class with certain attributes. Analogously, the variable `b` having `Boolean` type means that the value of `b` belongs to the set of values `{false, true}`. The built-in types `Real`, `Integer`, `String`, `Boolean` are considered to be distinct types.

The *subtype* relation between types is analogous to the subset relation between sets. A type `A1` being a subtype of type `A` means that the set of values corresponding to type `A1` is a subset of the set of values corresponding to type `A`.

The type `Integer` is not a subtype of `Real` in Modelica even though the set of primitive integer values is a subset of the primitive real values since there are some attributes of `Real` that are not part of `Integer` (section 4.8).

The concept of *interface* as defined in section 6.3 and used in this document is equivalent to the notion of type based on sets in the following sense:

An element is characterized by its interface defined by some attributes (section 6.3). The *type* of the element is the set of values having the same interface, i.e. the same attributes.

A *subtype* `A1` in relation to another type `A`, means that the elements of the set corresponding to `A1` is a subset of the set corresponding to `A`, characterized by the elements of that subset having additional properties.

[Example: A record `R`: `record R Boolean b; Real x; end R`;

Another record called `R2`: `R2 Boolean b; Real x; Real y; end R2`;

An instance `r`: `R r`;

An instance `r2`: `R2 r2`;

The type `R` of `r` can be viewed as the set of all record values having the attributes defined by the interface of `R`, e.g. the infinite set `{R(b=false, x=1.2), R(b=false, x=3.4), R(b=true, x=1.2), R(b=true, x=1.2, y=2), R(b=true, x=1.2, a=2), ...}`. The statement that `r` has the type (or interface) `R` means that the value of `r` is to this infinite set.

The type `R2` is a subtype of `R` since its instances fulfill the additional property of having the component `Real y`; in all its values.

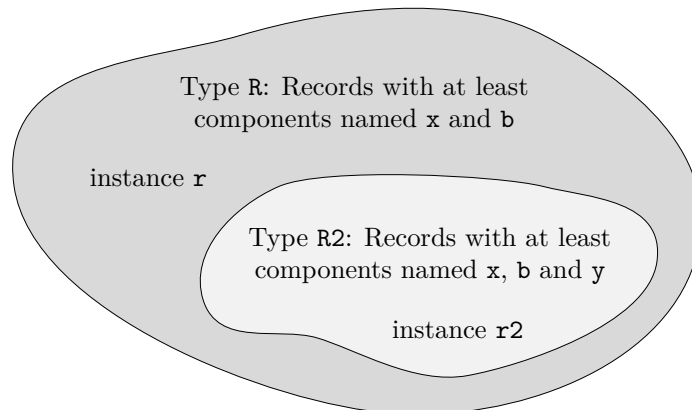


Figure 6.1: The type `R` can be defined as the set of record values containing `x` and `b`. The subtype `R2` is the subset of values that all contain `x`, `b`, and `y`.

]

6.3 Interface or Type

Based on a flattened class or component we can construct an interface for that flattened class or component. The interface or type (the terms *interface* and *type* are equivalent and can be used interchangeably) is defined as the following information about the flattened element itself:

- Whether it is replaceable or not.
- Whether the class itself or the class of the component is transitively non-replaceable (section 6.3.1), and if not, the reference to the replaceable class it refers to.
- Whether it is a component or a class.
- Additional information about the element:
 - The **flow** or **stream** prefix.
 - Declared variability (**constant**, **parameter**, **discrete**).
 - The prefixes **input** and **output**.
 - The prefixes **inner** and/or **outer**.
 - Whether the declaration is **final**, and in that case its semantics contents.
 - Array sizes (if any).
 - Condition of conditional components (if any).
 - Which kind of specialized class.
 - For an enumeration type or component of enumeration type the names of the enumeration literals in order.
 - Whether it is a built-in type and the built-in type (**RealType**, **IntegerType**, **StringType** or **BooleanType**).
- Only for an **operator record** class and classes derived from **ExternalObject**: the full name of the operator record base-class (i.e. the one containing the operations), or the derived class. See chapter 14 and section 12.9.7.

The following item does not apply for an **operator record** class or class derived from **ExternalObject**, since the type is already uniquely defined by the full name.

- For each named public element of the class or component (including both local and inherited named elements) a tuple comprised of:
 - Name of the element.
 - Interface or type of the element.

[This might have been modified by modifiers and is thus not necessarily identical to the interface of the original declaration.]

The corresponding *constraining* interface is constructed based on the *constraining* type (section 7.3.2) of the declaration (if replaceable – otherwise same as actual type) and with the *constraining* interface for the named elements.

In a class all references to elements of that class should be limited to their constraining interface.

[The constraining interface consists of only the public elements, and if the declaration is replaceable the element is limited to the constraining interface.]

[The public interface does not contain all of the information about the class or component. When using a class as a base-class we also need protected elements, and for internal type-checking we need e.g. import-elements. However, the information is sufficient for checking compatibility and for using the class to flatten components.]

6.3.1 Transitively non-Replaceable

[In several cases it is important that no new elements can be added to the interface of a class, especially considering short class definitions. Such classes are defined as transitively non-replaceable.]

A class reference is transitively non-replaceable iff (i.e. *if and only if*) all parts of the name satisfy the following:

- If the class definition is long it is transitively non-replaceable if not declared replaceable.
- If the class definition is short (i.e. `class A = P.B`) it is transitively non-replaceable if it is non-replaceable and equal to class reference (`P.B`) that is transitively non-replaceable.

[According to section 7.1.4, for a hierarchical name all parts of the name must be transitively non-replaceable, i.e. in `extends A.B.C` this implies that `A.B.C` must be transitively non-replaceable, as well as `A` and `A.B`, with the exception of the class extends redeclaration mechanism see section 7.3.1.]

6.3.2 Inheritance Interface or Class Type

For inheritance the interface also must include protected elements; this is the only change compared to above.

Based on a flattened class we can construct an inheritance interface or class type for that flattened class. The inheritance interface or class type is defined as the following information about the flattened element itself:

- Whether it is replaceable or not.
- Whether the class itself or the class of the component is transitively non-replaceable (section 6.3.1), and if not the reference to replaceable class it refers to.
- For each named element of the class (including both local and inherited named elements) a tuple comprised of:
 - Name of the element.
 - Whether the element is component or a class.
 - For elements that are classes: Inheritance interface or class type of the element.
 [This might have been modified by modifiers and is thus not necessarily identical to the interface of the original declaration.]
 - For elements that are components: interface or type of the element.
 [This might have been modified by modifiers and is thus not necessarily identical to the interface of the original declaration.]
- Additional information about the element:
 - The `flow` or `stream` prefix.
 - Declared variability (`constant`, `parameter`, `discrete`).
 - The prefixes `input` and `output`.
 - The prefixes `inner` and/or `outer`.
 - Whether the declaration is `final`, and in that case its semantics contents.
 - Array sizes (if any).
 - Condition of conditional components (if any).
 - Which kind of specialized class.
 - For an enumeration type or component of enumeration type the names of the enumeration literals in order.
 - Whether it is a built-in type and the built-in type (`RealType`, `IntegerType`, `StringType` or `BooleanType`).

- Visibility (`public` or `protected`).

6.4 Interface Compatibility or Subtyping

An interface of a class or component A is compatible with an interface of a class or component B (or the constraining interface of B), or equivalently that the type of A is a subtype of the type of B, iff:

- A is a class if and only if B is a class (and thus: A is a component if and only if B is a component).
- If A has an `operator record` base-class then B must also have one and it must be the same. If A does not have an operator record base-class then B shall not have one. See chapter 14.
- If A is derived from `ExternalObject`, then B must also be derived from `ExternalObject` and have the same full name. If A is not derived from `ExternalObject` then B shall not be derived from `ExternalObject`. See section 12.9.7.
- If B is not replaceable then A shall not be replaceable.
- If B is transitively non-replaceable then A must be transitively non-replaceable (section 6.3.1). For all elements of the inheritance interface of B there must exist a compatible element with the same name and visibility in the inheritance interface of A. The interface of A shall not contain any other elements.

[We might even extend this to say that A and B should have the same contents, as in the additional restrictions below.]

- If B is replaceable then for all elements of the component interface of B there must exist a plug-compatible element with the same name in the component interface of A.
- If B is neither transitively non-replaceable nor replaceable then A must be linked to the same class, and for all elements of the component interface of B there must thus exist a plug-compatible element with the same name in the component interface of A.
- Additional restrictions on the additional information. These elements should either match or have a natural total order:

- If B is a non-replaceable long class definition A must also be a long class definition.
- The `flow` or `stream` prefix should be matched for compatibility.
- Declared variability is ordered `constant` < `parameter` < `discrete` < continuous-time (`Real` without prefix), and A is only compatible with B if the declared variability in A is less than or equal the variability in B.

[For a redeclaration of an element the variability prefix is as default inherited by the redeclaration (i.e. no need to repeat `parameter` when redeclaring a parameter).]

- The `input` and `output` prefixes must be matched. This ensures that the rules regarding inputs/outputs for matching connectors and (non-connector inputs) are preserved, as well as the restriction on blocks.

[For a redeclaration of an element the `input` or `output` prefix is inherited from the original declaration.]

- The `inner` and/or `outer` prefixes should be matched.

[For a redeclaration of an element the `inner` and/or `outer` prefixes are inherited from the original declaration (since it is not possible to have `inner` and/or `outer` as part of a redeclare).]

- If B is final A must also be final and have the same semantic contents.
- The number of array dimensions in A and B must be matched.
- Conditional components are only compatible with conditional components. The conditions must have equivalent contents (similar to array sizes, except there is no `:` for conditional components).

[For a redeclaration of an element the conditional part is inherited from the original.]

- A **function** class is only compatible with a **function** class, a **package** class only compatible with a **package** class, a **connector** class only with a **connector** class, a **model** or **block** class only compatible with a **model** or **block** class, and a **type** or **record** class only compatible with a **type** or **record** class.
- If B is an enumeration type A must also be an enumeration type and vice versa. If B is an enumeration type not defined as (:) then A must have the same enumeration literals in the same order; if B is an enumeration type defined as (:) then there is no restriction on the enumeration type A.
- If B is a built-in type then A must also be of the same built-in type and vice versa.

[Intuitively, that the type A is a subtype of the type of B means that all important elements of B are present in A.]

Plug-compatibility is a further restriction of compatibility (subtyping) defined in section 6.5, and further restricted for functions, see section 6.6. For a replaceable declaration or modifier the default class must be compatible with the constraining class.

For a modifier the following must apply:

- The modified element should exist in the element being modified.
- The modifier should be compatible with the element being modified, and in most cases also plug-compatible, section 6.5.

[If the original constraining flat class is legal (no references to unknown elements and no illegal use of class/component), and modifiers legal as above, then the resulting flat class will be legal (no references to unknown elements and no illegal use of class/component and compatible with original constraining class) and references refer to similar entities.]

6.5 Plug-Compatibility or Restricted Subtyping

[If a sub-component is redeclared, see section 7.3, it is impossible to connect to any new connector. A connector with **input** prefix must be connected to, and since one cannot connect across hierarchies, one should not be allowed to introduce such a connector at a level where a connection is not possible. Therefore all public components present in the interface A that are not present in B must be connected by default.]

Definition 6.6. Plug-compatibility (= restricted subtyping). An interface A is plug-compatible with (a restricted subtype of) an interface B (or the constraining interface of B) iff:

- A is compatible with (subtype of) B.
- All public components present in A but not in B must be default-connectable (as defined below).

□

Definition 6.7. Default connectable. A component of an interface is default-connectable iff:

- All of its components are default connectable.
- A connector component must not be an **input**.
[Otherwise a connection to the input will be missing.]
- A connector component must not be of an expandable connector class.
[The expandable connector does potentially have inputs.]
- A parameter, constant, or non-connector input must either have a binding equation or all of its sub-components must have binding equations.

□

Based on the above definitions, there are the following restrictions:

- A redeclaration of an inherited top-level component must be *compatible with* (subtype of) the constraining interface of the element being redeclared.

- In all other cases redeclarations must be *plug-compatible* with the constraining interface of the element being redeclared.

[The reason for the difference is that for an inherited top-level component it is possible to connect to the additional connectors, either in this class or in a derived class.]

Example:

```

partial model TwoFlanges
  Modelica.Mechanics.Rotational.Interfaces.Flange_a flange_a;
  Modelica.Mechanics.Rotational.Interfaces.Flange_b flange_b;
end TwoFlanges;

partial model FrictionElement
  extends TwoFlanges;
  ...
end FrictionElement;

model Clutch "compatible - but not plug-compatible with FrictionElement"
  Modelica.Blocks.Interfaces.RealInput pressure;
  extends FrictionElement;
  ...
end Clutch;

model DriveLineBase
  extends TwoFlanges;
  Inertia J1;
  replaceable FrictionElement friction;
equation
  connect(flange_a, J1.flange_a);
  connect(J1.flange_b, friction.flange_a);
  connect(friction.flange_b, flange_b);
end DriveLineBase;

model DriveLine
  extends DriveLineBase(redeclare Clutch friction);
  Constant const;
equation
  connect(const.y, friction.pressure);
  // Legal connection to new input connector.
end DriveLine;

model UseDriveLine "illegal model"
  DriveLineBase base(redeclare Clutch friction);
  // Cannot connect to friction.pressure
end UseDriveLine;

```

If a subcomponent is redeclared, it is impossible to connect to any new connector. Thus any new connectors must work without being connected, i.e., the default connection of flow variables. That fails for inputs (and expandable connectors may contain inputs). For parameters and non-connector inputs it would be possible to provide bindings in a derived class, but that would require hierarchical modifiers and it would be bad modeling practice that a hierarchical modifier must be used in order to make a model valid. A replaceable class might be used as the class for a sub-component, therefore *plug-compatibility* is required not only for replaceable sub-components, but also for replaceable classes.]

6.6 Function-Compatibility or Function-Subtyping for Functions

[Functions may be called with either named or positional arguments, and thus both the name and order is significant. If a function is redeclared, see section 7.3, any new arguments must have defaults (and be at the end) in order to preserve the meaning of existing calls.]

Definition 6.8. Function-compatibility or function-subtyping for functions. A `function` class A is *function-compatible* with or a *function subtype* of `function` class B iff (the terms *function-compatible*

and *function subtype* of are synonyms and used interchangeably):

- A is compatible to (subtype of) B.
- All public input components of B have correspondingly named public input components of A in the same order and preceding any additional public input components of A.
- All public output components of B have correspondingly named public output components of A in the same order and preceding any additional public output components of A.
- A public input component of A must have a binding assignment if the corresponding named element has a binding assignment in B.
- A public input component of A not present in B must have a binding assignment.

□

Based on the above definition the following restriction holds:

- The interface of a redeclared function must be *function-compatible with or a function subtype of* the constraining interface of the function being redeclared.

[Example: Demonstrating a redeclaration using a function-compatible function

```
function GravityInterface
  input Modelica.Units.SI.Position position[3];
  output Modelica.Units.SI.Acceleration acceleration[3];
end GravityInterface;

function PointMassGravity
  extends GravityInterface;
  input Modelica.Units.SI.Mass m;
algorithm
  acceleration := -Modelica.Constants.G*m*position/(position*position)^1.5;
end PointMassGravity;

model Body
  model UseDriveLine "illegal model"
    DriveLineBase base(redeclare Clutch friction);
    // Cannot connect to friction.pressure
  end UseDriveLine;
  Modelica.Mechanics.MultiBody.Interface.Frame_a frame_a;
  replaceable function gravity = GravityInterface;
equation
  frame_a.f = gravity(frame_a.r0);
  // or gravity(position=frame_a.r0);
  frame_a.t = zeros(3);
end Body;

model PlanetSimulation
  function sunGravity = PointMassGravity (m=2e30);
  Body planet1(redeclare function gravity = sunGravity);
  Body planet2(redeclare function gravity = PointMassGravity (m=2e30));
  ...
end PlanetSimulation;
```

Note: *PointMassGravity* is not function-compatible with *GravityInterface* (no default for *m*), but *sunGravity* inside *PlanetSimulation* is function-compatible with *GravityInterface*.]

6.7 Type Compatible Expressions

Certain expressions consist of an operator applied to two or more type compatible sub-expressions (A and B), including binary operators, e.g. $A + B$, if-expressions, e.g. `if x then A else B`, and array expressions, e.g. $\{A, B\}$. The resulting type of the expression in case of two type compatible subexpressions A and B is defined as follows:

- If **A** is a record-expression **B** must also be a record-expression with the same named elements. The type compatible expression is a record comprised of named elements that are compatible with the corresponding named elements of both **A** and **B**.
- If **A** is an array expression then **B** must also be an array expression, and $\text{ndims}(\text{A}) = \text{ndims}(\text{B})$. The type compatible expression is an array expression with elements compatible with the elements of both **A** and **B**. If both $\text{size}(\text{A})$ and $\text{size}(\text{B})$ are known and $\text{size}(\text{A}) = \text{size}(\text{B})$ then this defines the size of the type compatible expression, otherwise the size of the expression is not known until the expression is about to be evaluated. In case of an if-expression the size of the type compatible expression is defined based on the branch selected, and for other cases $\text{size}(\text{A}) = \text{size}(\text{B})$ must hold at this point.
- If **A** is a scalar expression of a simple type **B** must also be a scalar expression of a simple type.
- If **A** is a **Real** expression then **B** must be a **Real** or **Integer** expression and the type compatible expression is **Real**, compare section 10.6.13.
- If **A** is an **Integer** expression then **B** must be a **Real** or **Integer** expression. For exponentiation and division the type compatible expression is **Real** (even if both **A** and **B** are **Integer**) see section 10.6.7 and section 10.6.5, in other cases the type compatible expression is **Real** or **Integer** (same as **B**), compare section 10.6.13.
- If **A** is a **Boolean** expression then **B** must be a **Boolean** expression and the type compatible expression is **Boolean**.
- If **A** is a **String** expression then **B** must be a **String** expression and the type compatible expression is **String**.
- If **A** is an enumeration expression then **B** must be an enumeration expression and the type compatible expression is enumeration expression, and all enumeration expressions must be defined in terms of an enumeration type with the same enumeration literals in the same order.
- If **A** has an **operator record** base-class then **B** must also have an **operator record** base-class, and it must be the same, and otherwise neither **A** nor **B** may have an **operator record** base-class. This is also the **operator record** base-class for the expression e.g. for **if (cond) then A else B**.
- If **A** is derived from **ExternalObject** then **B** must also be derived from **ExternalObject** and they must have the same full name; and otherwise neither **A** nor **B** may be derived from **ExternalObject**. The common full name also defines the type of the expression, e.g. for **if (cond) then A else B**.

Chapter 7

Inheritance, Modification, and Redeclaration

One of the major benefits of object-orientation is the ability to *extend* the behavior and properties of an existing class. The original class, known as the *base class*, is extended to create a more specialized version of that class, known as the *derived class* or *subclass*. In this process, the data and behavior of the original class in the form of variable declarations, equations, and certain other contents are reused, or *inherited*, by the subclass. In fact, the inherited contents is copied from the superclass into the subclass, but before copying certain operations, such as type expansion, checking, and modification, are performed on the inherited contents when appropriate. This chapter describes the inheritance concept in Modelica, together with the related concepts modification and redeclaration.

7.1 Inheritance – Extends Clause

The extends-clause is used to specify inheritance from a base class into an (enclosing) class containing the extends-clause. The syntax of the extends-clause is as follows:

```
extends-clause :  
  extends name [ class-modification ] [ annotation ]
```

The name of the base class is looked up in the partially flattened enclosing class (section 5.2) of the extends-clause. The found base class is flattened with a new environment and the partially flattened enclosing class of the extends-clause. The new environment is the result of merging

- arguments of all enclosing class environments that match names in the flattened base class
- the optional `class-modification` of the extends-clause

in that order.

[Example:

```
class A  
  parameter Real a, b;  
end A;  
  
class B  
  extends A(b = 2);  
end B;  
  
class C  
  extends B(a = 1);  
end C;
```

]

The elements of the flattened base class become elements of the flattened enclosing class, and are added at the place of the extends-clause: specifically components and classes, the equation sections, algorithm sections, optional external-clause, and the contents of the annotation at the end of the class, but excluding import-clauses.

[From the example above we get the following flattened class:

```
class Cinstance
  parameter Real a = 1;
  parameter Real b = 2;
end Cinstance;
```

The ordering of the merging rules ensures that, given classes A and B defined above,

```
class C2
  B bcomp(b = 3);
end C2;
```

yields an instance with `bcomp.b = 3`, which overrides `b = 2`.]

The declaration elements of the flattened base class shall either:

- Not already exist in the partially flattened enclosing class (i.e., have different names).
- The new element is a long form of redeclare or uses the `class extends A` syntax, see section 7.3.
- Be exactly identical to any element of the flattened enclosing class with the same name and the same level of protection (public or protected) and same contents. In this case, the first element in order (can be either inherited or local) is kept. It is recommended to give a warning for this case; unless it can be guaranteed that the identical contents will behave in the same way.

Otherwise the model is incorrect.

[Clarifying order:

```
function A
  input Real a;
  input Real b;
end A;

function B
  extends A;
  input Real a;
end B;
// The inputs of B are "a, b" in that order; and the "input Real a;" is ignored
.
```

]

Equations of the flattened base class that are syntactically equivalent to equations in the flattened enclosing class are discarded. This feature is deprecated, and it is recommended to give a warning when discarding them and for the future give a warning about all forms of equivalent equations due to inheritance.

[Equations that are mathematically equivalent but not syntactically equivalent are not discarded, hence yield an overdetermined system of equations.]

7.1.1 Multiple Inheritance

Multiple inheritance is possible since multiple extends-clauses can be present in a class.

7.1.2 Inheritance of Protected and Public Elements

If an extends-clause is used under the `protected` heading, all elements of the base class become protected elements of the current class. If an extends-clause is a public element, all elements of the base class are inherited with their own protection. The eventual headings `protected` and `public` from the base class

do not affect the consequent elements of the current class (i.e., headings `protected` and `public` are not inherited).

7.1.3 Restrictions on the Kind of Base Class

Since specialized classes of different kinds have different properties, see section 4.6, only specialized classes that are *in some sense compatible* to each other can be derived from each other via inheritance. The following table shows which kind of specialized class can be used in an extends clause of another kind of specialized class (the grey cells mark the few exceptional cases, where a specialized class can be derived from a specialized class of another kind):

Derived Class	Base Class											
	package	operator	function	operator function	type	record	operator record	expandable connector	connector	block	model	class
package	yes											yes
operator		yes										yes
function			yes									yes
operator function			yes	yes								yes
type					yes							yes
record						yes						yes
operator record							yes					yes
expandable connector								yes				yes
connector					yes	yes	yes		yes			yes
block						yes				yes		yes
model						yes				yes	yes	yes
class												yes

If a derived class is inherited from another type of specialized class, then the result is a specialized class of the derived class type.

[For example, if a `block` inherits from a `record`, then the result is a `block`.]

All specialized classes can be derived from `class`, provided that the resulting class fulfills the restriction of the specialized class. A `class` may only contain class-definitions, annotations, and extends-clauses (having any other contents is deprecated).

[It is recommended to use the most specific specialized class.]

The specialized classes `package`, `operator`, `function`, `type`, `record`, `operator record`, and `expandable connector` can only be derived from their own kind and from `class`.

[E.g. a `package` can only be base class for packages. All other kinds of classes can use the import-clause to use the contents of a package.]

[Example:

```

record RecordA
  ...
end RecordA;

package PackageA
  ...
end PackageA;

package PackageB
  extends PackageA; // fine
end PackageB;

model ModelA
  extends RecordA; // fine
end ModelA;

model ModelB
  extends PackageA; // error, inheritance not allowed
end ModelB;
    
```

]

7.1.4 Restrictions on Base Classes and Constraining Types to be Transitively Non-Replaceable

The class name used after `extends` for base-classes and for constraining classes must use a class reference considered transitively non-replaceable, see definition in section 6.3.1. For a replaceable component declaration without constraining clause the class must use a class reference considered transitively non-replaceable.

[The requirement to use a transitively non-replaceable name excludes the long form of `redeclare`, i.e. `redeclare model extends M...` where M must be an inherited replaceable class.]

[The rule for a replaceable component declaration without constraining clause implies that constraining classes are always transitively non-replaceable – both if explicitly given or implicitly by the declaration.]

7.2 Modifications

There are three kinds of constructs in the Modelica language in which modifications can occur:

- Variable declarations.
- Short class declarations.
- Extends-clauses.

A modifier modifies one or more declarations from a class by changing some aspect(s) of the declarations. The most common kind of modifier just changes the *default value* or the *start value* in a binding equation; the value and/or start-value should be compatible with the variable according to section 6.7.

[Example: Modifying the default `start` value of the `altitude` variable:

```
Real altitude(start= 59404);
```

]

A modification (e.g. `C1 c1(x = 5)`) is considered a modification equation, if the modified variable (here: `c1.x`) is a non-parameter variable.

[The modification equation is created, if the modified component (here: `c1`) is also created (see section 4.5). In most cases a modification equation for a non-parameter variable requires that the variable was declared with a declaration equation, see section 4.7; in those cases the declaration equation is replaced by the modification equation.]

A more dramatic change is to modify the *type* and/or the *prefixes* and possibly the *dimension sizes* of a declared element. This kind of modification is called a *redeclaration* (section 7.3) and requires the special keyword `redeclare` to be used in the modifier in order to reduce the risk for accidental modeling errors. In most cases a declaration that can be redeclared must include the prefix `replaceable` (section 7.3). The modifier value (and class for redeclarations) is found in the context in which the modifier occurs, see also section 5.3.1.

[Example: Scope for modifiers:

```
model B
  parameter Real x;
  package Medium = Modelica.Media.PartialMedium;
end B;

model C
  parameter Real x = 2;
  package Medium = Modelica.Media.PartialMedium;
  B b(x = x, redeclare package Medium = Medium);
  // The 'x' and 'Medium' being modified are declared in the model B.
  // The modifiers '= x' and '= Medium' are found in the model C.
end C;
```

```

model D
  parameter Real x = 3;
  package Medium = Modelica.Media.PartialMedium;
  C c(b(x = x, redeclare package Medium = Medium));
  // The 'x' and 'Medium' being modified are declared in the model B.
  // The modifiers '= x' and '= Medium' are found in the model D.
end D;

```

]

When present, the description-string of a modifier overrides the existing description.

7.2.1 Syntax of Modifications and Redeclarations

The syntax is defined in the grammar, appendix B.2.5.

7.2.2 Modification Environment

The modification environment contains arguments which modify elements of the class (e.g., parameter changes). The modification environment is built by merging class modifications, where outer modifications override inner modifications.

[This should not be confused with *inner outer* prefixes described in section 5.4.]

7.2.3 Merging of Modifications

Merging of modifiers means that outer modifiers override inner modifiers. The merging is hierarchical, and a value for an entire non-simple component overrides value modifiers for all components, and it is an error if this overrides a **final** prefix for a component, or if value for a simple component would override part of the value of a non-simple component. When merging modifiers each modification keeps its own **each** prefix.

[Example: The following larger example demonstrates several aspects:

```

class C1
  class C11
    parameter Real x;
  end C11;
end C1;

class C2
  class C21
    ...
  end C21;
end C2;

class C3
  extends C1;
  C11 t(x = 3); // ok, C11 has been inherited from C1
  C21 u; // ok, even though C21 is inherited below
  extends C2;
end C3;

```

The following example demonstrates overriding part of non-simple component:

```

record A
  parameter Real x;
  parameter Real y;
end A;

model B
  parameter A a = A(2, 3);
end B;

```

```

model C
  B b1(a(x = 4)); // Error since attempting to override value for a.x when a
                 // has a value.
end C;
    
```

The modification environment of the declaration of `t` is (`x = 3`). The modification environment is built by merging class modifications, as shown by:

```

class C1
  parameter Real a;
end C1;

class C2
  parameter Real b;
  parameter Real c;
end C2;

class C3
  parameter Real x1; // No default value
  parameter Real x2 = 2; // Default value 2
  parameter C1 x3; // No default value for x3.a
  parameter C2 x4(b = 4); // x4.b has default value 4
  parameter C1 x5(a = 5); // x5.a has default value 5
  extends C1; // No default value for inherited element a
  extends C2(b = 6, c = 77); // Inherited b has default value 6
end C3;

class C4
  extends C3(x2 = 22, x3(a = 33), x4(c = 44), x5 = x3, a = 55, b = 66);
end C4;
    
```

Outer modifications override inner modifications, e.g., `b = 66` overrides the nested class modification of `extends C2(b = 6)`. This is known as merging of modifications: `merge((b = 66), (b = 6))` becomes `(b = 66)`.

A flattening of class `C4` will give an object with the following variables:

Variable	Default value
<code>x1</code>	<i>none</i>
<code>x2</code>	<i>22</i>
<code>x3.a</code>	<i>33</i>
<code>x4.b</code>	<i>4</i>
<code>x4.c</code>	<i>44</i>
<code>x5.a</code>	<i>x3.a</i>
<code>a</code>	<i>55</i>
<code>b</code>	<i>66</i>
<code>c</code>	<i>77</i>

]

7.2.4 Single Modification

Two arguments of a modification shall not modify the same element, attribute, or description-string. When using qualified names the different qualified names starting with the same identifier are merged into one modifier. If a modifier with a qualified name has the `each` or `final` prefix, that prefix is only seen as applied to the final part of the name.

[Example:

```

class C1
  Real x[3];
end C1;
class C2 = C1(x = ones(3), x = ones(3)); // Error: x designated twice
    
```

```

class C3
  class C4
    Real x;
  end C4;
  C4 a(final x.unit = "V", x.displayUnit = "mV", x = 5.0);
  // Ok, different attributes designated (unit, displayUnit and value)
  // identical to:
  C4 b(x(final unit = "V", displayUnit = "mV") = 5.0));
end C3;

```

The following examples are incorrect:

```

m1(r = 1.5, r = 1.6) // Multiple modifier for r (its value)
m1(r = 1.5, r = 1.5) // Multiple modifier for r (its value) – even if identical
m1(r.start = 2, r(start = 3)) // Multiple modifier for r.start
m1(x.r = 1.5 "x", x.r(start = 2.0) "y") // Multiple description–string for x.r
m1(r = R(), r(y = 2)) // Multiple modifier for r.y – both direct value and
                        // part of record

```

The following examples are correct:

```

m1(r = 1.5, r(start = 2.0))
m1(r = 1.6, r "x")
m1(r = R(), r(y(min = 2)))

```

]

7.2.5 Modifiers for Array Elements

The following rules apply to modifiers:

- The **each** keyword on a modifier requires that it is applied in an array declaration/modification, and the modifier is applied individually to each element of the enclosing array (with regard to the position of **each**). In case of nested modifiers this implies it is applied individually to each element of each element of the enclosing array; see example. If the modified element is a vector and the modifier does not contain the **each**-prefix, the modification is split such that the first element in the vector is applied to the first element of the vector of elements, the second to the second element, until the last element of the vector-expression is applied to the last element of the array; it is an error if these sizes do not match. Matrices and general arrays of elements are treated by viewing those as vectors of vectors etc.
- If a nested modifier is split, the split is propagated to all elements of the nested modifier, and if they are modified by the **each**-keyword the split is inhibited for those elements. If the nested modifier that is split in this way contains re-declarations that are split it is illegal.

[Example:

```

model C
  parameter Real a [3];
  parameter Real d;
end C;

model B
  C c[5](each a = {1, 2, 3}, d = {1, 2, 3, 4, 5});
  parameter Real b = 0;
end B;

```

This implies $c[i].a[j] = j$ and $c[i].d = i$.

```

model D
  B b(each c.a = {3, 4, 5}, c.d = {2, 3, 4, 5, 6});
  // Equivalent to:
  B b2(c(each a = {3, 4, 5}, d = {2, 3, 4, 5, 6}));
end D;

```

This implies $b.c[i].a[j] = 2+j$ and $b.c[i].d = 1+i$.

```

model E
  B b[2](each c(each a = {1, 2, 3}, d = {1, 2, 3, 4, 5}), p = {1, 2});
  // Without the first each one would have to use:
  B b2[2](c(each a = {1, 2, 3}, d = fill({1, 2, 3, 4, 5}, 2)), p = {1, 2});
end E;
    
```

This implies $b[k].c[i].a[j] = j$, $b[k].c[i].d = i$, and $b[k].p = k$. For $c.a$ the additional (outer) `each` has no effect, but it is necessary for $c.d$.

Specifying array dimensions after the type works the same as specifying them after the variable name.

```

model F
  Real fail1[2](each start = {1, 2}); // Illegal
  Real work1[2](each start = 1);     // Legal
  Real[2] fail2(each start = {1, 2}); // Illegal
  Real[2] work2(each start = 2);     // Legal
end F;
    
```

]

7.2.6 Final Element Modification Prevention

An element defined as final by the `final` prefix in an element modification or declaration cannot be modified by a modification or by a redeclaration. All elements of a final element are also final.

[Setting the value of a parameter in an experiment environment is conceptually treated as a modification. This implies that a final modification equation of a parameter cannot be changed in a simulation environment.]

[Example: Final component modification.

```

type Angle =
  Real(final quantity = "Angle", final unit = "rad", displayUnit = "deg");

model TransferFunction
  parameter Real b[:] = {1} "numerator coefficient vector";
  parameter Real a[:] = {1, 1} "denominator coefficient vector";
  ...
end TransferFunction;

model PI "PI controller"
  parameter Real k = 1 "gain";
  parameter Real T = 1 "time constant";
  TransferFunction tf(final b = k * {T, 1}, final a = {T, 0});
end PI;

model Test
  PI c1(k = 2, T = 3); // fine, will indirectly change tf.b to 2 * {3, 1}
  PI c2(tf(b = {1})); // error, b is declared as final
end Test;
    
```

]

[Example: Final class declaration.

```

model Test2
  final model MyTF = TransferFunction(b = {1, 2});
  /* Equivalently:
  final model MyTF = TransferFunction(final a, final b = {1, 2});
  */
  MyTF tf1; // fine
  MyTF tf2(a = {1, 2}); // error, all elements in MyTF are final
  model M = MyTF(a = {4}); // error, all elements in MyTF are final
  model TFX
    
```

```

    extends MyTF;           // fine
    Real foo = 1.0;
  end TFX;
  TFX tfx(foo = 2.0);      // fine, foo is not from MyRF
  TFX tfx2(a = {1, 3});   // error, all elements from MyTF are final
  model TFX3 = TFX(a = {1, 4}); // error, all elements from MyTF are final
end Test2;

```

7.3 Redeclaration

A **redeclare** construct in a modifier replaces the declaration of a local class or component with another declaration. A **redeclare** construct as an element replaces the declaration of a local class or component with another declaration. Both **redeclare** constructs work in the same way. The **redeclare** construct as an element requires that the element is inherited, and cannot be combined with a modifier of the same element in the extends-clause. For modifiers the redeclare of classes uses a special short-class-definition construct; that is a subset of normal class definitions and semantically behave as the corresponding class-definition.

A modifier with the keyword **replaceable** is automatically seen as being a **redeclare**.

In redeclarations some parts of the original declaration is automatically inherited by the new declaration. This is intended to make it easier to write declarations by not having to repeat common parts of the declarations, and does in particular apply to prefixes that must be identical. The inheritance only applies to the declaration itself and not to elements of the declaration.

The general rule is that if no prefix within one of the following groups is present in the new declaration the old prefixes of that kind are preserved.

The groups that are valid for both classes and components:

- **public, protected**
- **inner, outer**
- constraining type according to rules in section 7.3.2.

The groups that are only valid for components:

- **flow, stream**
- **discrete, parameter, constant**
- **input, output**
- array dimensions

Note that if the old declaration was a short class definition with array dimensions the array dimensions are not automatically preserved, and thus have to be repeated in the few cases they are used.

Replaceable component array declarations with array sizes on the left of the component are seen as syntactic sugar for having all arrays sizes on the right of the component; and thus can be redeclared in a consistent way.

[Note: The inheritance is from the original declaration. In most cases replaced or original does not matter. It does matter if a user redeclares a variable to be a parameter and then redeclares it without parameter.]

```

[
model HeatExchanger
  replaceable parameter GeometryRecord geometry;
  replaceable input Real u[2];
end HeatExchanger;

HeatExchanger (

```

```

/*redeclare*/ replaceable /*parameter*/ GeoHorizontal geometry,
redeclare /*input*/ Modelica.Units.SI.Angle u /*[2]*/);
// The semantics ensure that parts in /*.*/* are automatically added
// from the declarations in HeatExchanger.
    
```

Example of arrays on the left of the component name:

```

model M
  replaceable Real [4] x[2];
  // Seen as syntactic sugar for "replaceable Real x[2, 4];"
  // Note the order.
end M;
M m(redeclare Modelica.Units.SI.Length x[2, 4]); // Valid redeclare of the type
    
```

7.3.1 The class extends Redeclaration Mechanism

A class declaration of the type `redeclare class extends B(...)`, where `class` as usual can be replaced by any other specialized class, replaces the inherited class `B` with another declaration that extends the inherited class where the optional class-modification is applied to the inherited class. Inherited `B` here means that the class containing `redeclare class extends B(...)` should also inherit another declaration of `B` from one of its extends-clauses. The new declaration should explicitly include `redeclare`.

[Since the rule about applying the optional class-modification implies that all declarations are inherited with modifications applied, there is no need to apply modifiers to the new declaration.]

For `redeclare class extends B(...)` the inherited class is subject to the same restrictions as a redeclare of the inherited element, and the original class `B` should be replaceable, and the new element is only replaceable if the new definition is replaceable. In contrast to normal extends it is not subject to the restriction that `B` should be transitively non-replaceable (since `B` should be replaceable).

The syntax rule for `class extends` construct is in the definition of the `class-specifier` nonterminal (see also class declarations in section 4.5):

```

class-definition :
  [ encapsulated ] class-prefixes
  class-specifier

class-specifier : long-class-specifier | ...

long-class-specifier : ...
  | extends IDENT [ class-modification ] description-string
  composition end IDENT
    
```

The nonterminal `class-definition` is referenced in several places in the grammar, including the following case which is used in some examples below, including `package extends` and `model extends`:

```

element :
  import-clause |
  extends-clause |
  [ redeclare ]
  [ final ]
  [ inner ] [ outer ]
  ( ( class-definition | component-clause ) |
  replaceable ( class-definition | component-clause )
  [constraining-clause comment])
    
```

Example to extend from existing packages:

```

package PowerTrain // library from someone else
  replaceable package GearBoxes
  ...
end GearBoxes;
end PowerTrain;
    
```



```

package MyPowerTrain
  extends PowerTrain; // use all classes from PowerTrain
  redeclare package extends GearBoxes // add classes to sublibrary
  ...
end GearBoxes;
end MyPowerTrain;
    
```

Example for an advanced type of package structuring with constraining types:

```

partial package PartialMedium "Generic medium interface"
  constant Integer nX "number of substances";
  replaceable partial model BaseProperties
    Real X[nX];
    ...
  end BaseProperties;

  replaceable partial function dynamicViscosity
    input Real p;
    output Real eta;...
  end dynamicViscosity;
end PartialMedium;

package MoistAir "Special type of medium"
  extends PartialMedium(nX=2);

  redeclare model extends BaseProperties(T(stateSelect = StateSelect.prefer))
    // replaces BaseProperties by a new implementation and
    // extends from BaseProperties with modification
    // note, nX = 2 (!)
  equation
    X = {0, 1};
    ...
  end BaseProperties;

  redeclare function extends dynamicViscosity
    // replaces dynamicViscosity by a new implementation and
    // extends from dynamicViscosity
  algorithm
    eta := 2 * p;
  end dynamicViscosity;
end MoistAir;
    
```

Note, since MoistAir extends from PartialMedium, constant $nX = 2$ in package MoistAir and the model BaseProperties and the function dynamicViscosity is present in MoistAir. By the following definitions, the available BaseProperties model is replaced by another implementation which extends from the BaseProperties model that has been temporarily constructed during the extends of package MoistAir from PartialMedium. The redeclared BaseProperties model references constant nX which is 2, since by construction the redeclared BaseProperties model is in a package with $nX = 2$.

This definition is compact but is difficult to understand. At a first glance an alternative exists that is more straightforward and easier to understand:

```

package MoistAir2 "Alternative definition that does not work"
  extends PartialMedium(nX=2,
    redeclare model BaseProperties = MoistAir_BaseProperties,
    redeclare function dynamicViscosity = MoistAir_dynamicViscosity);

  model MoistAir_BaseProperties
    // wrong model since nX has no value
    extends PartialMedium.BaseProperties;
  equation
    X = {1, 0};
  end MoistAir_BaseProperties;
    
```

```

model MoistAir_dynamicViscosity
  extends PartialMedium.dynamicViscosity;
algorithm
  eta := p;
end MoistAir_dynamicViscosity;
end MoistAir2;

```

Here, the usual approach is used to extend (here from `PartialMedium`) and in the modifier perform all redeclarations. In order to perform these redeclarations, corresponding implementations of all elements of `PartialMedium` have to be given under a different name, such as `MoistAir2.MoistAir_BaseProperties`, since the name `BaseProperties` already exists due to `extends PartialMedium`. Then it is possible in the modifier to redeclare `PartialMedium.BaseProperties` to `MoistAir2.MoistAir_BaseProperties`. Besides the drawback that the namespace is polluted by elements that have different names but the same implementation (e.g. `MoistAir2.BaseProperties` is identical to `MoistAir2.MoistAir_BaseProperties`) the whole construction does not work if arrays are present that depend on constants in `PartialMedium`, such as `X[nX]`: The problem is that `MoistAir_BaseProperties` extends from `PartialMedium.BaseProperties` where the constant `nX` does not yet have a value. This means that the dimension of array `X` is undefined and model `MoistAir_BaseProperties` is wrong. With this construction, all constant definitions have to be repeated whenever these constants shall be used, especially in `MoistAir_BaseProperties` and `MoistAir_dynamicViscosity`. For larger models this is not practical and therefore the only practically useful definition is the complicated construction in the previous example with `redeclare model extends BaseProperties`.

To detect this issue the rule on lookup of composite names (section 5.3.2) ensures that `PartialMedium.dynamicViscosity` is incorrect in a simulation model.]

7.3.2 Constraining Type

In a replaceable declaration the optional `constraining-clause` defines a constraining type. Any modifications following the constraining type name are applied both for the purpose of defining the actual constraining type and they are automatically applied in the declaration and in any subsequent redeclaration. The precedence order is that declaration modifiers override constraining type modifiers.

If the `constraining-clause` is not present in the original declaration (i.e., the non-redeclared declaration):

- The type of the declaration is also used as a constraining type.
- The modifiers for subsequent redeclarations and constraining type are the modifiers on the component or short-class-definition if that is used in the original declaration, otherwise empty.

The syntax of a constraining-clause is as follows:

```

constraining-clause :
  constrainedby name [ class-modification ]

```

[Example: Merging of modifiers:

```

class A
  parameter Real x;
end A;

class B
  parameter Real x = 3.14, y; // B is a subtype of A
end B;

class C
  replaceable A a(x = 1);
end C;

class D
  extends C(redeclare B a(y = 2));
end D;

```

which is equivalent to defining D as

```
class D
  B a(x = 1, y = 2);
end D;
```

A modification of the constraining type is automatically applied in subsequent redeclarations:

```
model ElectricalSource
  replaceable SineSource source constrainedby M0(final n=5);
  ...
end ElectricalSource;

model TrapezoidalSource
  extends ElectricalSource(
    redeclare Trapezoidal source); // source.n=5
end TrapezoidalSource;
```

A modification of the base type without a constraining type is automatically applied in subsequent redeclarations:

```
model Circuit
  replaceable model NonlinearResistor = Resistor(R=100);
  ...
end Circuit;

model Circuit2
  extends Circuit(
    redeclare replaceable model NonlinearResistor
      = ThermoResistor(T0 = 300));
    // As a result of the modification on the base type,
    // the default value of R is 100
end Circuit2;

model Circuit3
  extends Circuit2(
    redeclare replaceable model NonlinearResistor
      = Resistor(R = 200));
    // The T0 modification is not applied because it did not
    // appear in the original declaration
end Circuit3;
```

Circuit2 is intended to illustrate that a user can still select any resistor model (including the original one, as is done in Circuit3), since the constraining type is kept from the original declaration if not specified in the redeclare. Thus it is easy to select an advanced resistor model, without limiting the possible future changes.

A redeclaration can redefine the constraining type:

```
model Circuit4
  extends Circuit2(
    redeclare replaceable model NonlinearResistor
      = ThermoResistor constrainedby ThermoResistor);
end Circuit4;

model Circuit5
  extends Circuit4(
    redeclare replaceable model NonlinearResistor = Resistor); // illegal
end Circuit5;
```

]

The class or type of component shall be a subtype of the constraining type. In a redeclaration of a replaceable element, the class or type of a component must be a subtype of the constraining type. The constraining type of a replaceable redeclaration must be a subtype of the constraining type of the

declaration it redeclares. In an element modification of a replaceable element, the modifications are applied both to the actual type and to the constraining type.

In an element-redeclaration of a replaceable element the modifiers of the replaced constraining type are merged to both the new declaration and to the new constraining type, using the normal rules where outer modifiers override inner modifiers.

When a class is flattened as a constraining type, the flattening of its replaceable elements will use the constraining type and not the actual default types.

The number of dimension in the constraining type should correspond to the number of dimensions in the type-part. Similarly the type used in a redeclaration must have the same number of dimensions as the type of redeclared element.

[*Example:*

```
replaceable T1 x[n] constrainedby T2;
replaceable type T=T1[n] constrainedby T2;
replaceable T1[n] x constrainedby T2;
```

In these examples the number of dimensions must be the same in T1 and T2, as well as in a redeclaration. Normally T1 and T2 are scalar types, but both could also be defined as array types (with the same number of dimensions). Thus if T2 is a scalar type (e.g. type T2 = Real) then T1 must also be a scalar type, and if T2 is defined as vector type (e.g. type T2 = Real[3]) then T1 must also be vector type.]

7.3.2.1 Constraining-clause annotations

Description and annotations on the constraining-clause are applied to the entire declaration, and it is an error if they also appear on the definition.

[The intent is that the description and/or annotation are at the end of the declaration, but it is not straightforward to specify this in the grammar.]

[*Example:*

```
replaceable model Load1 =
  Resistor constrainedby TwoPin "The Load"; // Recommended
replaceable model Load2 =
  Resistor "The Load" constrainedby TwoPin; // Identical to Load1
replaceable model Load3 =
  Resistor "The Load" constrainedby TwoPin "The Load"; // Error

replaceable Resistor load1
  constrainedby TwoPin "The Load"; // Recommended
replaceable Resistor load2
  "The Load" constrainedby TwoPin; // Identical to load1
replaceable Resistor load3
  "The Load" constrainedby TwoPin "The Load!"; // Error
```

]

See also the examples in section 7.3.4.

7.3.3 Restrictions on Redeclarations

The following additional constraints apply to redeclarations (after prefixes are inherited, section 7.3):

- Only classes and components declared as replaceable can be redeclared with a new type, which must have an interface compatible with the constraining interface of the original declaration, and to allow further redeclarations one must use `redeclare replaceable`.

[Redeclaration with the same type can be used to restrict variability and/or change array dimensions.]

- An element declared as `constant` cannot be redeclared.
- An element declared as `final` shall not be modified, and thus not redeclared.

- Modelica does not allow a protected element to be redeclared as public, or a public element to be redeclared as protected.
- Array dimensions may be redeclared; provided the sub-typing rules in section 6.4 are satisfied.

[*This is one example of redeclaration of non-replaceable elements.*]

7.3.4 Annotation Choices for Suggested Redeclarations and Modifications

A declaration can have an annotation `choices` containing modifiers on `choice`, where each of them indicates a suitable redeclaration or modifications of the element.

This is a hint for users of the model, and can also be used by the user interface to suggest reasonable redeclaration, where the string comments on the choice declaration can be used as textual explanations of the choices. The annotation is not restricted to replaceable elements but can also be applied to non-replaceable elements, enumeration types, and simple variables. For a `Boolean` variable, a `choices` annotation may contain the definition `checkBox = true`, meaning to display a checkbox to input the values `false` or `true` in the graphical user interface.

The annotation `choicesAllMatching = true` on a replaceable element indicates that tools should automatically construct a menu with choices of elements usable for replacing it. Exact criteria for inclusion in such a menu are not defined, but there shall be a way to at least get a selection of classes, `A.B....X.Z`, that are either directly or indirectly derived by inheritance from the constraining class of the declaration, where `A` to `X` are non-partial packages, and `Z` is non-partial. This menu can be disabled using annotation `choicesAllMatching = false`.

[*The behavior when `choicesAllMatching` is not specified; ideally it should present (at least) the same choices as for `choicesAllMatching = true`; but if it takes (too long) time to present the list it is better to have `choicesAllMatching = false`.*]

[*Example:*

```
replaceable model MyResistor = Resistor
  annotation(choices(
    choice(redeclare model MyResistor=lib2.Resistor(a={2}) "..."),
    choice(redeclare model MyResistor=lib2.Resistor2 "...")));

replaceable Resistor Load(R = 2) constrainedby TwoPin
  annotation(choices(
    choice(redeclare lib2.Resistor Load(a={2}) "..."),
    choice(redeclare Capacitor Load(L=3) "...")));

replaceable FrictionFunction a(func = exp) constrainedby Friction
  annotation(choices(
    choice(redeclare ConstantFriction a(c=1) "..."),
    choice(redeclare TableFriction a(table="...") "..."),
    choice(redeclare FunctionFriction a(func=exp) "...")));

replaceable package Medium = Modelica.Media.Water.ConstantPropertyLiquidWater
  constrainedby Modelica.Media.Interfaces.PartialMedium
  annotation (choicesAllMatching = true);
```

It can also be applied to nonreplaceable declarations, e.g. to describe enumerations.

```
type KindOfController = Integer(min = 1, max = 3)
  annotation(choices(
    choice = 1 "P",
    choice = 2 "PI",
    choice = 3 "PID"));

model A
  parameter KindOfController x;
end A;
A a(x = 3 "PID");
```

It can also be applied to Boolean variables to define a check box.

```
parameter Boolean useHeatPort = false annotation(choices(checkBox = true));
```

```
]
```

Chapter 8

Equations

8.1 Equation Categories

Equations in Modelica can be classified into different categories depending on the syntactic context in which they occur:

- *Normal equality equations* occurring in equation sections, including connect-equations and other equation types of special syntactic form (section 8.3).
- *Declaration equations*, which are part of variable, parameter, or constant declarations (section 4.4.2.1).
- *Modification equations*, which are commonly used to modify attributes of classes (section 7.2).
- *Binding equations*, which include both declaration equations and modification equations (for the value itself).
- *Initial equations*, which are used to express equations for solving initialization problems (section 8.6).

8.2 Flattening and Lookup in Equations

A flattened equation is identical to the corresponding nonflattened equation.

Names in an equation shall be found by looking up in the partially flattened enclosing class of the equation.

8.3 Equations in Equation Sections

The following kinds of equations may occur in equation sections. The syntax is defined as follows:

```
equation :  
( simple-expression "=" expression  
  | if-equation  
  | for-equation  
  | connect-clause  
  | when-equation  
  | component-reference function-call-args )  
comment
```

No statements are allowed in equation sections, including the assignment statement using the := operator.

8.3.1 Simple Equality Equations

Simple equality equations are the traditional kinds of equations known from mathematics that express an equality relation between two expressions. There are two syntactic forms of such equations in Modelica.

The first form below is *equality* equations between two expressions, whereas the second form is used when calling a function with *several* results. The syntax for simple equality equations is as follows:

```
simple-expression "=" expression
```

The types of the left-hand-side and the right-hand-side of an equation need to be compatible in the same way as two arguments of binary operators (section 6.7).

Three examples:

- `simple_expr1 = expr2;`
- `(if pred then alt1 else alt2) = expr2;`
- `(out1, out2, out3) = function_name(inexpr1, inexpr2);`

[Note: According to the grammar the if-then-else expression in the second example needs to be enclosed in parentheses to avoid parsing ambiguities. Also compare with section 11.2.1.1 about calling functions with several results in assignment statements.]

8.3.2 For-Equations – Repetitive Equation Structures

The syntax of a for-equation is as follows:

```
for for-indices loop
  { equation ";" }
end for ";"
```

For-equations may optionally use several iterators (for-indices), see section 11.2.2.3 for more information:

```
for-indices:
  for-index {" ," for-index}

for-index:
  IDENT [ in expression ]
```

The following is one example of a prefix of a for-equation:

```
for IDENT in expression loop
```

8.3.2.1 Explicit Iteration Ranges of For-Equations

The **expression** of a for-equation shall be a vector expression, where more general array expressions are treated as vector of vectors or vector of matrices. It is evaluated once for each for-equation, and is evaluated in the scope immediately enclosing the for-equation. The expression of a for-equation shall be a parameter expression. The iteration range of a for-equation can also be specified as **Boolean** or as an enumeration type, see section 11.2.2.2 for more information. The loop-variable (**IDENT**) is in scope inside the loop-construct and shall not be assigned to. For each element of the evaluated vector expression, in the normal order, the loop-variable gets the value of that element and that is used to evaluate the body of the for-loop.

[Example:

```
for i in 1 : 10 loop           // i takes the values 1, 2, 3, ..., 10
for r in 1.0 : 1.5 : 5.5 loop // r takes the values 1.0, 2.5, 4.0, 5.5
for i in {1, 3, 6, 7} loop    // i takes the values 1, 3, 6, 7
for i in TwoEnums loop       // i takes the values TwoEnums.one, TwoEnums.two
                             // for TwoEnums = enumeration(one,two)
```

The loop-variable may hide other variables as in the following example. Using another name for the loop-variable is, however, strongly recommended.

```
constant Integer j = 4;
Real x[j]
equation
  for j in 1:j loop // The loop-variable j takes the values 1,2,3,4
```



```
x[j] = j;           // Uses the loop-variable j
end for;
```

]

8.3.2.2 Implicit Iteration Ranges of For-Equations

The iteration range of a loop-variable may sometimes be inferred from its use as an array index. See section 11.2.2.1 for more information.

[Example:

```
Real x[n],y[n];
for i loop           // Same as: for i in 1:size(x ,1) loop
  x[i] = 2*y[i];
end for;
```

]

8.3.3 Connect-Equations

A connect-equation has the following syntax:

```
connect "(" component-reference "," component-reference ")" ";"
```

These can be placed inside for-equations and if-equations; provided the indices of the for-loop and conditions of the if-clause are parameter expressions that do not depend on `cardinality`, `rooted`, `Connections.rooted`, or `Connections.isRoot`. The for-equations/if-equations are expanded. Connect-equations are described in detail in section 9.1.

The same restrictions apply to `Connections.branch`, `Connections.root`, and `Connections.potentialRoot`; which after expansion are handled according to section 9.4.

8.3.4 If-Equations

If-equations have the following syntax:

```
if expression then
  { equation ";" }
{ elseif expression then
  { equation ";" } }
[ else
  { equation ";" }
]
end if ";"
```

The `expression` of an if- or elseif-clause must be a scalar `Boolean` expression. One if-clause, and zero or more elseif-clauses, and an optional else-clause together form a list of branches. One or zero of the bodies of these if-, elseif- and else-clauses is selected, by evaluating the conditions of the if- and elseif-clauses sequentially until a condition that evaluates to true is found. If none of the conditions evaluate to true the body of the else-clause is selected (if an else-clause exists, otherwise no body is selected). In an equation section, the equations in the body are seen as equations that must be satisfied. The bodies that are not selected have no effect on that model evaluation.

If-equations in equation sections which do not have exclusively parameter expressions as switching conditions shall have the same number of equations in each branch (a missing else is counted as zero equations and the number of equations is defined after expanding the equations to scalar equations).

[If this condition is violated, the single assignment rule would not hold, because the number of equations may change during simulation although the number of unknowns remains the same.]

8.3.5 When-Equations

When-equations have the following syntax:

```

when expression then
  { equation ";" }
{ elsewhen expression then
  { equation ";" } }
end when ";"

```

The **expression** of a when-equation shall be a discrete-time **Boolean** scalar or vector expression. The statements within a when-equation are activated when the scalar expression or any of the elements of the vector expression becomes true.

[Example: The order between the equations in a when-equation does not matter, e.g.

```

equation
  when x > 2 then
    y3 = 2*x + y1 + y2; // Order of y1 and y3 equations does not matter
    y1 = sin(x);
  end when;
  y2 = sin(y1);

```

]

8.3.5.1 Defining When-Equations by If-Expressions in Equality Equations

A when-equation:

```

equation
  when x > 2 then
    v1 = expr1;
    v2 = expr2;
  end when;

```

is conceptually equivalent to the following equations containing special if-expressions

```

// Not correct Modelica
Boolean b(start = x.start > 2);
equation
  b = x > 2;
  v1 = if edge(b) then expr1 else pre(v1);
  v2 = if edge(b) then expr2 else pre(v2);

```

[The equivalence is conceptual since `pre(...)` of a non discrete-time Real variable or expression can only be used within a when-clause. Example:

```

/* discrete */ Real x;
input Real u;
output Real y;
equation
  when sample() then
    x = a * pre(x) + b * pre(u);
  end when;
  y = x;

```

Here, `x` is a discrete-time variable (whether it is declared with the `discrete` prefix or not), but `u` and `y` cannot be discrete-time variables (since they are not assigned in when-clauses). However, `pre(u)` is legal within the when-clause, since the body of the when-clause is only evaluated at events, and thus all expressions are discrete-time expressions.]

The start-values of the introduced **Boolean** variables are defined by the taking the start-value of the when-condition, as above where `b` is a parameter variable. The start-value of the special functions `initial`, `terminal`, and `sample` is `false`.

8.3.5.2 Restrictions on Where a When-Equation may Occur

- When-equations shall not occur inside initial equations.

- When-equations cannot be nested.
- When-equations can only occur within if-equations and for-equations if the controlling expressions are exclusively parameter expressions.

[Example: The following when-equation is invalid:

```
when x > 2 then
  when y1 > 3 then
    y2 = sin(x);
  end when;
end when;
```

]

8.3.5.3 Restrictions on Equations within When-Equations

The equations within the when-equation must have one of the following forms:

- $v = \text{expr};$
- $(\text{out1}, \text{out2}, \text{out3}, \dots) = \text{function_call_name}(\text{in1}, \text{in2}, \dots);$
- Operators `assert`, `terminate`, `reinit`.
- For- and if-equations if the equations within the for- and if-equations satisfy these requirements.
- The different branches of `when/elsewhen` must have the same set of component references on the left-hand side.
- The branches of an if-then-else clause inside when-equations must have the same set of component references on the left-hand side, unless the if-then-else have exclusively parameter expressions as switching conditions.

Any left hand side reference, $(v, \text{out1}, \dots)$, in a when-clause must be a component reference, and any indices must be parameter expressions.

[The needed restrictions on equations within a when-equation becomes apparent with the following example:

```
Real x, y;
equation
  x + y = 5;
  when condition then
    2 * x + y = 7; // error: not valid Modelica
  end when;
```

When the equations of the when-equation are not activated it is not clear which variable to hold constant, either x or y . A corrected version of this example is:

```
Real x,y;
equation
  x + y = 5;
  when condition then
    y = 7 - 2 * x; // fine
  end when;
```

Here, variable y is held constant when the when-equation is deactivated and x is computed from the first equation using the value of y from the previous event instant.]

[Example: The restrictions for if-equations mean that both of the following variants are illegal:

```
Real x, y;
equation
  if time < 1 then
    when sample(1, 2) then
      x = time;
    end when;
```

```

else
  when sample(1, 3) then
    y = time;
  end when;
end if;

when sample(1, 2) then
  if time < 1 then
    y = time;
  else
    x = time;
  end if;
end when;

```

whereas the restriction to parameter-expression is intended to allow:

```

parameter Boolean b = true;
parameter Integer n = 3;
Real x[n];
equation
  if b then
    for i in 1 : n loop
      when sample(i, i) then
        x[i] = time;
      end when;
    end for;
  end if;

```

]

8.3.5.4 Application of the Single-assignment Rule to When-Equations

The Modelica single-assignment rule (section 8.4) has implications for when-equations:

- Two when-equations shall *not* define the same variable.

[Without this rule this may actually happen for the erroneous model `DoubleWhenConflict` below, since there are two equations (`close = true;` `close = false;`) defining the same variable `close`. A conflict between the equations will occur if both conditions would become `true` at the same time instant.

```

model DoubleWhenConflict
  Boolean close; // Erroneous model: close defined by two equations!
equation
  ...
  when condition1 then
    ...
    close = true;
  end when;
  when condition2 then
    close = false;
  end when;
  ...
end DoubleWhenConflict;

```

One way to resolve the conflict would be to give one of the two when-equations higher priority. This is possible by rewriting the when-equation using `elsewhen`, as in the `WhenPriority` model below or using the statement version of the when-construct, see section 11.2.7.]

- When-equations involving `elsewhen`-parts can be used to resolve assignment conflicts since the first of the when/elsewhen parts are given higher priority than later ones:

[Below it is well defined what happens if both conditions become `true` at the same time instant since `condition1` with associated conditional equations has a higher priority than `condition2`.

```

model WhenPriority
  Boolean close; // Correct model: close defined by two equations!
algorithm
  ...
  when condition1 then
    close = true;
  elseif condition2 then
    close = false;
  end when;
  ...
end WhenPriority;

```

8.3.6 reinit

`reinit` can only be used in the body of a when-equation. It has the following syntax:

```
reinit(x, expr);
```

The operator reinitializes `x` with `expr` at an event instant. `x` is a `Real` variable (or an array of `Real` variables) that must be selected as a state (resp., states), i.e. `reinit` on `x` implies `stateSelect = StateSelect.always` on `x`. `expr` needs to be type-compatible with `x`. `reinit` can for the same variable (resp. array of variables) only be applied (either as an individual variable or as part of an array of variables) in one equation (having `reinit` of the same variable in when and else-when of the same variable is allowed). In case of `reinit` active during initialization (due to when initial), see section 8.6.

`reinit` does not break the single assignment rule, because `reinit(x, expr)` in equations evaluates `expr` to a value, then at the end of the current event iteration step it assigns this value to `x` (this copying from values to reinitialized state(s) is done after all other evaluations of the model and before copying `x` to `pre(x)`).

[Example: If a higher index system is present, i.e., constraints between state variables, some state variables need to be redefined to non-state variables. During simulation, non-state variables should be chosen in such a way that variables with an applied `reinit` are selected as states at least when the corresponding when-clauses become active. If this is not possible, an error occurs, since otherwise `reinit` would be applied to a non-state variable.]

Example for the usage of `reinit` (bouncing ball):

```

der(h) = v;
der(v) = if flying then -g else 0;
flying = not (h <= 0 and v <= 0);
when h < 0 then
  reinit(v, -e * pre(v));
end when

```

8.3.7 assert

An equation or statement of one of the following forms is an assertion:

```

assert(condition, message); // Uses level=AssertionLevel.error
assert(condition, message, assertionLevel);
assert(condition, message, level = assertionLevel);

```

Here, `condition` is a `Boolean` expression, `message` is a `String` expression, and `assertionLevel` is an optional parameter expression of the built-in enumeration type `AssertionLevel`. It can be used in equation sections or algorithm sections.

[This means that `assert` can be called as if it were a function with three formal parameters, the third formal parameter has the name `level` and the default value `AssertionLevel.error`.]

[A parameter expression is required for `level` since it shall be evaluated at compile time.]

If the `condition` of an assertion is true, `message` is not evaluated and the procedure call is ignored. If the `condition` evaluates to false, different actions are taken depending on the `level` input:

- `level = AssertionLevel.error`: The current evaluation is aborted. The simulation may continue with another evaluation. If the simulation is aborted, `message` indicates the cause of the error.

[Ways to continue simulation with another evaluation include using a shorter step-size, or changing the values of iteration variables.]

Failed assertions take precedence over successful termination, such that if the model first triggers the end of successful analysis by reaching the stop-time or explicitly with `terminate`, but the evaluation with `terminal()=true` triggers an assert, the analysis failed.

- `level = AssertionLevel.warning`: The current evaluation is not aborted. `message` indicates the cause of the warning.

[It is recommended to report the warning only once when the condition becomes false, and it is reported that the condition is no longer violated when the condition returns to true. The `assert` statement shall have no influence on the behavior of the model. For example, by evaluating the condition and reporting the message only after accepted integrator steps. `condition` needs to be implicitly treated with `noEvent` since otherwise events might be triggered that can lead to slightly changed simulation results.]

[The `AssertionLevel.error` case can be used to avoid evaluating a model outside its limits of validity; for instance, a function to compute the saturated liquid temperature cannot be called with a pressure lower than the triple point value.

The `AssertionLevel.warning` case can be used when the boundary of validity is not hard: for instance, a fluid property model based on a polynomial interpolation curve might give accurate results between temperatures of 250 K and 400 K, but still give reasonable results in the range 200 K and 500 K. When the temperature gets out of the smaller interval, but still stays in the largest one, the user should be warned, but the simulation should continue without any further action. The corresponding code would be:

```
assert(T > 250 and T < 400, "Medium model outside full accuracy range",
  AssertionLevel.warning);
assert(T > 200 and T < 500, "Medium model outside feasible region");
```

]

8.3.8 terminate

The `terminate` equation or statement (using function syntax) successfully terminates the analysis which was carried out, see also section 8.3.7. The termination is not immediate at the place where it is defined since not all variable results might be available that are necessary for a successful stop. Instead, the termination actually takes place when the current integrator step is successfully finalized or at an event instant after the event handling has been completed before restarting the integration.

The `terminate` statement has a string argument indicating the reason for the success.

[Example: The intention of the `terminate` statement is to give more complex stopping criteria than a fixed point in time:

```
model ThrowingBall
  Real x(start=0);
  Real y(start=1);
equation
  der(x) = ...;
  der(y) = ...;
algorithm
  when y < 0 then
    terminate("The ball touches the ground");
  end when;
end ThrowingBall;
```

]

8.3.9 Equation Operators for Overconstrained Connection-Based Equation Systems

See section 9.4 for a description of this topic.

8.4 Synchronous Data-flow Principle and Single Assignment Rule

Modelica is based on the synchronous data flow principle and the single assignment rule, which are defined in the following way:

1. Discrete-time variables keep their values until these variables are explicitly changed. Differentiated variables have `der(x)` corresponding to the time-derivative of `x`, and `x` is continuous, except when `reinit` is triggered, see section 8.3.6. Variable values can be accessed at any time instant during continuous integration and at event instants.
2. At every time instant, during continuous integration and at event instants, the equations express relations between variables which have to be fulfilled concurrently.
3. Computation and communication at an event instant does not take time.

[If computation or communication time has to be simulated, this property has to be explicitly modeled.]

4. There must exist a perfect matching of variables to equations after flattening, where a variable can only be matched to equations that can contribute to solving for the variable (*perfect matching rule* – previously called *single assignment rule*); see also globally balanced section 4.7.

8.5 Events and Synchronization

The integration is halted and an event occurs whenever an event generation expression, e.g. `x > 2` or `floor(x)`, changes its value. An event generating expression has an internal buffer, and the value of the expression can only be changed at event instants. If the evaluated expression is inconsistent with the buffer, that will trigger an event and the buffer will be updated with a new value at the event instant. During continuous integration event generation expression has the constant value of the expression from the last event instant.

[A root finding mechanism is needed which determines a small time interval in which the expression changes its value; the event occurs at the right side of this interval.]

[Example:

```
y = if u > uMax then uMax else if u < uMin then uMin else u;
```

During continuous integration always the same if-branch is evaluated. The integration is halted whenever `u-uMax` or `u-uMin` crosses zero. At the event instant, the correct if-branch is selected and the integration is restarted.

Numerical integration methods of order n ($n \geq 1$) require continuous model equations which are differentiable up to order n . This requirement can be fulfilled if `Real` elementary relations are not treated literally but as defined above, because discontinuous changes can only occur at event instants and no longer during continuous integration.]

[It is a quality of implementation issue that the following special relations

```
time >= discrete expression
time < discrete expression
```

trigger a time event at `time = discrete expression`, i.e., the event instant is known in advance and no iteration is needed to find the exact event instant.]

Relations are taken literally also during continuous integration, if the relation or the expression in which the relation is present, are the argument of `noEvent`. `smooth` also allows relations used as argument to be taken literally. The `noEvent` feature is propagated to all subrelations in the scope of the `noEvent` application. For `smooth` the liberty to not allow literal evaluation is propagated to all subrelations, but the smoothness property itself is not propagated.

[Example:

```
x = if noEvent(u > uMax) then uMax elseif noEvent(u < uMin) then uMin else u;
y = noEvent( if u > uMax then uMax elseif u < uMin then uMin else u);
z = smooth(0, if u > uMax then uMax elseif u < uMin then uMin else u);
```

In this case `x = y = z`, but a tool might generate events for `z`. The `if-expression` is taken literally without inducing state events.

The `smooth` operator is useful, if e.g. the modeler can guarantee that the used `if-clauses` fulfill at least the continuity requirement of integrators. In this case the simulation speed is improved, since no state event iterations occur during integration. The `noEvent` operator is used to guard against outside domain errors, e.g. `y = if noEvent(x >= 0) then sqrt(x) else 0.`]

All equations and assignment statements within when-clauses and all assignment statements within `function` classes are implicitly treated with `noEvent`, i.e., relations within the scope of these operators never induce state or time events.

[Using state events in when-clauses is unnecessary because the body of a when-clause is not evaluated during continuous integration.]

[Example:

```
Limit1 = noEvent(x1 > 1); // Error since Limit1 is a discrete-time variable
when noEvent(x1>1) or x2>10 then // error, when-conditions is not a discrete-
    time expression
    Close = true;
end when;
```

]

Modelica is based on the synchronous data flow principle (section 8.4).

[The rules for the synchronous data flow principle guarantee that variables are always defined by a unique set of equations. It is not possible that a variable is e.g. defined by two equations, which would give rise to conflicts or non-deterministic behavior. Furthermore, the continuous and the discrete parts of a model are always automatically “synchronized”. Example:

```
equation // Illegal example
  when condition1 then
    close = true;
  end when;

  when condition2 then
    close = false;
  end when;
```

This is not a valid model because rule 4 is violated since there are two equations for the single unknown variable `close`. If this would be a valid model, a conflict occurs when both conditions become true at the same time instant, since no priorities between the two equations are assigned. To become valid, the model has to be changed to:

```
equation
  when condition1 then
    close = true;
  elseif condition2 then
    close = false;
```



```
end when;
```

Here, it is well-defined if both conditions become true at the same time instant (condition1 has a higher priority than condition2).]

There is no guarantee that two different events occur at the same time instant.

[As a consequence, synchronization of events has to be explicitly programmed in the model, e.g. via counters. Example:

```

Boolean fastSample, slowSample;
Integer ticks(start=0);
equation
  fastSample = sample(0,1);
algorithm
  when fastSample then
    ticks      := if pre(ticks) < 5 then pre(ticks)+1 else 0;
    slowSample := pre(ticks) == 0;
  end when;
algorithm
  when fastSample then // fast sampling
    ...
  end when;
algorithm
  when slowSample then // slow sampling (5-times slower)
    ...
  end when;
```

The `slowSample` when-clause is evaluated at every 5th occurrence of the `fastSample` when-clause.]

[The single assignment rule and the requirement to explicitly program the synchronization of events allow a certain degree of model verification already at compile time.]

8.6 Initialization, initial equation, and initial algorithm

Before any operation is carried out with a Modelica model (e.g., simulation or linearization), initialization takes place to assign consistent values for all variables present in the model. During this phase, also the derivatives (`der`), and the pre-variables (`pre`), are interpreted as unknown algebraic variables. The initialization uses all equations and algorithms that are utilized in the intended operation (such as simulation or linearization).

The equations of a when-clause are active during initialization, if and only if they are explicitly enabled with `initial()`, and only in one of the two forms `when initial() then` or `when {..., initial(), ...} then` (and similarly for `elsewhen` and algorithms see below). In this case, the when-clause equations remain active during the whole initialization phase. In case of a `reinit(x, expr)` being active during initialization (due to being inside `when initial()`) this is interpreted as adding `x = expr` (the `reinit`-equation) as an initial equation.

[If a when-clause equation `v = expr`; is not active during the initialization phase, the equation `v = pre(v)` is added for initialization. This follows from the mapping rule of when-clause equations. If the condition of the when-clause contains `initial()`, but not in one of the specific forms, the when-clause is not active during initialization: `when not initial() then print("simulation started"); end when`];

The algorithmic statements within a when-statement are active during initialization, if and only they are explicitly enabled with `initial()`, and only in one of the two forms `when initial() then` or `when {..., initial(), ...} then`. In this case, the algorithmic statements within the when-statement remain active during the whole initialization phase.

An active when-clause inactivates the following `elsewhen` (similarly as for when-clauses during simulation), but apart from that the first `elsewhen initial() then` or `elsewhen {..., initial(), ...} then` is similarly active during initialization as `when initial() then` or `when {..., initial(), ...} then`.

[That means that any subsequent `elsewhen initial()` has no effect, similarly as `when false then`.]

[There is no special handling of inactive when-statements during initialization, instead variables assigned in when-statements are initialized using `v := pre(v)` before the body of the algorithm (since they are discrete), see section 11.1.2.]

Further constraints, necessary to determine the initial values of all variables, can be defined in the following ways:

1. As equations in an `initial equation` section or as assignments in an `initial algorithm` section. The equations and assignments in these initial sections are purely algebraic, stating constraints between the variables at the initial time instant. It is not allowed to use when-clauses in these sections.
2. For a non-discrete (that is continuous-time) `Real` variable `vc`, the equation `pre(vc) = vc` is added to the initialization equations.

[If `pre(vc)` is not present in the flattened model, a tool may choose not to introduce this equation, or if it was introduced it can eliminate it (to avoid the introduction of many dummy variables `pre(vc)`).]
3. Implicitly by using the `start` attribute for variables with `fixed = true`. With `start` given by `startExpression`:
 - For a non-discrete-time (that is continuous-time) `Real` variable `vc`, the equation `vc = startExpression` is added to the initialization equations.
 - For a discrete-time variable `vd`, the equation `pre(vd) = startExpression` is added to the initialization equations.
 - For a variable declared as `constant` or `parameter`, no equation is added to the initialization equations.

For constants and parameters, the attribute `fixed` defaults to `true`, which is the only allowed value for a constant. For other variables, `fixed` defaults to `false`.

`start`-values of variables having `fixed = false` can be used as initial guesses, in case iterative solvers are used in the initialization phase.

[In case of iterative solver failure, it is recommended to specially report those variables for which the solver needs an initial guess, but which only have the default value of the `start` attribute as defined in section 4.8, since the lack of appropriate initial guesses is a likely cause of the solver failure.]

If a parameter has a modifier for the `start` attribute, does not have `fixed = false`, and neither has a binding equation nor is part of a record having a binding equation, the modifier for the `start` attribute can be used to add a parameter binding equation assigning the parameter to that `start` value. In this case a diagnostic message is recommended in a simulation model.

[This is used in libraries to give non-zero defaults so that users can quickly combine models and simulate without setting parameters; but still easily find the parameters that need to be set.]

All variables declared as `parameter` having `fixed = false` are treated as unknowns during the initialization phase, i.e. there must be additional equations for them – and the `start`-value can be used as a guess-value during initialization.

[In the case a parameter has both a binding equation and `fixed = false` a diagnostics is recommended, but the parameter should be solved from the binding equation.

Non-discrete (that is continuous-time) `Real` variables `vc` have exactly one initialization value since the rules above assure that during initialization `vc = pre(vc) = vc.startExpression` (if `fixed = true`).

Before the start of the integration, it must be guaranteed that for all variables `v`, `v = pre(v)`. If this is not the case for some variables `vi`, `pre(vi) := vi` must be set and an event iteration at the initial time must follow, so the model is re-evaluated, until this condition is fulfilled.

A Modelica translator may first transform the continuous equations of a model, at least conceptually, to state space form. This may require to differentiate equations for index reduction, i.e., additional equations and, in some cases, additional unknown variables are introduced. This whole set of equations, together

with the additional constraints defined above, should lead to an algebraic system of equations where the number of equations and the number of all variables (including `der` and `pre` variables) is equal. Often, this is a nonlinear system of equations and therefore it may be necessary to provide appropriate guess values (i.e., `start` values and `fixed = false`) in order to compute a solution numerically.

It may be difficult for a user to figure out how many initial equations have to be added, especially if the system has a higher index. A tool may add or remove initial equations automatically such that the resulting system is structurally nonsingular. In these cases diagnostics are appropriate since the result is not unique and not necessarily what the user expects. A missing initial value of a discrete-time variable which does not influence the simulation result, may be automatically set to the `start` value or its default without informing the user. For example, variables assigned in a `when`-clause which are not accessed outside of the `when`-clause and where `pre` is not explicitly used on these variables, do not have an effect on the simulation.]

[Example: Continuous time controller initialized in steady-state:

```
Real y(fixed = false); // fixed=false is redundant
equation
  der(y) = a * y + b * u;
initial equation
  der(y) = 0;
```

This has the following solution at initialization:

```
der(y) = 0;
y = - b / a * u;
```

]

[Example: Continuous time controller initialized either in steady-state or by providing a `start` value for state `y`:

```
parameter Boolean steadyState = true;
parameter Real y0 = 0 "start value for y, if not steadyState";
Real y;
equation
  der(y) = a * y + b * u;
initial equation
  if steadyState then
    der(y) = 0;
  else
    y = y0;
  end if;
```

This can also be written as follows (this form is less clear):

```
parameter Boolean steadyState = true;
Real y (start = 0, fixed = not steadyState);
Real der_y(start = 0, fixed = steadyState) = der(y);
equation
  der(y) = a * y + b * u;
```

]

[Example: Discrete time controller initialized in steady-state:

```
discrete Real y;
equation
  when {initial(), sampleTrigger} then
    y = a * pre(y) + b * u;
  end when;
initial equation
  y = pre(y);
```

This leads to the following equations during initialization:

```
y = a * pre(y) + b * u;
y = pre(y);
```

with the solution:

```
y := (b * u) / (1 - a);
pre(y) := y;
```

]

8.6.1 The Number of Equations Needed for Initialization

[In general, for the case of a pure (first order) ordinary differential equation (ODE) system with n state variables and m output variables, we will have $n + m$ unknowns in the simulation problem. The ODE initialization problem has n additional unknowns corresponding to the derivative variables. At initialization of an ODE we will need to find the values of $2n + m$ variables, in contrast to just $n + m$ variables to be solved for during simulation.]

[Example: Consider the following simple equation system:

```
der(x1) = f1(x1);
der(x2) = f2(x2);
y = x1+x2+u;
```

Here we have three variables with unknown values: two dynamic variables that also are state variables, $x1$ and $x2$, i.e., $n = 2$, one output variable y , i.e., $m = 1$, and one input variable u with known value. A consistent solution of the initial value problem providing initial values for $x1$, $x2$, $der(x1)$, $der(x2)$, and y needs to be found. Two additional initial equations thus need to be provided to solve the initialization problem.

Regarding DAEs, only that at most n additional equations are needed to arrive at $2n + m$ equations in the initialization system. The reason is that in a higher index DAE problem the number of dynamic continuous-time state variables might be less than the number of state variables n . As noted in section 8.6 a tool may add/remove initial equations to fulfill this requirement, if appropriate diagnostics are given.]

8.6.2 Recommended selection of start-values

In general many variables have start-values that are not fixed and selecting a sub-set of these can give a consistent set of start-values close to the user-expectations. The following gives a non-normative procedure for finding such a sub-set.

[A model has a hierarchical component structure. Each component of a model can be given a unique model component hierarchy level number. The top level model has a level number of 1. The level number increases by 1 for each level down in the model component hierarchy. The model component hierarchy level number is used to give start values a confidence number, where a lower number means that the start value is more confident. Loosely, if the start value is set or modified on level i then the confidence number is i . If a start value is set by a possibly hierarchical modifier at the top level, then this start value has the highest confidence, namely 1 irrespectively on what level, the variable itself is declared.]

Chapter 9

Connectors and Connections

This chapter covers connectors, connect-equations, and connections.

Connectors and connect-equations are designed so that different components can be connected graphically with well-defined semantics. However, the graphical part is optional and found in chapter 18.

9.1 Connect-Equations and Connectors

Connections between objects are introduced by connect-equations in the equation part of a class. A connect-equation has the following syntax:

```
connect(component-reference, component-reference);
```

[A connector is an instance of a `connector` class.]

The connect-equation construct takes two references to connectors, each of which is either of the following forms:

- $c_1.c_2\dots c_n$, where c_1 is a connector of the class, $n \geq 1$ and c_{i+1} is a connector element of c_i for $i=1:(n-1)$.
- $m.c$, where m is a non-connector element in the class and c is a connector element of m .

There may optionally be array subscripts on any of the components; the array subscripts shall be parameter expressions or the special operator `.`. If the connect construct references array of connectors, the array dimensions must match, and each corresponding pair of elements from the arrays is connected as a pair of scalar connectors.

[Example: Array usage:

```
connector InPort = input Real;
connector OutPort = output Real;
block MatrixGain
  input InPort u[size(A,2)];
  output OutPort y[size(A,1)];
  parameter Real A[:,:] = [1];
equation
  y=A*u;
end MatrixGain;
Modelica.Blocks.Sources.Sine sinSource[5];
MatrixGain gain (A = 5*identity(5));
MatrixGain gain2(A = ones(2,5));
OutPort x[2];
equation
  connect(sinSource.y, gain.u); // Legal
  connect(gain.y, gain2.u); // Legal
  connect(gain2.y, x); // Legal
```

]

The three main tasks are to:

- Elaborate expandable connectors.
- Build connection sets from connect-equations.
- Generate equations for the complete model.

9.1.1 Connection Sets

A connection set is a set of variables connected by means of connect-equations. A connection set shall contain either only flow variables or only non-flow variables.

9.1.2 Inside and Outside Connectors

In an element instance M , each connector element of M is called an *outside connector* with respect to M . Any other connector elements that is hierarchically inside M , but not in one of the outside connectors of M , is called an *inside connector* with respect to M . This is done before resolving **outer** elements to corresponding **inner** ones.

[Example:




Figure 9.1: Example for inside and outside connectors.

The figure visualizes the following `connect` equations to the connector `c` in the models m_i . Consider the following `connect` equations found in the model for component `m0`:

```
connect(m1.c, m3.c); // m1.c and m3.c are inside connectors
connect(m2.c, m3.c); // m2.c and m3.c are inside connectors
```

and in the model for component `m3` (`c.x` is a sub-connector inside `c`):

```
connect(c, m4.c); // c is an outside
connector, m4.c is an inside connector
connect(c.x, m5.c); // c.x is an outside
connector, m5.c is an inside connector
connect(c, d); // c is an outside connector, d is an outside connector
```

and in the model for component `m6`:

```
connect(d, m7.c); // d is an outside connector, m7.c is an inside connector
```

]

9.1.3 Expandable Connectors

If the `expandable` qualifier is present on a connector definition, all instances of that connector are referred to as expandable connectors. Instances of connectors that do not possess this qualifier will be referred to as non-expandable connectors.

Before generating connection equations non-parameter scalar variables and non-parameter array elements declared in expandable connectors are marked as only being potentially present. A non-parameter array element may be declared with array dimensions `:` indicating that the size is unknown (note that the size of such a dimension cannot be determined using `size`, see section 10.3.1). This applies to both variables of simple types, and variables of structured types.

Then connections containing expandable connectors are elaborated:

- One connector in the connect equation must reference a declared component, and if the other connector is an undeclared element in a declared expandable connector it is handled as follows (elements that are only potentially present are not seen as declared):
 - The expandable connector instance is automatically augmented with a new component having the used name and corresponding type.
 - If the undeclared component is subscripted, an array variable is created, and a connection to the specific array element is performed. Introducing elements in an array gives an array with at least the specified elements, other elements are either not created or have a default value (i.e. as if they were only potentially present, and the same note regarding the use of `size` also applies here).
 - If the variable on the other side of the connect-equation is `input` or `output` the new component will be either `input` or `output` to satisfy the restrictions in section 9.3 for a non-expandable connector.

[If the existing side refers to an inside connector (i.e. a connector of a component) the new variable will copy its causality, i.e. `input` if `input` and `output` if `output`, since the expandable connector must be an outside connector.]

For an array the `input/output` property can be deduced separately for each array element.

- When two expandable connectors are connected, each is augmented with the variables that are only declared in the other expandable connector (the new variables are neither `input` nor `output`). This is repeated until all connected expandable connector instances have matching variables.

[I.e. each of the connector instances is expanded to be the union of all connector variables.]

- The variables introduced in the elaboration follow additional rules for generating connection sets (given in section 9.2).
- If a variable appears as an `input` in one expandable connector, it should appear as a non-`input` in at least one other expandable connector instance in the same augmentation set. An augmentation set is defined as the set of connected expandable connector instances that through the elaboration will have matching variables.

[Example:

```

expandable connector EngineBus
end EngineBus;

block Sensor
  RealOutput speed; // Output, i.e., non-input
end Sensor;
block Actuator
  RealInput speed; // Input
end Actuator;

model Engine
  EngineBus bus;
  Sensor sensor;
  Actuator actuator;

```

```

equation
  connect(bus.speed, sensor.speed); // provides the non-input from sensor.
  speed
  connect(bus.speed, actuator.speed);
end Engine;

```

]

- All components in an expandable connector are seen as connector instances even if they are not declared as such.

[I.e. it is possible to connect to e.g. a Real variable.]

[Example:

```

expandable connector EngineBus // has predefined signals
  import Modelica.Units.SI;
  SI.AngularVelocity speed;
  SI.Temperature T;
end EngineBus;

block Sensor
  RealOutput speed;
end Sensor;

model Engine
  EngineBus bus;
  Sensor sensor;
equation
  connect(bus.speed, sensor.speed);
  // connection to non-connector speed is possible
  // in expandable connectors
end Engine;

```

]

- An expandable connector shall not contain a component declared with the prefix `flow`, but may contain non-expandable connector components with `flow` components.

[Example:

```

import Interfaces=Modelica.Electrical.Analog.Interfaces;
expandable connector ElectricalBus
  Interfaces.PositivePin p12, n12; // OK
  flow Modelica.Units.SI.Current i; // Error
end ElectricalBus;

model Battery
  Interfaces.PositivePin p42, n42;
  ElectricalBus bus;
equation
  connect(p42, bus.p42); // Adds new electrical pin
  connect(n42, bus.n42); // Adds another pin
end Battery;

```

]

- expandable connectors can only be connected to other expandable connectors.

If a connect equation references a potentially present variable, or variable element, in an expandable connector the variable or variable element is marked as being present, and due to the paragraphs above it is possible to deduce whether the bus variable shall be treated as input, or shall be treated as output in the connect equation. That `input` or `output` prefix is added if no `input/output` prefix is present on the declaration.

[Example:


```

expandable connector EmptyBus
end EmptyBus;

model Controller
  EmptyBus bus1;
  EmptyBus bus2;
  RealInput speed;
equation
  connect(speed, bus1.speed); // OK; only one undeclared and not subscripted.
  connect(bus1.pressure, bus2.pressure); // Error; both undeclared.
  connect(speed, bus2.speed[2]); // Introduces speed array (with element [2]).
end Controller;

```

]

An expandable connector array component for which `size` is not defined (see section 10.3.1) is referred to as a *sizeless array component*. Such a component shall not be used without subscripts, and the subscripts must slice the array so that the sizeless dimensions are removed.

[*Example: Valid and invalid uses of sizeless array components:*

```

expandable connector EngineBus
end EngineBus;

block Sensor
  RealOutput speed;
end Sensor;

model Engine
  parameter Integer n = 1;
  EngineBus bus;
  Sensor sensor;
  RealOutput sensorSpeeds[:];
equation
  /* Comments below refer to the use of sizeless array bus.speed. */
  connect(bus.speed[n], sensor.speed); // OK; subscript to scalar component.
  connect(bus.speed, sensorSpeeds); // Error; missing subscripts.
public
  Real s[:] = bus.speed; // Error; missing subscripts.
  Real s[2] = bus.speed[{1, 3}]; // OK; subscript selects fixed size sub-array.
end Engine;

```

]

After this elaboration the expandable connectors are treated as normal connector instances, and the connections as normal connections, and all potentially present variables and array elements that are not actually present are undefined. It is an error if there are expressions referring to potentially present variables or array elements that are not actually present or non-declared variables. This elaboration implies that expandable connectors can be connected even if they do not contain the same components.

[*A tool may remove undefined variables in an expandable connector, or set them to the default value, e.g. zero for Real variables.*]

[*Expressions can only “read” variables from the bus that are actually declared and present in the connector, in order that the types of the variables can be determined in the local scope.*]

[*Note that the introduction of variables, as described above, is conceptual and does not necessarily impact the flattening hierarchy in any way. Furthermore, it is important to note that these elaboration rules must consider:*

1. *Expandable connectors nested hierarchically. This means that both outside and inside connectors must be included at every level of the hierarchy in this elaboration process.*
2. *When processing an expandable connector that possesses the `inner` scope qualifier, all outer instances must also be taken into account during elaboration.*

]

[Example: Engine system with sensors, controllers, actuator and plant that exchange information via a bus (i.e. via expandable connectors):

```

import Modelica.Units.SI;
import Modelica.Blocks.Interfaces.RealInput;
// Plant Side
model SparkPlug
  RealInput spark_advance;
  ...
end SparkPlug;

expandable connector EngineBus
  // No minimal set
end EngineBus;

expandable connector CylinderBus
  Real spark_advance;
end CylinderBus;

model Cylinder
  CylinderBus cylinder_bus;
  SparkPlug spark_plug;
  ...
equation
  connect(spark_plug.spark_advance,
    cylinder_bus.spark_advance);
end Cylinder;

model I4
  EngineBus engine_bus;
  Modelica.Mechanics.Rotational.Sensors.SpeedSensor speed_sensor;
  Modelica.Thermal.HeatTransfer.Sensors.TemperatureSensor temp_sensor;
  parameter Integer nCylinder = 4 "Number of cylinders";
  Cylinder cylinder[nCylinder];
equation
  // adds engine_speed (as output)
  connect(speed_sensor.w, engine_bus.engine_speed);
  // adds engine_temp (as output)
  connect(temp_sensor.T, engine_bus.engine_temp);
  // adds cylinder_bus1 (a nested bus)
  for i in 1:nCylinder loop
    connect(cylinder[i].cylinder_bus,
      engine_bus.cylinder_bus[i]);
  end for;
end I4;

```

Due to the above connection, conceptually a connector consisting of the union of all connectors is introduced.

The engine_bus contains the following variable declarations:

```

RealOutput engine_speed;
RealOutput engine_temp;
CylinderBus cylinder_bus[1];
CylinderBus cylinder_bus[2];
CylinderBus cylinder_bus[3];
CylinderBus cylinder_bus[4];

```

]

9.2 Generation of Connection Equations

When generating connection equations, **outer** elements are resolved to the corresponding **inner** elements in the instance hierarchy (see instance hierarchy name lookup section 5.4). The arguments to each **connect**-equation are resolved to two connector elements.

For every use of the **connect**-equation

```
connect(a, b);
```

the primitive components of **a** and **b** form a connection set, together with an indication of whether they are from an inside or an outside connector. The primitive elements are of simple types or of types defined as **operator record** (i.e. a component of an **operator record** type is not split into sub-components). The elements of the connection sets are tuples of primitive variables together with an indication of inside or outside; if the same tuple belongs to two connection sets those two sets are merged, until every tuple is only present in one set. Composite connector types are broken down into primitive components. The **outer** components are handled by mapping the objects to the corresponding **inner** components, and the inside indication is not influenced. The outer connectors are handled by mapping the objects to the corresponding inner connectors, and they are always treated as outside connectors.

[Rationale: The inside/outside as part of the connection sets ensure that connections from different hierarchical levels are treated separately. Connection sets are formed from the primitive elements and not from the connectors; this handles connections to parts of hierarchical connectors and also makes it easier to generate equations directly from the connection sets. All variables in one connection set will either be flow variables or non-flow variables due to restriction on connect-equations. The mapping from an outer to an inner element must occur before merging the sets in order to get one zero-sum equation, and ensures that the equations for the outer elements are all given for one side of the connector, and the inner element can define the other side.]

The following connection sets with just one member are also present (and merged):

- Each primitive flow variable as inside connector.
- Each flow variable *added* during augmentation of expandable connectors, both as inside and as outside.

[Note that the flow variable is not directly in the expandable connector, but in a connector inside the expandable connector.]

[Rationale: If these variables are not connected they will generate a set comprised only of this element, and thus they will be implicitly set to zero (see below). If connected, this set will be merged and adding this at the start has no impact.]

Each connection set is used to generate equations for potential and flow (zero-sum) variables of the form

- $a_1 = a_2 = \dots = a_n$ (neither flow nor stream variables)
- $z_1 + z_2 + (-z_3) + \dots + z_n = \mathbf{0}$ (flow variables)

The bold-face **0** represents an array or scalar zero of appropriate dimensions (i.e. the same size as z).

For an **operator record** type this uses the operator '0' – which must be defined in the operator record – and all of the flow variables for the **operator record** must be of the same **operator record** type. This implies that in order to have flow variables of an **operator record** type the **operator record** must define addition, negation, and '0'; and these operations should define an additive group.

In order to generate equations for flow variables (using the **flow** prefix), the sign used for the connector variable z_i above is +1 for inside connectors and -1 for outside connectors (z_3 in the example above).

[Example: Simple example:]

```
model Circuit
  Ground ground;
  Load load;
  Resistor resistor;
equation
  connect(load.p , ground.p);
```

```

    connect(resistor.p, ground.p);
end Circuit;

model Load
  extends TwoPin;
  Resistor resistor;
equation
  connect(p, resistor.p);
  connect(resistor.n, n);
end Load;

```

The connection sets are before merging (note that one part of the load and resistor is not connected):

```

{<load.p.i, inside>}
{<load.n.i, inside>}
{<ground.p.i, inside>}
{<load.resistor.p.i, inside>}
{<load.resistor.n.i, inside>}
{<resistor.p.i, inside>}
{<resistor.n.i, inside>}
{<resistor.p.i, inside>, <ground.p.i, inside>}
{<resistor.p.v, inside>, <ground.p.v, inside>}
{<load.p.i, inside>, <ground.p.i, inside>}
{<load.p.v, inside>, <ground.p.v, inside>}
{<load.p.i, outside>, <load.resistor.p.i, inside>}
{<load.p.v, outside>, <load.resistor.p.v, inside>}
{<load.n.i, outside>, <load.resistor.n.i, inside>}
{<load.n.v, outside>, <load.resistor.n.v, inside>}

```

After merging this gives:

```

{<load.p.i, outside>, <load.resistor.p.i, inside>}
{<load.p.v, outside>, <load.resistor.p.v, inside>}
{<load.n.i, outside>, <load.resistor.n.i, inside>}
{<load.n.v, outside>, <load.resistor.n.v, inside>}
{<load.p.i, inside>, <ground.p.i, inside>, <resistor.p.i, inside>}
{<load.p.v, inside>, <ground.p.v, inside>, <resistor.p.v, inside>}
{<load.n.i, inside>}
{<resistor.n.i, inside>}

```

And thus the equations:

```

load.p.v = load.resistor.p.v;
load.n.v = load.resistor.n.v;
load.p.v = ground.p.v;
load.p.v = resistor.p.v;
0 = (-load.p.i) + load.resistor.p.i;
0 = (-load.n.i) + load.resistor.n.i;
0 = load.p.i + ground.p.i + resistor.p.i;
0 = load.n.i;
0 = resistor.n.i;

```

]

[Example: Outer component example:

```

model Circuit
  Ground ground;
  Load load;
  inner Resistor resistor;
equation
  connect(load.p, ground.p);
end Circuit;

```

```

model Load
  extends TwoPin;
  outer Resistor resistor;
equation
  connect(p, resistor.p);
  connect(resistor.n, n);
end Load;

```

The connection sets are before merging (note that one part of the load and resistor is not connected):

```

{<load.p.i, inside>}
{<load.n.i, inside>}
{<ground.p.i, inside>}
{<resistor.p.i, inside>}
{<resistor.n.i, inside>}
{<load.p.i, inside>, <ground.p.i, inside>}
{<load.p.v, inside>, <ground.p.v, inside>}
{<load.p.i, outside>, <resistor.p.i, inside>}
{<load.p.v, outside>, <resistor.p.v, inside>}
{<load.n.i, outside>, <resistor.n.i, inside>}
{<load.n.v, outside>, <resistor.n.v, inside>}

```

After merging this gives:

```

{<load.p.i, outside>, <resistor.p.i, inside>}
{<load.p.v, outside>, <resistor.p.v, inside>}
{<load.n.i, outside>, <resistor.n.i, inside>}
{<load.n.v, outside>, <resistor.n.v, inside>}
{<load.p.i, inside>, <ground.p.i, inside>}
{<load.p.v, inside>, <ground.p.v, inside>}
{<load.n.i, inside>}

```

And thus the equations:

```

load.p.v = resistor.p.v;
load.n.v = resistor.n.v;
load.p.v = ground.p.v;
0 = (-load.p.i) + resistor.p.i;
0 = (-load.n.i) + resistor.n.i;
0 = load.p.i + ground.p.i;
0 = load.n.i;

```

This corresponds to a direct connection of the resistor.]

9.3 Restrictions of Connections and Connectors

- The connect-equations (and the special functions for overdetermined connectors) can only be used in equations, and shall not be used inside if-equations with conditions that are not parameter expressions, or in when-equations.

[For-equations always have parameter expressions for the array expression.]

- A connector component shall not be declared with the prefix **parameter** or **constant**. In the connect-equation the primitive components may only connect parameter variables to parameter variables and constant variables to constant variables.
- The connect-equation construct only accepts forms of connector references as specified in section 9.1.
- In a connect-equation the two connectors must have the same named component elements with the same dimensions; recursively down to the primitive components. The primitive components with the same name are matched and belong to the same connection set.

- The matched primitive components of the two connectors must have the same primitive types, and flow variables may only connect to other flow variables, stream variables only to other stream variables, and causal variables (**input/output**) only to causal variables (**input/output**).
- A connection set of causal variables (**input/output**) may at most contain variables from one inside **output** connector (for state-machines extended as specified in section 17.3.6) or one public outside **input** connector.

[I.e., a connection set may at most contain one source of a signal.]

- At least one of the following must hold for a connection set containing causal variables generated for a non-partial model or block:
 1. the connection set includes variables from an outside public expandable connector,
 2. the set contains variables from protected outside connectors,
 3. it contains variables from one inside **output** connector, or
 4. from one public outside **input** connector, or
 5. the set is comprised solely of one variable from one inside **input** connector that is not part of an expandable connector.

[I.e., a connection set must – unless the model or block is partial – contain one source of a signal (item 5 covers the case where a connector of a component is left unconnected and the source given textually).]

- Variables from a protected outside connector must be part of a connection set containing at least one inside connector or one declared public outside connector (i.e. it shall not be an implicitly defined part of an expandable connector).

[Otherwise it would not be possible to deduce the causality for the expandable connector element.]

- In a connection set all variables having non-empty quantity attribute must have the same quantity attribute.
- A **connect** equation shall not (directly or indirectly) connect two connectors of **outer** elements.

*[Indirectly is similar to them being part of the same connection set. However, connections to **outer** elements are “moved up” before forming connection sets. Otherwise the connection sets could contain redundant information breaking the equation count for locally balanced models and blocks.]*

- Subscripts in a connector reference shall be parameter expressions or the special operator **::**.
- Constants or parameters in connected components yield the appropriate assert statements to check that they have the same value; connections are not generated.
- For conditional connectors, see section 4.4.5.

9.3.1 Balancing Restriction and Size of Connectors

For each non-partial non-expandable connector class the number of flow variables shall be equal to the number of variables that are neither **parameter**, **constant**, **input**, **output**, **stream** nor **flow**. The *number of variables* is the number of all elements in the connector class after expanding all records and arrays to a set of scalars of primitive types. The number of variables of an overdetermined type or record class (see section 9.4.1) is the size of the output argument of the corresponding `equalityConstraint()` function.

[Expandable connector classes are excluded from this, since their component declarations are only a form of constraint.]

[Example:

```
connector Pin // a physical connector of
  Modelica.Electrical.Analog
  Real v;
  flow Real i;
end Pin;
```

```

connector Plug // a hierarchical connector of
  Modelica.Electrical.MultiPhase
  parameter Integer m=3;
  Pin p[m];
end Plug;

connector InputReal = input Real; // A causal input connector
connector OutputReal = output Real; // A causal output connector

connector Frame_Illegal
  Modelica.Units.SI.Position r0[3] "Position vector of frame origin";
  Real S[3, 3] "Rotation matrix of frame";
  Modelica.Units.SI.Velocity v[3] "Abs. velocity of frame origin";
  Modelica.Units.SI.AngularVelocity w[3] "Abs. angular velocity of frame";
  Modelica.Units.SI.Acceleration a[3] "Abs. acc. of frame origin";
  Modelica.Units.SI.AngularAcceleration z[3] "Abs. angular acc. of frame";
  flow Modelica.Units.SI.Force f[3] "Cut force";
  flow Modelica.Units.SI.Torque t[3] "Cut torque";
end Frame_Illegal;

```

The `Frame_Illegal` connector (intended to be used in a simple multi-body package without over-determined connectors) is illegal since the number of flow and non-flow variables do not match. The solution is to create two connector classes, where two 3-vectors (e.g., `a` and `z`) are acausal `Real` and the other variables are matching pairs of `input` and `output`. This ensures that the models can only be connected in a tree-structure or require a “loop-breaker” joint for every closed kinematic loop:

```

connector Frame_a "correct connector"
  input Modelica.Units.SI.Position r0[3];
  input Real S[3, 3];
  input Modelica.Units.SI.Velocity v[3];
  input Modelica.Units.SI.AngularVelocity w[3];
  Modelica.Units.SI.Acceleration a[3];
  Modelica.Units.SI.AngularAcceleration z[3];
  flow Modelica.Units.SI.Force f[3];
  flow Modelica.Units.SI.Torque t[3];
end Frame_a;

connector Frame_b "correct connector"
  output Modelica.Units.SI.Position r0[3];
  output Real S[3, 3];
  output Modelica.Units.SI.Velocity v[3];
  output Modelica.Units.SI.AngularVelocity w[3];
  Modelica.Units.SI.Acceleration a[3];
  Modelica.Units.SI.AngularAcceleration z[3];
  flow Modelica.Units.SI.Force f[3];
  flow Modelica.Units.SI.Torque t[3];
end Frame_b;

```

The subsequent connectors `Plug_Expanded` and `PlugExpanded2` are correct, but `Plug_Expanded_Illegal` is illegal since the number of non-flow and flow variables is different if `n` and `m` are different. It is not clear how a tool can detect in general that connectors such as `Plug_Expanded_Illegal` are illegal. However, it is always possible to detect this defect after actual values of parameters and constants are provided in the simulation model.

```

connector Plug_Expanded "correct connector"
  parameter Integer m=3;
  Real v[m];
  flow Real i[m];
end Plug_Expanded;

connector Plug_Expanded2 "correct connector"
  parameter Integer m=3;
  final parameter Integer n=m;

```

```

Real v[m];
flow Real i[n];
end Plug_Expanded2;

connector Plug_Expanded_Illegal "connector is illegal"
parameter Integer m=3;
parameter Integer n=m;
Real v[m];
flow Real i[n];
end Plug_Expanded_Illegal;

```

9.4 Equation Operators for Overconstrained Connection-Based Equation Systems

There is a special problem regarding equation systems resulting from *loops* in connection graphs where the connectors contain *non-flow* (i.e., potential) variables *dependent* on each other. When a loop structure occurs in such a graph, the resulting equation system will be *overconstrained*, i.e., have more equations than variables, since there are implicit constraints between certain non-flow variables in the connector in addition to the connection equations around the loop. At the current state-of-the-art, it is not possible to automatically eliminate the unneeded equations from the resulting equation system without additional information from the model designer.

This section describes a set of equation operators for such overconstrained connection-based equation systems, that makes it possible for the model designer to specify enough information in the model to allow a Modelica environment to automatically remove the superfluous equations.

[Connectors may contain redundant variables. For example, the orientation between two coordinate systems in 3 dimensions can be described by 3 independent variables. However, every description of orientation with 3 variables has at least one singularity in the region where the variables are defined. It is therefore not possible to declare only 3 variables in a connector. Instead n variables ($n > 3$) have to be used. These variables are no longer independent from each other and there are $n - 3$ constraint equations that have to be fulfilled. A proper description of a redundant set of variables with constraint equations does no longer have a singularity. A model that has loops in the connection structure formed by components and connectors with redundant variables, may lead to a differential algebraic equation system that has more equations than unknown variables. The superfluous equations are usually consistent with the rest of the equations, i.e., a unique mathematical solution exists. Such models cannot be treated with the currently known symbolic transformation methods. To overcome this situation, operators are defined in order that a Modelica translator can remove the superfluous equations. This is performed by replacing the equality equations of non-flow variables from connection sets by a reduced number of equations in certain situations.]

This section handles a certain class of overdetermined systems due to connectors that have a redundant set of variables. There are other causes of overdetermined systems, e.g., explicit zero-sum equations for flow variables, that are not handled by the method described below.]

9.4.1 Overconstrained Equation Operators for Connection Graphs

A type or record declaration may have an optional definition of function `equalityConstraint` that shall have the following prototype:

```

type Type // overdetermined type
extends <base type>;
function equalityConstraint // non-redundant equality
input Type T1;
input Type T2;
output Real residue[ <n> ];
algorithm
  residue := ...
end equalityConstraint;

```



```

end Type;

record Record
  < declaration of record fields >
  function equalityConstraint // non-redundant equality
    input Record R1;
    input Record R2;
    output Real residue[ <n> ];
  algorithm
    residue := ...
  end equalityConstraint;
end Record;

```

The `residue` output of the `equalityConstraint` function shall have known size, say constant n . The function shall express the equality between the two type instances `T1` and `T2` or the record instances `R1` and `R2`, respectively, with a non-redundant number $n \geq 0$ of equations. The residues of these equations are returned in vector `residue` of size n . The set of n non-redundant equations stating that `R1 = R2` is given by the equation (0 represents a vector of zeros of appropriate size):

```

Record R1, R2;
equation
  0 = Record.equalityConstraint(R1, R2);

```

[If the elements of a record `Record` are not independent from each other, the equation `R1 = R2` contains redundant equations.]

A type class with an `equalityConstraint` function declaration is called *overdetermined type*. A record class with an `equalityConstraint` function definition is called *overdetermined record*. A connector that contains instances of overdetermined type and/or record classes is called *overdetermined connector*. An overdetermined type or record may neither have flow components nor may be used as a type of flow components. If an array is used as argument to any of the `Connections.*` functions it is treated as one unit – unlike `connect`, there is no special treatment of this case, compare section 9.1.

Every instance of an overdetermined type or record in an overdetermined connector is a node in a virtual connection graph that is used to determine when the standard equation `R1 = R2` or when the equation `0 = equalityConstraint(R1, R2)` has to be used for the generation of `connect` equations. The edges of the virtual connection graph are implicitly defined by `connect` and explicitly by `Connections.branch` statements, see table below. `Connections` is a built-in package in global scope containing built-in operators. Additionally, corresponding nodes of the virtual connection graph have to be defined as roots or as potential roots with functions `Connections.root` and `Connections.potentialRoot`, respectively.

The overconstrained equation operators for connection graphs are listed below. Here, `A` and `B` are connector instances that may be hierarchically structured, e.g., `A` may be an abbreviation for `EnginePort.Frame`.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>connect(A, B)</code>	Optional spanning-tree edge	Operator 9.1
<code>Connections.branch(A.R, B.R)</code>	Required spanning-tree edge	Operator 9.2
<code>Connections.root(A.R)</code>	Definite root node	Operator 9.3
<code>Connections.potentialRoot(A.R, ...)</code>	Potential root node	Operator 9.4
<code>Connections.isRoot(A.R)</code>	Predicate for being selected as root	Operator 9.5
<code>Connections.rooted(A.R)</code>	Predicate for being closer to root	Operator 9.6

Operator 9.1 `connect`

```
connect(A, B)
```

Defines *optional spanning-tree edge* from the overdetermined type or record instances in connector instance `A` to the corresponding overdetermined type or record instances in connector instance `B` for a virtual connection graph. The types of the corresponding overdetermined type or record instances shall be the same.

Operator 9.2 `Connections.branch`

```
Connections.branch(A.R, B.R)
```

Defines a *required spanning-tree edge* from the overdetermined type or record instance R in connector instance A to the corresponding overdetermined type or record instance R in connector instance B for a virtual connection graph. This function can be used at all places where a `connect` statement is allowed.

[E.g., it is not allowed to use this function in a when-clause. This definition shall be used if in a model with connectors A and B the overdetermined records A.R and B.R are algebraically coupled in the model, e.g., due to $B.R = f(A.R, \langle \text{other unknowns} \rangle)$.]

Operator 9.3 Connections.root

```
Connections.root(A.R)
```

The overdetermined type or record instance R in connector instance A is a (definite) *root node* in a virtual connection graph.

[This definition shall be used if in a model with connector A the overdetermined record A.R is (consistently) assigned, e.g., from a parameter expressions.]

Operator 9.4 Connections.potentialRoot

```
Connections.potentialRoot(A.R)
Connections.potentialRoot(A.R, priority=p)
```

The overdetermined type or record instance R in connector instance A is a *potential root node* in a virtual connection graph with priority p ($p \geq 0$). If no second argument is provided, the priority is zero. p shall be a parameter expression of type `Integer`. In a virtual connection subgraph without a `Connections.root` definition, one of the potential roots with the lowest priority number is selected as root.

[This definition may be used if in a model with connector A the overdetermined record A.R appears differentiated – `der(A.R)` – together with the constraint equations of A.R, i.e., a non-redundant subset of A.R maybe used as states.]

Operator 9.5 Connections.isRoot

```
Connections.isRoot(A.R)
```

Returns true, if the overdetermined type or record instance R in connector instance A is selected as a root in the virtual connection graph.

Operator 9.6 Connections.rooted

```
Connections.rooted(A.R)
rooted(A.R) // deprecated!
```

If the operator `Connections.rooted(A.R)` is used, or the equivalent but deprecated operator `rooted(A.R)`, then there must be exactly one statement `Connections.branch(A.R, B.R)` involving A.R (the argument of `Connections.rooted` must be the first argument of `Connections.branch`). In that case `Connections.rooted(A.R)` returns true, if A.R is closer to the root of the spanning tree than B.R; otherwise false is returned.

[This operator can be used to avoid equation systems by providing analytic inverses, see `Modelica.Mechanics.MultiBody.Parts.FixedRotation`.]

[Note, that `Connections.branch`, `Connections.root`, `Connections.potentialRoot` do not generate equations. They only generate nodes and edges in the virtual graph for analysis purposes.]

9.4.2 Converting the Connection Graph into Trees and Generating Connection Equations

Before `connect` equations are generated, the virtual connection graph is transformed into a set of spanning trees by removing optional spanning tree edges from the graph. This is performed in the following way:

1. Every root node defined via the `Connections.root` statement is a definite root of one spanning tree.
2. The virtual connection graph may consist of sets of subgraphs that are not connected together. Every subgraph in this set shall have at least one root node or one potential root node in a simulation model. If a graph of this set does not contain any root node, then one potential root node in this subgraph that has the lowest priority number is selected to be the root of that subgraph. The selection can be inquired in a class with function `Connections.isRoot`, see table above.
3. If there are n selected roots in a subgraph, then optional spanning-tree edges have to be removed such that the result shall be a set of n spanning trees with the selected root nodes as roots.

After this analysis, the connection equations are generated in the following way:

1. For every optional spanning-tree edge (i.e., a `connect(A, B)` equation), in one of the spanning trees, the connection equations are generated according to section 9.2.
2. For every optional spanning-tree edge not in any of the spanning trees, the connection equations are generated according to section 9.2, except for overdetermined type or record instances R . Here the equations `0 = R.equalityConstraint(A.R, B.R)` are generated instead of `A.R = B.R`.

9.4.3 Examples of Overconstrained Connection Graphs

[Example:




Figure 9.2: Example of a virtual connection graph.

]

9.4.3.1 An Overdetermined Connector for Power Systems

[An overdetermined connector for power systems based on the transformation theory of Park may be defined as:

```

type AC_Angle "Angle of source, e.g., rotor of generator"
  extends Modelica.Units.SI.Angle; // AC_Angle is a Real number
  // with unit = "rad"
function equalityConstraint
  input AC_Angle theta1;
  input AC_Angle theta2;
  output Real residue[0] "No constraints";
algorithm
  /* make sure that theta1 and theta2 from joining edges are identical */
  assert(abs(theta1 - theta2) < 1.e-10, "Consistent angles");
end equalityConstraint;
end AC_Angle;

connector AC_Plug "3-phase alternating current connector"
  import Modelica.Units.SI;
    
```

```

AC_Angle theta;
SI.Voltage v[3] "Voltages resolved in AC_Angle frame";
flow SI.Current i[3] "Currents resolved in AC_Angle frame";
end AC_Plug;

```

The currents and voltages in the connector are defined relatively to the harmonic, high-frequency signal of a power source that is essentially described by angle θ of the rotor of the source. This allows much faster simulations, since the basic high frequency signal of the power source is not part of the differential equations. For example, when the source and the rest of the line operates with constant frequency (= nominal case), then `AC_Plug.v` and `AC_Plug.i` are constant. In this case a variable step integrator can select large time steps. An element, such as a 3-phase inductor, may be implemented as:

```

model AC_Inductor
  parameter Real X[3,3], Y[3,3]; // component constants
  AC_Plug p;
  AC_Plug n;
  Real omega;
equation
  Connections.branch(p.theta,n.theta); //edge in virtual graph
  // since n.theta = p.theta
  n.theta = p.theta; // pass angle theta between plugs
  omega = der(p.theta); // frequency of source
  zeros(3) = p.i + n.i;
  X*der(p.i) + omega*Y*p.i = p.v - n.v;
end AC_Inductor

```

At the place where the source frequency, i.e., essentially variable θ , is defined, a `Connections.root` must be present:

```

  AC_Plug p;
equation
  Connections.root(p.theta);
  der(p.theta) = 2*Modelica.Constants.pi*50 // 50 Hz;

```

The graph analysis performed with the virtual connection graph identifies the connectors, where the `AC_Angle` needs not to be passed between components, in order to avoid redundant equations.]

9.4.3.2 An Overdetermined Connector for 3-dimensional Mechanical Systems

[An overdetermined connector for 3-dimensional mechanical systems may be defined as:

```

type TransformationMatrix = Real[3,3];
type Orientation "Orientation from frame 1 to frame 2"
  extendsTransformationMatrix;
function equalityConstraint
  input Orientation R1 "Rotation from inertial frame to frame 1";
  input Orientation R2 "Rotation from inertial frame to frame 2";
  output Real residue[3];
  protected
  Orientation R_rel "Relative Rotation from frame 1 to frame 2";
algorithm
  R_rel := R2*transpose(R1);
  /*
  If frame_1 and frame_2 are identical, R_rel must be
  the unit matrix. If they are close together, R_rel can be
  linearized yielding:
    R_rel = [ 1, phi3, -phi2;
             -phi3, 1, phi1;
             phi2, -phi1, 1 ];
  where phi1, phi2, phi3 are the small rotation angles around
  axis x, y, z of frame 1 to rotate frame 1 into frame 2.
  The atan2 is used to handle large rotation angles, but does not
  modify the result for small angles.
  */

```

```

    residue := { Modelica.Math.atan2(R_rel[2, 3], R_rel[1, 1]),
                Modelica.Math.atan2(R_rel[3, 1], R_rel[2, 2]),
                Modelica.Math.atan2(R_rel[1, 2], R_rel[3, 3])};
    end equalityConstraint;
end Orientation;

connector Frame "3-dimensional mechanical connector"
    import Modelica.Units.SI;
    SI.Position r[3] "Vector from inertial frame to Frame";
    Orientation R "Orientation from inertial frame to Frame";
    flow SI.Force f[3] "Cut-force resolved in Frame";
    flow SI.Torque t[3] "Cut-torque resolved in Frame";
end Frame;

```

A fixed translation from a frame A to a frame B may be defined as:

```

model FixedTranslation
    parameter Modelica.Units.SI.Position r[3];
    Frame frame_a, frame_b;
equation
    Connections.branch(frame_a.R, frame_b.R);
    frame_b.r = frame_a.r + transpose(frame_a.R)*r;
    frame_b.R = frame_a.R;
    zeros(3) = frame_a.f + frame_b.f;
    zeros(3) = frame_a.t + frame_b.t + cross(r, frame_b.f);
end FixedTranslation;

```

Since the transformation matrix `frame_a.R` is algebraically coupled with `frame_b.R`, an edge in the virtual connection graph has to be defined. At the inertial system, the orientation is consistently initialized and therefore the orientation in the inertial system connector has to be defined as root:

```

model InertialSystem
    Frame frame_b;
equation
    Connections.root(frame_b.R);
    frame_b.r = zeros(3);
    frame_b.R = identity(3);
end InertialSystem;

```

]

Chapter 10

Arrays

An array can be regarded as a collection of values, all of the same type. Modelica arrays can be multidimensional and are “rectangular”, which in the case of matrices has the consequence that all rows in a matrix have equal length, and all columns have equal length.

Each array has a certain dimensionality, i.e., number of dimensions. The degenerate case of a scalar variable is not really an array, but can be regarded as an array with zero dimensions. Vectors have one dimension, matrices have two dimensions, etc.

So-called row vectors and column vectors do not exist in Modelica and cannot be distinguished since vectors have only one dimension. If distinguishing these is desired, row matrices and column matrices are available, being the corresponding two-dimensional entities. However, in practice this is seldom needed since the usual matrix arithmetic and linear algebra operations have been defined to give the expected behavior when operating on Modelica vectors and matrices.

Modelica is a strongly typed language, which also applies to array types. The number of dimensions of an array is fixed and cannot be changed at run-time. However, the sizes of array dimensions can be computed at run-time.

The fixed number of array dimensions permits strong type checking and efficient implementation. The non-fixed sizes of array dimensions on the other hand, allow fairly generic array manipulation code to be written as well as interfacing to standard numeric libraries implemented in other programming languages.

An array is allocated by declaring an array variable or calling an array constructor. Elements of an array can be indexed by `Integer`, `Boolean`, or `enumeration` values.

10.1 Array Declarations

The Modelica type system includes scalar number, vector, matrix (number of dimensions, `ndim=2`), and arrays of more than two dimensions.

[There is no distinction between a row and column vector.]

The following table shows the two possible forms of declarations and defines the terminology. `C` is a placeholder for any class, including the built-in type classes `Real`, `Integer`, `Boolean`, `String`, and enumeration types. The type of a dimension upper bound expression, e.g. `n`, `m`, `p`, ... in the table below, need to be a subtype of `Integer` or `EB` for a class `EB` that is an enumeration type or subtype of the `Boolean` type.

Colon (`:`) indicates that the dimension upper bound is unknown and is a subtype of `Integer`. The size of such a variable can be determined from its binding equation, or the size of any of its array attributes, see also section 12.4.5. The size cannot be determined from other equations or algorithm.

Upper and lower array dimension index bounds are described in section 10.1.1.

An array indexed by `Boolean` or enumeration type can only be used in the following ways:

- Subscripted using expressions of the appropriate type (i.e. `Boolean` or the enumerated type).

- Binding equations of the form $\mathbf{x1} = \mathbf{x2}$ as well as declaration assignments of the form $\mathbf{x1} := \mathbf{x2}$ are allowed for arrays independent of whether the index types of dimensions are subtypes of `Integer`, `Boolean`, or enumeration types.

Table 10.1: General forms of declaration of arrays. The notation `EB` stands for an enumeration type or `Boolean`. The general array can have one or more dimensions ($k \geq 1$).

Modelica form 1	Modelica form 2	# dims	Designation	Explanation
<code>C x;</code>	<code>C x;</code>	0	Scalar	Scalar
<code>C[n] x;</code>	<code>C x[n];</code>	1	Vector	n -vector
<code>C[EB] x;</code>	<code>C x[EB]</code>	1	Vector	Vector indexed by EB
<code>C[n, m] x;</code>	<code>C x[n, m];</code>	2	Matrix	$n \times m$ matrix
<code>C[n₁, n₂, ..., n_k] x;</code>	<code>C x[n₁, n₂, ..., n_k];</code>	k	Array	General array

[Example: The number of dimensions and the dimensions sizes are part of the type, and shall be checked for example at redeclarations. Declaration form 1 displays clearly the type of an array, whereas declaration form 2 is the traditional way of array declarations in languages such as Fortran, C, C++.

```
Real[:] v1, v2 // vectors v1 and v2 have unknown sizes. The actual sizes may be different.
```

It is possible to mix the two declaration forms although it might be confusing.

```
Real[3, 2] x[4, 5]; // x has type Real[4, 5, 3, 2];
```

The reason for this order is given by examples such as:

```
type R3 = Real[3];
R3 a;
R3 b[1] = {a};
Real[3] c[1] = b;
```

Using a type for `a` and `b` in this way is normal, and substituting a type by its definition allows `c`.

A vector `y` indexed by enumeration values

```
type TwoEnums = enumeration(one, two);
Real[TwoEnums] y;
```

]

Zero-valued dimensions are allowed, so: `C x[0];` declares an empty vector, and: `C x[0, 3];` an empty matrix. Some examples of array dimensions of size one are given in table 10.2.

Table 10.2: Special cases of declaration of arrays as 1-vectors, row-vectors, or column-vectors of arrays.

Modelica form 1	Modelica form 2	# dims	Designation	Explanation
<code>C[1] x;</code>	<code>C x[1];</code>	1	Vector	1-vector, representing a scalar
<code>C[1, 1] x;</code>	<code>C x[1, 1];</code>	2	Matrix	(1×1) -matrix, representing a scalar
<code>C[n, 1] x;</code>	<code>C x[n, 1];</code>	2	Matrix	$(n \times 1)$ -matrix, representing a column
<code>C[1, n] x;</code>	<code>C x[1, n];</code>	2	Matrix	$(1 \times n)$ -matrix, representing a row

The type of an array of array is the multidimensional array which is constructed by taking the first dimensions from the component declaration and subsequent dimensions from the maximally expanded component type. A type is maximally expanded, if it is either one of the built-in types (`Real`, `Integer`, `Boolean`, `String`, enumeration type) or it is not a type class. Before operator overloading is applied, a type class of a variable is maximally expanded.

[Example:

```

type Voltage = Real(unit = "V");
type Current = Real(unit = "A");
connector Pin
  Voltage v; // type class of v = Voltage, type of v = Real
  flow Current i; // type class of i = Current, type of i = Real
end Pin;
type MultiPin = Pin[5];
MultiPin[4] p; // type class of p is MultiPin, type of p is Pin[4, 5];
type Point = Real[3];
Point p1[10];
Real p2[10, 3];
    
```

The components `p1` and `p2` have identical types.

```

p2[5] = p1[2] + p2[4]; // equivalent to p2[5, :] = p1[2, :] + p2[4, :]
Real r[3] = p1[2]; // equivalent to r[3] = p1[2, :]
    
```

[Automatic assertions at simulation time:

Let `A` be a declared array and `i` be the declared maximum dimension size of the `di`-dimension, then an assert statement `assert(i >= 0, ...)` is generated provided this assertion cannot be checked at compile time. It is a quality of implementation issue to generate a good error message if the assertion fails.

Let `A` be a declared array and `i` be an index accessing an index of the `di`-dimension. Then for every such index-access an assert statement `assert(1 <= i and i <= size(A, di), ...)` is generated, provided this assertion cannot be checked at compile time.

For efficiency reasons, these implicit assert statements may be optionally suppressed.]

10.1.1 Array Dimension Lower and Upper Index Bounds

The lower and upper index bounds for a dimension of an array indexed by `Integer`, `Boolean`, or `enumeration` values are as follows:

- An array dimension indexed by `Integer` values has a lower bound of 1 and an upper bound being the size of the dimension.
- An array dimension indexed by `Boolean` values has the lower bound `false` and the upper bound `true`.
- An array dimension indexed by `enumeration` values of the type `E = enumeration(e1, e2, ..., en)` has the lower bound `E.e1` and the upper bound `E.en`.

10.2 Flexible Array Sizes

Regarding flexible array sizes and resizing of arrays in functions, see section 12.4.5.

10.3 Built-in Array Functions

Modelica provides a number of built-in functions that are applicable to arrays.

The `promote` function listed below is utilized to define other array operators and functions.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>promote(A, n)</code>	Append dimensions of size 1 (not Modelica)	Operator 10.1

Operator 10.1 promote

```
promote(A, n) /* Not available in Modelica. */
```

Fills dimensions of size 1 from the right to array `A` upto dimension `n`, where $n \geq \text{ndims}(A)$ is required.

Let $C = \text{promote}(A, n)$, with $n_A = \text{ndims}(A)$, then $\text{ndims}(C) = n$, $\text{size}(C, j) = \text{size}(A, j)$ for $1 \leq j \leq n_A$, $\text{size}(C, j) = 1$ for $n_A + 1 \leq j \leq n$, $C[i_1, \dots, i_{n_A}, 1, \dots, 1] = A[i_1, \dots, i_{n_A}]$

The argument n must be a constant that can be evaluated during translation, as it determines the number of dimensions of the returned array.

[An n that is not a constant that can be evaluated during translation for `promote` complicates matrix handling as it can change matrix-equations in subtle ways (e.g. changing inner products to matrix multiplication).]

[Some examples of using the functions defined in the following section 10.3.1 to section 10.3.5:

```
Real x[4, 1, 6];
size(x, 1) = 4;
size(x); // vector with elements 4, 1, 6
size(2 * x + x) = size(x);
Real [3] v1 = fill(1.0, 3);
Real [3, 1] m = matrix(v1);
Real [3] v2 = vector(m);
Boolean check[3, 4] = fill(true, 3, 4);
```

]

10.3.1 Array Dimension and Size Functions

The functions listed below operate on the array dimensions of the type of an expression:

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>ndims(A)</code>	Number of dimensions	Operator 10.2
<code>size(A, i)</code>	Size of single array dimension	Operator 10.3
<code>size(A)</code>	Sizes of all array dimensions	Operator 10.4

Operator 10.2 `ndims`

`ndims(A)`

Returns the number of dimensions k of expression A , with $k \geq 0$.

Operator 10.3 `size`

`size(A, i)`

Returns the size of dimension i of array expression A where $0 \leq i \leq \text{ndims}(A)$.

If A refers to a component of an expandable connector, then the component must be a declared component of the expandable connector, and it must not use colon (`:`) to specify the array size of dimension i .

Operator 10.4 `size`

`size(A)`

Returns a vector of length `ndims(A)` containing the dimension sizes of A .

If A refers to a component of an expandable connector, then the component must be a declared component of the expandable connector, and it must not use colon (`:`) to specify the size of any array dimension.

10.3.2 Dimensionality Conversion Functions

The conversion functions listed below convert scalars, vectors, and arrays to scalars, vectors, or matrices by adding or removing 1-sized dimensions.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>scalar(A)</code>	Extract only element	Operator 10.5
<code>vector(A)</code>	Vector of all elements	Operator 10.6
<code>matrix(A)</code>	Two-dimensional array	Operator 10.7

Operator 10.5 scalar

```
scalar(A)
```

Returns the single element of array A . `size(A, i) = 1` is required for $1 \leq i \leq \text{ndims}(A)$.

Operator 10.6 vector

```
vector(A)
```

Returns a 1-vector if A is a scalar, and otherwise returns a vector containing all the elements of the array, provided there is at most one dimension size > 1 .

Operator 10.7 matrix

```
matrix(A)
```

Returns `promote(A, 2)` if A is a scalar or vector, and otherwise returns the elements of the first two dimensions as a matrix. `size(A, i) = 1` is required for $2 < i \leq \text{ndims}(A)$.

10.3.3 Specialized Array Constructor Functions

An array constructor function constructs and returns an array computed from its arguments. Most of the constructor functions listed below construct an array by filling in values according to a certain pattern, in several cases just giving all array elements the same value. The general array constructor with syntax `array(...)` or `{...}` is described in section 10.4.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>identity(n)</code>	Identity matrix	Operator 10.8
<code>diagonal(v)</code>	Diagonal matrix	Operator 10.9
<code>zeros(n₁, n₂, n₃, ...)</code>	Array with all elements being 0	Operator 10.10
<code>ones(n₁, n₂, n₃, ...)</code>	Array with all elements being 1	Operator 10.11
<code>fill(s, n₁, n₂, n₃, ...)</code>	Array with all elements equal	Operator 10.12
<code>linspace(x₁, x₂, n)</code>	Vector with equally spaced elements	Operator 10.13

Operator 10.8 identity

```
identity(n)
```

Returns the $n \times n$ **Integer** identity matrix, with ones on the diagonal and zeros at the other places.

Operator 10.9 diagonal

```
diagonal(v)
```

Returns a square matrix with the elements of vector v on the diagonal and all other elements zero.

Operator 10.10 zeros

```
zeros(n1, n2, n3, ...)
```

Returns the $n_1 \times n_2 \times n_3 \times \dots$ **Integer** array with all elements equal to zero ($n_i \geq 0$). The function needs one or more arguments, that is, `zeros()` is not legal.

Operator 10.11 ones

```
ones(n1, n2, n3, ...)
```

Return the $n_1 \times n_2 \times n_3 \times \dots$ **Integer** array with all elements equal to one ($n_i \geq 0$). The function needs one or more arguments, that is, `ones()` is not legal.

Operator 10.12 fill

```
fill(s, n1, n2, n3, ...)
```

Returns the $n_1 \times n_2 \times n_3 \times \dots$ array with all elements equal to scalar or array expression s ($n_i \geq 0$). The returned array has the same type as s .

Recursive definition: $\text{fill}(s, n_1, n_2, n_3, \dots) = \text{fill}(\text{fill}(s, n_2, n_3, \dots), n_1); \text{fill}(s, n) = \{s, s, \dots, s\}$.

The function needs two or more arguments; that is, $\text{fill}(s)$ is not legal.

Operator 10.13 linspace

`linspace(x1, x2, n)`

Returns a **Real** vector with n equally spaced elements, such that $v = \text{linspace}(x_1, x_2, n)$ results in

$$v[i] = x_1 + (x_2 - x_1) \frac{i - 1}{n - 1} \quad \text{for } 1 \leq i \leq n$$

It is required that $n \geq 2$. The arguments x_1 and x_2 shall be numeric scalar expressions.

10.3.4 Reduction Functions and Operators

The reduction functions listed below “reduce” an array (or several scalars) to one value (normally a scalar, but the `sum` reduction function may give an array as result and also be applied to an operator record). Note that none of these operators (particularly `min` and `max`) generate events themselves (but arguments could generate events). The restriction on the type of the input in section 10.3.4.1 for reduction expressions also applies to the array elements/scalar inputs for the reduction operator with the same name.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>min(A)</code>	Least element or array	Operator 10.14
<code>min(x, y)</code>	Least of two scalars	Operator 10.15
<code>min(... for ...)</code>	Reduction to least value	Operator 10.16
<code>max(A)</code>	Greatest element or array	Operator 10.17
<code>max(x, y)</code>	Greatest of two scalars	Operator 10.18
<code>max(... for ...)</code>	Reduction to greatest value	Operator 10.19
<code>sum(A)</code>	Sum of scalar array elements	Operator 10.20
<code>sum(... for ...)</code>	Sum reduction	Operator 10.21
<code>product(A)</code>	Product of scalar array elements	Operator 10.22
<code>product(... for ...)</code>	Product reduction	Operator 10.23

Operator 10.14 min

`min(A)`

Returns the least element of array expression A ; as defined by $<$.

Operator 10.15 min

`min(x, y)`

Returns the least element of the scalars x and y ; as defined by $<$.

Operator 10.16 min

`min(e(i, ..., j) for i in u, ..., j in v)`

Also described in section 10.3.4.1. Returns the least value (as defined by $<$) of the scalar expression $e(i, \dots, j)$ evaluated for all combinations of i in u, \dots, j in v .

Operator 10.17 max

`max(A)`

Returns the greatest element of array expression A ; as defined by $>$.

Operator 10.18 max

`max(x, y)`

Returns the greatest element of the scalars x and y ; as defined by $>$.

Operator 10.19 max

$$\max(e(i, \dots, j) \text{ for } i \text{ in } u, \dots, j \text{ in } v)$$

Also described in section 10.3.4.1. Returns the greatest value (as defined by $>$) of the scalar expression $e(i, \dots, j)$ evaluated for all combinations of i in u, \dots, j in v .

Operator 10.20 sum

$$\text{sum}(A)$$

Returns the scalar sum of all the elements of array expression A . Equivalent to sum reduction (see below, including application to operator records) over all array indices: $\text{sum}(A[j, k, \dots] \text{ for } j, k, \dots)$

Operator 10.21 sum

$$\text{sum}(e(i, \dots, j) \text{ for } i \text{ in } u, \dots, j \text{ in } v)$$

Also described in section 10.3.4.1. Returns the sum of the expression $e(i, \dots, j)$ evaluated for all combinations of i in u, \dots, j in v .

The **sum** reduction function (both variants) may be applied to an operator record, provided that the operator record defines '0' and '+'. It is then assumed to form an additive group.

For **Integer** indexing this is

$$\begin{aligned} &e(u[1], \dots, v[1]) + e(u[2], \dots, v[1]) + \dots \\ &+ e(u[\text{end}], \dots, v[1]) + \dots \\ &+ e(u[\text{end}], \dots, v[\text{end}]) \end{aligned}$$

For non-**Integer** indexing this uses all valid indices instead of 1..end.

The type of $\text{sum}(e(i, \dots, j) \text{ for } i \text{ in } u, \dots, j \text{ in } v)$ is the same as the type of $e(i, \dots, j)$.

Operator 10.22 product

$$\text{product}(A)$$

Returns the scalar product of all the elements of array expression A . Equivalent to product reduction (see below) over all array indices: $\text{product}(A[j, k, \dots] \text{ for } j, k, \dots)$

Operator 10.23 product

$$\text{product}(e(i, \dots, j) \text{ for } i \text{ in } u, \dots, j \text{ in } v)$$

Also described in section 10.3.4.1. Returns the product of the expression $e(i, \dots, j)$ evaluated for all combinations of i in u, \dots, j in v .

For **Integer** indexing this is

$$\begin{aligned} &e(u[1], \dots, v[1]) * e(u[2], \dots, v[1]) * \dots \\ &* e(u[\text{end}], \dots, v[1]) * \dots \\ &* e(u[\text{end}], \dots, v[\text{end}]) \end{aligned}$$

For non-**Integer** indexing this uses all valid indices instead of 1..end.

The type of $\text{product}(e(i, \dots, j) \text{ for } i \text{ in } u, \dots, j \text{ in } v)$ is the same as the type of $e(i, \dots, j)$.

10.3.4.1 Reduction Expressions

An expression:

$$\text{function-name } "(" \text{ expression1 for iterators } ")"$$

is a reduction-expression. The expressions in the iterators of a reduction-expression shall be vector expressions. They are evaluated once for each reduction-expression, and are evaluated in the scope immediately enclosing the reduction-expression.

For an iterator:

IDENT in expression2

the loop-variable, IDENT, is in scope inside `expression1`. The loop-variable may hide other variables, as in for-clauses. The result depends on the `function-name`, and currently the only legal function-names are the built-in operators `array`, `sum`, `product`, `min`, and `max`. For `array`, see section 10.4. If `function-name` is `sum`, `product`, `min`, or `max` the result is of the same type as `expression1` and is constructed by evaluating `expression1` for each value of the loop-variable and computing the `sum`, `product`, `min`, or `max` of the computed elements. For deduction of ranges, see section 11.2.2.1; and for using types as ranges see section 11.2.2.2.

Table 10.3: Reduction expressions with iterators. (The least and greatest values of `Real` are available as `-Modelica.Constants.inf` and `Modelica.Constants.inf`, respectively.)

Reduction	Restriction on <code>expression1</code>	Result for empty <code>expression2</code>
<code>sum</code>	Integer or Real	<code>zeros(...)</code>
<code>product</code>	Scalar Integer or Real	1
<code>min</code>	Scalar enumeration, Boolean, Integer or Real	Greatest value of type
<code>max</code>	Scalar enumeration, Boolean, Integer or Real	Least value of type

[Example:

```
sum(i for i in 1:10) // Gives  $\sum_{i=1}^{10} i = 1+2+\dots+10=55$ 
// Read it as: compute the sum of i for i in the range 1 to 10.
sum(i^2 for i in {1,3,7,6}) // Gives  $\sum_{i \in \{1, 3, 7, 6\}} i^2 = 1+9+49+36=95$ 
{product(j for j in 1:i) for i in 0:4} // Gives {1,1,2,6,24}
max(i^2 for i in {3,7,6}) // Gives 49
```

]

10.3.5 Matrix and Vector Algebra Functions

Functions for matrix and vector algebra are listed below. The function `transpose` can be applied to any matrix. The functions `outerProduct`, `symmetric`, `cross` and `skew` require `Real` vector(s) or matrix as input(s) and return a `Real` vector or matrix.

Expression	Description	Details
<code>transpose(A)</code>	Matrix transpose	Operator 10.24
<code>outerProduct(x, y)</code>	Vector outer product	Function 10.1
<code>symmetric(A)</code>	Symmetric matrix, keeping upper part	Function 10.2
<code>cross(x, y)</code>	Cross product	Function 10.3
<code>skew(x)</code>	Skew symmetric matrix associated with vector	Function 10.4

Operator 10.24 `transpose`

```
transpose(A)
```

Permutates the first two dimensions of array `A`. It is an error if array `A` does not have at least 2 dimensions.

Function 10.1 `outerProduct`

```
outerProduct(x, y)
```

Returns the outer product of vectors `x` and `y`, that is: `matrix(x) * transpose(matrix(y))`

Function 10.2 `symmetric`

```
symmetric(A)
```

Returns a symmetric matrix which is identical to the square matrix `A` on and above the diagonal.

That is, if $B := \text{symmetric}(A)$, then B is given by:

$$B[i,j] = \begin{cases} A[i,j] & \text{if } i \leq j \\ A[j,i] & \text{if } i > j \end{cases}$$

Function 10.3 cross

```
cross(x, y)
```

Returns the cross product of the 3-vectors x and y :

```
vector([ x[2] * y[3] - x[3] * y[2] ;
        x[3] * y[1] - x[1] * y[3] ;
        x[1] * y[2] - x[2] * y[1] ])
```

Function 10.4 skew

```
skew(x)
```

Returns the 3×3 skew symmetric matrix associated with a 3-vector, i.e., $\text{cross}(x, y) = \text{skew}(x) * y$. Equivalently, $\text{skew}(x)$ is given by:

```
[ 0,    -x[3], x[2] ;
  x[3], 0,    -x[1] ;
 -x[2], x[1], 0     ]
```

10.4 Vector, Matrix and Array Constructors

The constructor function `array(A, B, C, ...)` constructs an array from its arguments according to the following rules:

- Size matching: All arguments must have the same sizes, i.e., `size(A) = size(B) = size(C) = ...`
- All arguments must be type compatible expressions (section 6.7) giving the type of the elements. The data type of the result array is the maximally expanded type of the arguments. `Real` and `Integer` subtypes can be mixed resulting in a `Real` result array where the `Integer` numbers have been transformed to `Real` numbers.
- Each application of this constructor function adds a one-sized dimension to the left in the result compared to the dimensions of the argument arrays, i.e., `ndims(array(A, B, C)) = ndims(A) + 1 = ndims(B) + 1, ...`
- `{A, B, C, ...}` is a shorthand notation for `array(A, B, C, ...)`.
- There must be at least one argument.

[The reason `array()` or `{}` is not defined is that at least one argument is needed to determine the type of the resulting array.]

[Example:

```
{1, 2, 3} is a 3-vector of type Integer.
{{11, 12, 13}, {21, 22, 23}} is a 2 x 3 matrix of type Integer
{{{1.0, 2.0, 3.0}}}
```

```
Real[3] v = array(1, 2, 3.0);
type Angle = Real(unit="rad");
parameter Angle alpha = 2.0; // type of alpha is Real.
// array(alpha, 2, 3.0) or {alpha, 2, 3.0} is a 3-vector of type Real.
Angle[3] a = {1.0, alpha, 4}; // type of a is Real[3].
```

]

10.4.1 Array Constructor with Iterators

An expression:

```
"{" expression for iterators }"
```

or

```
array "(" expression for iterators ")"
```

is an array constructor with iterators. The expressions inside the iterators of an array constructor shall be vector expressions. They are evaluated once for each array constructor, and are evaluated in the scope immediately enclosing the array constructor.

For an iterator:

```
IDENT in array_expression
```

the loop-variable, IDENT, is in scope inside expression in the array construction. The loop-variable may hide other variables, as in for-clauses. The loop-variable has the same type as the type of the elements of array_expression; and can be simple type as well as a record type. The loop-variable will have the same type for the entire loop – i.e., for an array_expression {1, 3.2} the iterator will have the type of the type-compatible expression (Real) for all iterations. For deduction of ranges, see section 11.2.2.1; and for using types as range see section 11.2.2.2.

10.4.1.1 Array Constructor with One Iterator

If only one iterator is used, the result is a vector constructed by evaluating expression for each value of the loop-variable and forming an array of the result.

[Example:

```
array(i for i in 1:10)
// Gives the vector 1:10={1,2,3,...,10}

{r for r in 1.0 : 1.5 : 5.5}
// Gives the vector 1.0:1.5:5.5={1.0, 2.5, 4.0, 5.5}

{i^2 for i in {1,3,7,6}}
// Gives the vector {1, 9, 49, 36}
```

]

10.4.1.2 Array Constructor with Several Iterators

The notation with several iterators is a shorthand notation for nested array constructors. The notation can be expanded into the usual form by replacing each ',' by '}' for' and prepending the array constructor with a '{'.

[Example:

```
Real toeplitz[:,:]={i-j for i in 1:n, j in 1:n};
Real toeplitz2[:,:]={i-j for i in 1:n} for j in 1:n};
```

]

10.4.2 Array Concatenation

The function cat(k , A, B, C, ...) concatenates arrays A, B, C, ... along dimension k according to the following rules:

- Arrays A, B, C, ... must have the same number of dimensions, i.e., $\text{ndims}(A) = \text{ndims}(B) = \dots$
- Arrays A, B, C, ... must be type compatible expressions (section 6.7) giving the type of the elements of the result. The maximally expanded types should be equivalent. Real and Integer subtypes can be mixed resulting in a Real result array where the Integer numbers have been transformed to Real numbers.

- k has to characterize an existing dimension, i.e., $1 \leq k \leq \text{ndims}(A) = \text{ndims}(B) = \text{ndims}(C)$; k shall be a parameter expression of `Integer` type.
- Size matching: Arrays A, B, C, \dots must have identical array sizes with the exception of the size of dimension k , i.e., $\text{size}(A, j) = \text{size}(B, j)$, for $1 \leq j \leq \text{ndims}(A)$ and $j \neq k$.

[Example:

```
Real [2,3] r1 = cat(1, {{1.0, 2.0, 3}}, {{4, 5, 6}});
Real [2,6] r2 = cat(2, r1, 2*r1);
```

]

Formally, the concatenation $R = \text{cat}(k, A, B, C, \dots)$ is defined as follows. Let $n = \text{ndims}(A) = \text{ndims}(B) = \text{ndims}(C) = \dots$. Then the size of R is given by

$$\begin{aligned} \text{size}(R, k) &= \text{size}(A, k) + \text{size}(B, k) + \text{size}(C, k) + \dots \\ \text{size}(R, j) &= \text{size}(A, j) = \text{size}(B, j) = \text{size}(C, j) = \dots \text{ for } 1 \leq j \leq n \text{ and } j \neq k \end{aligned}$$

and the array elements of R are given by

$$\begin{aligned} R[i_1, \dots, i_k, \dots, i_n] &= A[i_1, \dots, i_k, \dots, i_n] \\ &\text{for } 0 < i_k \leq \text{size}(A, k) \\ R[i_1, \dots, i_k, \dots, i_n] &= B[i_1, \dots, i_k - \text{size}(A, k), \dots, i_n] \\ &\text{for } \text{size}(A, k) < i_k \leq \text{size}(A, k) + \text{size}(B, k) \\ R[i_1, \dots, i_k, \dots, i_n] &= C[i_1, \dots, i_k - \text{size}(A, k) - \text{size}(B, k), \dots, i_n] \\ &\text{for } \text{size}(A, k) + \text{size}(B, k) < i_k \leq \text{size}(A, k) + \text{size}(B, k) + \text{size}(C, k) \\ &\dots \end{aligned}$$

where $1 \leq i_j \leq \text{size}(R, j)$ for $1 \leq j \leq n$.

10.4.2.1 Array Concatenation along First and Second Dimensions

For convenience, a special syntax is supported for the concatenation along the first and second dimensions:

- *Concatenation along first dimension:*
 $[A; B; C; \dots] = \text{cat}(1, \text{promote}(A, n), \text{promote}(B, n), \text{promote}(C, n), \dots)$ where $n = \max(2, \text{ndims}(A), \text{ndims}(B), \text{ndims}(C), \dots)$. If necessary, 1-sized dimensions are added to the right of A, B, C before the operation is carried out, in order that the operands have the same number of dimensions which will be at least two.
- *Concatenation along second dimension:*
 $[A, B, C, \dots] = \text{cat}(2, \text{promote}(A, n), \text{promote}(B, n), \text{promote}(C, n), \dots)$ where $n = \max(2, \text{ndims}(A), \text{ndims}(B), \text{ndims}(C), \dots)$. If necessary, 1-sized dimensions are added to the right of A, B, C before the operation is carried out, especially that each operand has at least two dimensions.
- The two forms can be mixed. $[\dots; \dots]$ has higher precedence than $[\dots; \dots]$, e.g., $[a, b; c, d]$ is parsed as $[[a, b]; [c, d]]$.
- $[A] = \text{promote}(A, \max(2, \text{ndims}(A)))$, i.e., $[A] = A$, if A has 2 or more dimensions, and it is a matrix with the elements of A , if A is a scalar or a vector.
- There must be at least one argument (i.e. $[]$ is not defined).

[Example:

```
Real s1, s2, v1[n1], v2[n2], M1[m1,n],
M2[m2,n], M3[n,m1], M4[n,m2], K1[m1,n,k],
K2[m2,n,k];
[v1;v2] is a (n1+n2) x 1 matrix
[M1;M2] is a (m1+m2) x n matrix
[M3,M4] is a n x (m1+m2) matrix
[K1;K2] is a (m1+m2) x n x k array
[s1;s2] is a 2 x 1 matrix
[s1,s1] is a 1 x 2 matrix
[s1] is a 1 x 1 matrix
[v1] is a n1 x 1 matrix
```



```

Real [3] v1 = array(1, 2, 3);
Real [3] v2 = {4, 5, 6};
Real [3,2] m1 = [v1, v2];
Real [3,2] m2 = [v1, [4;5;6]]; // m1 = m2
Real [2,3] m3 = [1, 2, 3; 4, 5, 6];
Real [1,3] m4 = [1, 2, 3];
Real [3,1] m5 = [1; 2; 3];

```

]

10.4.3 Vector Construction

Vectors can be constructed with the general array constructor, e.g.,

```
Real [3] v = {1, 2, 3};
```

The range vector operator or colon operator of simple-expression can be used instead of or in combination with this general constructor to construct `Real`, `Integer`, `Boolean` or enumeration type vectors. Semantics of the colon operator:

- $j : k$ is the `Integer` vector $\{j, j + 1, \dots, k\}$, if j and k are of type `Integer`.
- $j : k$ is the `Real` vector $\{j, j + 1.0, \dots, j + n\}$, with $n = \text{floor}(k - j)$, if j and/or k are of type `Real`.
- $j : k$ is a `Real`, `Integer`, `Boolean`, or enumeration type vector with zero elements, if $j > k$.
- $j : d : k$ is the `Integer` vector $\{j, j + d, \dots, j + nd\}$, with $n = \text{div}(k - j, d)$, if j , d , and k are of type `Integer`.
- $j : d : k$ is the `Real` vector $\{j, j + d, \dots, j + nd\}$, with $n = \text{floor}((k - j)/d)$, if j , d , or k are of type `Real`. In order to avoid rounding issues for the length it is recommended to use $\{j + d * i \text{ for } i \text{ in } 0 : n\}$ or `linspace(j, k, n + 1)` – if the number of elements are known.
- $j : d : k$ is a `Real` or `Integer` vector with zero elements, if $d > 0$ and $j > k$ or if $d < 0$ and $j < k$.
- `false : true` is the `Boolean` vector $\{\text{false}, \text{true}\}$.
- $j : j$ is $\{j\}$ if j is `Real`, `Integer`, `Boolean`, or enumeration type.
- `E.ei : E.ej` is the enumeration type vector $\{E.ei, \dots, E.ej\}$ where $E.ej > E.ei$, and ei and ej belong to some enumeration type $E = \text{enumeration}(\dots, ei, \dots, ej, \dots)$.

[Example:

```

Real v1[5] = 2.7 : 6.8;
Real v2[5] = {2.7, 3.7, 4.7, 5.7, 6.7}; // = same as v1
Boolean b1[2] = false:true;
Colors = enumeration (red,blue,green);
Colors ec[3] = Colors.red : Colors.green;

```

]

10.5 Array Indexing

The array indexing operator `name[...]` is used to access array elements for retrieval of their values or for updating these values. An indexing operation is subject to upper and lower array dimension index bounds (section 10.1.1). The indexing operator takes two or more operands, where the first operand is the array to be indexed and the rest of the operands are index expressions:

```
arrayname[indexexpr1, indexexpr2, ...]
```

A colon is used to denote all indices of one dimension. A vector expression can be used to pick out selected rows, columns and elements of vectors, matrices, and arrays. The number of dimensions of the

expression is reduced by the number of scalar index arguments. If the number of index arguments is smaller than the number of dimensions of the array, the trailing indices will use ‘:’.

It is also possible to use the array access operator to assign to element/elements of an array in algorithm sections. If the index is an array the assignments take place in the order given by the index array. For assignments to arrays and elements of arrays, the entire right-hand side and the index on the left-hand side are evaluated before any element is assigned a new value.

[An indexing operation is assumed to take constant time, i.e., largely independent of the size of the array.]

[Example:

```

a[:, j] is a vector of the j-th column of a,
a[j] is a vector of the j-th row of a: a[j, :]
a[j : k] is {[a[j], a[j+1], ... , a[k]}
a[:, j : k] is [a[:,j], a[:,j+1], ... , a[:,k]],
v[2:2:8] = v[ {2,4,6,8} ] .
v[{j,k}] := {2,3}; // Same as v[j] := 2; v[k] := 3;
v[{1,1}] := {2,3}; // Same as v[1] := 3;
    
```

If \mathbf{x} is a vector, $\mathbf{x}[1]$ is a scalar, but the slice $\mathbf{x}[1:5]$ is a vector (a vector-valued or colon index expression causes a vector to be returned).

Table 10.4: Examples of scalars vs. array slices created with the colon index. The examples make use of the array variables $\mathbf{x}[n,m]$, $\mathbf{v}[k]$, and $\mathbf{z}[i,j,p]$.

Expression	# dims	Description
$\mathbf{x}[1, 1]$	0	Scalar
$\mathbf{x}[:, 1]$	1	n -vector
$\mathbf{x}[1, :]$ or $\mathbf{x}[1]$	1	m -vector
$\mathbf{v}[1:p]$	1	p -vector
$\mathbf{x}[1:p, :]$	2	$p \times m$ matrix
$\mathbf{x}[1:1, :]$	2	$1 \times m$ “row” matrix
$\mathbf{x}[\{1, 3, 5\}, :]$	2	$3 \times m$ matrix
$\mathbf{x}[:, \mathbf{v}]$	2	$n \times k$ matrix
$\mathbf{z}[:, 3, :]$	2	$i \times p$ matrix
$\mathbf{x}[\text{scalar}([1]), :]$	1	m -vector
$\mathbf{x}[\text{vector}([1]), :]$	2	$1 \times m$ “row” matrix

10.5.1 Indexing with Boolean or Enumeration Values

Arrays can be indexed using values of enumeration types or the `Boolean` type, not only by `Integer`. The type of the index should correspond to the type used for declaring the dimension of the array.

[Example:

```

type ShirtSizes = enumeration(small, medium, large, xlarge);
Real[ShirtSizes] w;
Real[Boolean] b2;
algorithm
  w[ShirtSizes.large] := 2.28; // Assign a value to an element of w
  b2[true] := 10.0;
  b2[ShirtSizes.medium] := 4; // Error, b2 was declared with Boolean dimension
  w[1] := 3; // Error, w was declared with ShirtSizes dimension
    
```

]

10.5.2 Indexing with end

The expression `end` may only appear inside array subscripts, and if used in the i :th subscript of an array expression \mathbf{A} it is equivalent to `size(A, i)` provided indices to \mathbf{A} are a subtype of `Integer`. If used inside nested array subscripts it refers to the most closely nested array.

[Example:

```
A[end - 1, end] is A[size(A,1) - 1, size(A,2)]
A[v[end], end] is A[v[size(v,1)], size(A,2)] // First end is referring to end of v.
```

]

10.6 Scalar, Vector, Matrix, and Array Operator Functions

The mathematical operations defined on scalars, vectors, and matrices are the subject of linear algebra.

The term numeric or numeric class is used below for a subtype of the `Real` or `Integer` type classes. The standard type coercion defined in section 10.6.13 applies.

10.6.1 Equality and Assignment

Equality $a = b$ and assignment $a := b$ of scalars, vectors, matrices, and arrays is defined element-wise and require both objects to have the same number of dimensions and corresponding dimension sizes. See section 10.5 regarding assignments to array variables with vector of subscripts.

The operands need to be type equivalent. This is legal for the simple types and all types satisfying the requirements for a record, and is in the latter case applied to each component-element of the records.

Table 10.5: Equality and assignment of arrays and scalars. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of a	Size of b	Size of a = b	Operation
Scalar	Scalar	Scalar	$a = b$
n -vector	n -vector	n -vector	$a[j] = b[j]$
$n \times m$ matrix	$n \times m$ matrix	$n \times m$ matrix	$a[j, k] = b[j, k]$
$n \times m \times \dots$	$n \times m \times \dots$	$n \times m \times \dots$	$a[j, k, \dots] = b[j, k, \dots]$

10.6.2 Array Element-wise Addition, Subtraction, and String Concatenation

Addition $a + b$ and subtraction $a - b$ of numeric scalars, vectors, matrices, and arrays is defined element-wise and require $\text{size}(a) = \text{size}(b)$ and a numeric type for a and b . Unary plus and minus are defined element-wise. Addition $a + b$ of string scalars, vectors, matrices, and arrays is defined as element-wise string concatenation of corresponding elements from a and b , and require $\text{size}(a) = \text{size}(b)$.

Table 10.6: Array addition, subtraction, and string concatenation. In this table the symbolic operator \pm represents either $+$ or $-$. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of a	Size of b	Size of a \pm b	Operation $c := a \pm b$
Scalar	Scalar	Scalar	$c := a \pm b$
n -vector	n -vector	n -vector	$c[j] := a[j] \pm b[j]$
$n \times m$ matrix	$n \times m$ matrix	$n \times m$ matrix	$c[j, k] := a[j, k] \pm b[j, k]$
$n \times m \times \dots$	$n \times m \times \dots$	$n \times m \times \dots$	$c[j, k, \dots] := a[j, k, \dots] \pm b[j, k, \dots]$

Element-wise addition $a .+ b$ and subtraction $a .- b$ of numeric scalars, vectors, matrices or arrays a and b requires a numeric type class for a and b and either $\text{size}(a) = \text{size}(b)$ or scalar a or scalar b . Element-wise addition $a .+ b$ of string scalars, vectors, matrices, and arrays is defined as element-wise string concatenation of corresponding elements from a and b , and require either $\text{size}(a) = \text{size}(b)$ or scalar a or scalar b .

Table 10.7: Array element-wise addition, subtraction, and string concatenation. In this table the symbolic operator \pm represents either $+$ or $-$, and when preceded by a dot ($\dot{\pm}$), either $\dot{+}$ or $\dot{-}$. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of \mathbf{a}	Size of \mathbf{b}	Size of $\mathbf{a} \dot{\pm} \mathbf{b}$	Operation $\mathbf{c} := \mathbf{a} \dot{\pm} \mathbf{b}$
Scalar	Scalar	Scalar	$\mathbf{c} := \mathbf{a} \pm \mathbf{b}$
Scalar	$n \times m \times \dots$	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \mathbf{a} \pm \mathbf{b}[j, k, \dots]$
$n \times m \times \dots$	Scalar	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \mathbf{a}[j, k, \dots] \pm \mathbf{b}$
$n \times m \times \dots$	$n \times m \times \dots$	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \mathbf{a}[j, k, \dots] \pm \mathbf{b}[j, k, \dots]$

Table 10.8: Unary operators. In this table the symbolic operator \pm represents either unary $+$ or unary $-$. The element-wise ($\dot{+}$, $\dot{-}$) and normal ($+$, $-$) operators give the same results. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of \mathbf{a}	Size of $\pm \mathbf{a}$	Operation $\mathbf{c} := \pm \mathbf{a}$
Scalar	Scalar	$\mathbf{c} := \pm \mathbf{a}$
$n \times m \times \dots$	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \pm \mathbf{a}[j, k, \dots]$

10.6.3 Array Element-wise Multiplication

Scalar multiplication $\mathbf{s} * \mathbf{a}$ or $\mathbf{a} * \mathbf{s}$ with numeric scalar \mathbf{s} and numeric scalar, vector, matrix or array \mathbf{a} is defined element-wise:

Table 10.9: Scalar and scalar to array multiplication of numeric elements. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of \mathbf{s}	Size of \mathbf{a}	Size of $\mathbf{s} * \mathbf{a}$ and $\mathbf{a} * \mathbf{s}$	Operation $\mathbf{c} := \mathbf{s} * \mathbf{a}$ or $\mathbf{c} := \mathbf{a} * \mathbf{s}$
Scalar	Scalar	Scalar	$\mathbf{c} := \mathbf{s} * \mathbf{a}$
Scalar	n -vector	n -vector	$\mathbf{c}[j] := \mathbf{s} * \mathbf{a}[j]$
Scalar	$n \times m$ matrix	$n \times m$ matrix	$\mathbf{c}[j, k] := \mathbf{s} * \mathbf{a}[j, k]$
Scalar	$n \times m \times \dots$	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \mathbf{s} * \mathbf{a}[j, k, \dots]$

Element-wise multiplication $\mathbf{a} .* \mathbf{b}$ of numeric scalars, vectors, matrices or arrays \mathbf{a} and \mathbf{b} requires a numeric type class for \mathbf{a} and \mathbf{b} and either `size(a) = size(b)` or scalar \mathbf{a} or scalar \mathbf{b} .

Table 10.10: Array element-wise multiplication. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of \mathbf{a}	Size of \mathbf{b}	Size of $\mathbf{a} .* \mathbf{b}$	Operation $\mathbf{c} := \mathbf{a} .* \mathbf{b}$
Scalar	Scalar	Scalar	$\mathbf{c} := \mathbf{a} * \mathbf{b}$
Scalar	$n \times m \times \dots$	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \mathbf{a} * \mathbf{b}[j, k, \dots]$
$n \times m \times \dots$	Scalar	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \mathbf{a}[j, k, \dots] * \mathbf{b}$
$n \times m \times \dots$	$n \times m \times \dots$	$n \times m \times \dots$	$\mathbf{c}[j, k, \dots] := \mathbf{a}[j, k, \dots] * \mathbf{b}[j, k, \dots]$

10.6.4 Matrix and Vector Multiplication of Numeric Arrays

Multiplication $\mathbf{a} * \mathbf{b}$ of numeric vectors and matrices is defined only for the following combinations:

Table 10.11: Matrix and vector multiplication of arrays with numeric elements. The scalar Operation applies for all i in $1, \dots, l$ and j in $1, \dots, n$, and the summation over k goes from 1 to m .

Size of \mathbf{a}	Size of \mathbf{b}	Size of $\mathbf{a} * \mathbf{b}$	Operation $\mathbf{c} := \mathbf{a} * \mathbf{b}$
m -vector	m -vector	Scalar	$c := \sum_k a[k] * b[k]$
m -vector	$m \times n$ matrix	n -vector	$c[j] := \sum_k a[k] * b[k, j]$
$l \times m$ matrix	m -vector	l -vector	$c[i] := \sum_k a[i, k] * b[k]$
$l \times m$ matrix	$m \times n$ matrix	$l \times n$ matrix	$c[i, j] := \sum_k a[i, k] * b[k, j]$

[Example:

```
Real A[3, 3], x[3], b[3], v[3];
A * x = b;
x * A = b; // same as transpose([x])*A*b
[v] * transpose([v]) // outer product
v * A * v // scalar
transpose([v]) * A * v // vector with one element
```

]

10.6.5 Division of Scalars or Numeric Arrays by Numeric Scalars

Division \mathbf{a} / \mathbf{s} of numeric scalars, vectors, matrices, or arrays \mathbf{a} and numeric scalars \mathbf{s} is defined element-wise. The result is always of `Real` type. In order to get integer division with truncation, use the function `div`.

Table 10.12: Division of scalars and arrays by numeric elements. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of \mathbf{a}	Size of \mathbf{s}	Size of \mathbf{a} / \mathbf{s}	Operation $\mathbf{c} := \mathbf{a} / \mathbf{s}$
Scalar	Scalar	Scalar	$c := a / s$
n -vector	Scalar	n -vector	$c[k] := a[k] / s$
$n \times m$ matrix	Scalar	$n \times m$ matrix	$c[j, k] := a[j, k] / s$
$n \times m \times \dots$	Scalar	$n \times m \times \dots$	$c[j, k, \dots] := a[j, k, \dots] / s$

10.6.6 Array Element-wise Division

Element-wise division $\mathbf{a} ./ \mathbf{b}$ of numeric scalars, vectors, matrices or arrays \mathbf{a} and \mathbf{b} requires a numeric type class for \mathbf{a} and \mathbf{b} and either `size(a) = size(b)` or scalar \mathbf{a} or scalar \mathbf{b} . The result is always of `Real` type. In order to get integer division with truncation, use the function `div`.

Table 10.13: Element-wise division of arrays. The scalar Operation applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of \mathbf{a}	Size of \mathbf{b}	Size of $\mathbf{a} ./ \mathbf{b}$	Operation $\mathbf{c} := \mathbf{a} ./ \mathbf{b}$
Scalar	Scalar	Scalar	$c := a / b$
Scalar	$n \times m \times \dots$	$n \times m \times \dots$	$c[j, k, \dots] := a / b[j, k, \dots]$
$n \times m \times \dots$	Scalar	$n \times m \times \dots$	$c[j, k, \dots] := a[j, k, \dots] / b$
$n \times m \times \dots$	$n \times m \times \dots$	$n \times m \times \dots$	$c[j, k, \dots] := a[j, k, \dots] / b[j, k, \dots]$

[Example: Element-wise division by scalar (`./`) and division by scalar (`/`) are identical: $\mathbf{a} ./ \mathbf{s} = \mathbf{a} / \mathbf{s}$:

```
2./[1, 2; 3, 4] // error; same as 2.0 / [1, 2; 3, 4]
2 ./[1, 2; 3, 4] // fine; element-wise division
```

This is a consequence of the parsing rules, since ‘`2.`’ is a lexical unit. Using a space after the literal solves the problem.]

10.6.7 Exponentiation of Scalars of Numeric Elements

Exponentiation $a \wedge b$ is defined as `pow(double a, double b)` in the ANSI C library if both `a` and `b` are `Real` scalars. A `Real` scalar value is returned. If `a` or `b` are `Integer` scalars, they are automatically promoted to `Real`. Consequences of exceptional situations, such as ($a = 0.0$ and $b \leq 0.0$, $a < 0$ and `b` is not an integer) or overflow are undefined.

Element-wise exponentiation `a .^ b` of numeric scalars, vectors, matrices, or arrays `a` and `b` requires a numeric type class for `a` and `b` and either `size(a) = size(b)` or scalar `a` or scalar `b`.

Table 10.14: Element-wise exponentiation of arrays. The scalar *Operation* applies for all j in $1, \dots, n$ and k in $1, \dots, m$.

Size of <code>a</code>	Size of <code>b</code>	Size of <code>a .^ b</code>	Operation <code>c := a .^ b</code>
Scalar	Scalar	Scalar	<code>c := a ^ b</code>
Scalar	$n \times m \times \dots$	$n \times m \times \dots$	<code>c[j, k, ...] := a ^ b[j, k, ...]</code>
$n \times m \times \dots$	Scalar	$n \times m \times \dots$	<code>c[j, k, ...] := a[j, k, ...] ^ b</code>
$n \times m \times \dots$	$n \times m \times \dots$	$n \times m \times \dots$	<code>c[j, k, ...] := a[j, k, ...] ^ b[j, k, ...]</code>

[Example:

```
2.^[1, 2; 3, 4] // error; same as 2.0 ^ [1, 2; 3, 4]
2 .^[1, 2; 3, 4] // fine; element-wise exponentiation
```

This is a consequence of the parsing rules, i.e. since `2.` could be a lexical unit it seen as a lexical unit; using a space after literals solves the problem.]

10.6.8 Scalar Exponentiation of Square Matrices of Numeric Elements

Exponentiation `a ^ s` is defined if `a` is a square numeric matrix and `s` is a scalar as a subtype of `Integer` with $s \geq 0$. The exponentiation is done by repeated multiplication, e.g.:

```
a^3 = a * a * a;
a^0 = identity(size(a, 1));
assert(size(a, 1) == size(a, 2), "Matrix must be square");
a^1 = a;
```

[Non-Integer exponents are forbidden, because this would require computing the eigenvalues and eigenvectors of `a` and this is no longer an elementary operation.]

10.6.9 Slice Operation

The following holds for slice operations:

- If `a` is an array containing scalar components and `m` is a component of those components, the expression `a.m` is interpreted as a slice operation. It returns the array of components `{a[1].m, ...}`.
- If `m` is also an array component, the slice operation is valid only if `size(a[1].m) = size(a[2].m) = ...`
- The slicing operation can be combined with indexing, e.g. `a.m[1]`. It returns the array of components `{a[1].m[1], a[2].m[1], ...}`, and does not require that `size(a[1].m) = size(a[2].m)`. The number of subscripts on `m` must not be greater than the number of array dimension for `m` (the number can be smaller, in which case the missing trailing indices are assumed to be `:`), and is only valid if `size(a[1].m[...]) = size(a[2].m[...])`.

[Example: The size-restriction on the operand is only applicable if the indexing on the second operand uses vectors or colon as in the example:

```
constant Integer m=3;
Modelica.Blocks.Continuous.LowpassButterworth tf[m](n=2:(m+1));
Real y[m];
```

```

Real y2,y3;
equation
// Extract the x1 slice even though different x1's have different lengths
y = tf.x1[1] ; // Legal, = {tf[1].x1[1], tf[2].x1[1], ... tf[m].x1[1]};
y2 = sum(tf.x1[:]); // Illegal to extract all elements since they have
// different lengths. Does not satisfy:
// size(tf[1].x1[:]) = size(tf[2].x1[:]) = ... = size(tf[m].x1[:])
y3 = sum(tf.x1[1:2]); // Legal.
// Since x1 has at least 2 elements in all tf, and
// size(tf[1].x1[1:2]) = size(tf[2].x1[1:2]) = ... = size(tf[m].x1[1:2]) =
{2}

```

In this example the different `x1` vectors have different lengths, but it is still possible to perform some operations on them.]

10.6.10 Relational Operators

Relational operators `<`, `<=`, `>`, `>=`, `==`, `<>`, are only defined for scalar operands of simple types, not for arrays, see section 3.5

10.6.11 Boolean Operators

The operators `and` and `or` take expressions of `Boolean` type, which are either scalars or arrays of matching dimensions. The operator `not` takes an expression of `Boolean` type, which is either scalar or an array. The result is the element-wise logical operation. For short-circuit evaluation of `and` and `or`, see section 3.3.

10.6.12 Vectorized Calls of Functions

See section 12.4.6.

10.6.13 Standard Type Coercion

In all contexts that require an expression which is a subtype of `Real`, an expression which is a subtype of `Integer` can also be used; the `Integer` expression is automatically converted to `Real`.

This also applies to arrays of `Real`, and for fields of record expressions. There is no similar rule for sub-typing.

[Example:

```

record RealR
  Real x,y;
end RealR;
record IntegerR
  Integer x,y;
end IntegerR;
parameter Integer a = 1;
Real y(start=a); // Ok, a is automatically coerced to Real
RealR r1 = IntegerR(a, a); // Ok, record is automatically coerced
RealR r2 = RealR(a, a); // Ok, a is automatically coerced to Real

```

]

10.7 Empty Arrays

Arrays may have dimension sizes of 0. For example:

```

Real x[0]; // an empty vector
Real A[0, 3], B[5, 0], C[0, 0]; // empty matrices

```

Empty matrices can be constructed with the fill function. For example:

```
Real A[:, :] = fill(0.0, 0, 1); // a Real 0 x 1 matrix
Boolean B[:, :, :] = fill(false, 0, 1, 0); // a Boolean 0 x 1 x 0 matrix
```

It is not possible to access an element of an empty matrix, e.g. $v[j, k]$ cannot be evaluated if $v = []$ because the assertion fails that the index must be bigger than one.

Size-requirements of operations, such as $+$, $-$, must also be fulfilled if a dimension is zero. For example:

```
Real [3, 0] A, B;
Real [0, 0] C;
A + B // fine, result is an empty matrix
A + C // error, sizes do not agree
```

Multiplication of two empty matrices results in a zero matrix of corresponding numeric type if the result matrix has no zero dimension sizes, i.e.,

```
Real [0, m] * Real [m, n] = Real [0, n] (empty matrix)
Real [m, n] * Real [n, 0] = Real [m, 0] (empty matrix)
Real [m, 0] * Real [0, n] = fill(0.0, m, n) (non-empty matrix, with zero elements
).
```

[Example:

```
Real u[p], x[n], y[q], A[n, n], B[n, p], C[q, n], D[q, p];
der(x) = A * x + B * u
y = C * x + D * u
```

Assume $n = 0, p > 0, q > 0$: Results in $y = D * u$.]

Chapter 11

Statements and Algorithm Sections

Whereas equations are very well suited for physical modeling, there are situations where computations are more conveniently expressed as algorithms, i.e., sequences of statements. In this chapter we describe the algorithmic constructs that are available in Modelica.

Statements are imperative constructs allowed in algorithm sections.

11.1 Algorithm Sections

An algorithm section is comprised of the keyword `algorithm` followed by a sequence of statements. The formal syntax is as follows:

```
algorithm-section :  
[ initial ] algorithm { statement ";" | annotation ";" }
```

Equation equality = or any other kind of equation (see chapter 8) shall not be used in an algorithm section.

11.1.1 Initial Algorithm Sections

See section 8.6 for a description of both initial algorithm sections and initial equation sections.

11.1.2 Execution of an algorithm in a model

An algorithm section is conceptually a code fragment that remains together and the statements of an algorithm section are executed in the order of appearance. Whenever an algorithm section is invoked, all variables appearing on the left hand side of the assignment operator := are initialized (at least conceptually):

- A continuous-time variable is initialized with its start value (i.e. the value of the `start` attribute).
- A discrete-time variable `v` is initialized with `pre(v)`.
- If at least one element of an array appears on the left hand side of the assignment operator, then the complete array is initialized in this algorithm section.
- A parameter assigned in an initial algorithm, section 8.6, is initialized with its start-value (i.e. the value of the `start` attribute).

[Initialization is performed, in order that an algorithm section cannot introduce a “memory” (except in the case of discrete-time variables assigned in the algorithm), which could invalidate the assumptions of a numerical integration algorithm. Note, a Modelica tool may change the evaluation of an algorithm section, provided the result is identical to the case, as if the above conceptual processing is performed.]

An algorithm section is treated as an atomic vector-equation, which is sorted together with all other equations. For the sorting process (BLT), every algorithm section with N different left-hand side variables, is treated as an atomic N -dimensional vector-equation containing all variables appearing in the algorithm

section. This guarantees that all N equations end up in an algebraic loop and the statements of the algorithm section remain together.

Example:

```

model Test // wrong Modelica model (has 4 equations for 2 unknowns)
  Real x[2](start={-11, -22});
algorithm // conceptually: x = {1,-22}
  x[1] := 1;
algorithm // conceptually: x = {-11,2}
  x[2] := 2;
end Test;

```

The conceptual part indicate that if the variable is assigned unconditionally in the algorithm before it is used the initialization can be omitted. This is usually the case, except for algorithms with when-statements, and especially for initial algorithms.]

11.1.3 Execution of the algorithm in a function

See section 12.4.4 *Initialization and Binding Equations of Components in Functions*.

11.2 Statements

Statements are imperative constructs allowed in algorithm sections. A flattened statement is identical to the corresponding nonflattened statement.

Names in statements are found as follows:

- If the name occurs inside an expression: it is first found among the lexically enclosing reduction functions (see section 10.3.4) in order starting from the inner-most, and if not found it proceeds as if it were outside an expression:
- Names in a statement are first found among the lexically enclosing for-statements in order starting from the inner-most, and if not found:
- Names in a statement shall be found by looking up in the partially flattened enclosing class of the algorithm section.

The syntax of statements is as follows:

```

statement :
  ( component-reference ( ":" expression | function-call-args )
    | "(" output-expression-list ")" ":" component-reference
    function-call-args
    | break
    | return
    | if-statement
    | for-statement
    | while-statement
    | when-statement )
  comment

```

11.2.1 Simple Assignment Statements

The syntax of simple assignment statement is as follows:

```

component-reference ":" expression

```

The **expression** is evaluated. The resulting value is stored into the variable denoted by **component-reference**. Assignment to array variables with subscripts is described in section 10.5.

11.2.1.1 Assignments from Called Functions with Multiple Results

There is a special form of assignment statement that is used only when the right-hand side contains a call to a function with multiple results. The left-hand side contains a parenthesized, comma-separated list of variables receiving the results from the function call. A function with n results needs $m \leq n$ receiving variables on the left-hand side, and the variables are assigned from left to right.

```
(out1, out2, out3) := function_name(in1, in2, in3, in4);
```

It is possible to omit receiving variables from this list:

```
(out1, , out3) := function_name(in1, in2, in3, in4);
```

[Example: The function `f` called below has three results and two inputs:

```
(a, b, c) := f(1.0, 2.0);
(x[1], x[2], x[1]) := f(3, 4);
```

In the second example above `x[1]` is assigned twice: first with the first output, and then with the third output. For that case the following will give the same result:

```
(, x[2], x[1]) := f(3,4);
```

]

The syntax of an assignment statement with a call to a function with multiple results is as follows:

```
"(" output-expression-list ")" := component-reference function-call-args
```

[Also see section 8.3.1 regarding calling functions with multiple results within equations.]

11.2.1.2 Restrictions on assigned variables

Only components of the restricted classes `type`, `record`, `operator record`, and `connector` may appear as left-hand-side in algorithms. This applies both to simple assignment statements, and the parenthesized, comma-separated list of variables for functions with multiple results.

11.2.2 For-statement

The syntax of a for-statement is as follows:

```
for for-indices loop
  { statement ";" }
end for
```

For-statements may optionally use several iterators (`for-indices`), see section 11.2.2.3 for more information:

```
for-indices:
  for-index {" ," for-index}

for-index:
  IDENT [ in expression ]
```

The following is an example of a prefix of a for-statement:

```
for IDENT in expression loop
```

The rules for for-statements are the same as for for-expressions in section 8.3.2.1 – except that the `expression` of a for-statement is not restricted to a parameter expression.

[Example:

```
for i in 1 : 10 loop // i takes the values 1, 2, 3, ..., 10
for r in 1.0 : 1.5 : 5.5 loop // r takes the values 1.0, 2.5, 4.0, 5.5
for i in {1, 3, 6, 7} loop // i takes the values 1, 3, 6, 7
for i in TwoEnums loop // i takes the values TwoEnums.one, TwoEnums.two
  // for TwoEnums = enumeration(one, two)
```

The loop-variable may hide other variables as in the following example. Using another name for the loop-variable is, however, strongly recommended.

```

constant Integer j = 4;
Real x[j];
equation
  for j in 1:j loop // The loop-variable j takes the values 1, 2, 3, 4
    x[j] = j; // Uses the loop-variable j
  end for;

```

]

11.2.2.1 Implicit Iteration Ranges

An iterator IDENT **in** range-expr without the **in** range-expr requires that the IDENT appears as the subscript of one or several subscripted expressions, where the expressions are not part of an array in a component of an expandable connector. The dimension size of the array expression in the indexed position is used to deduce the range-expr as `1:size(array-expression, indexpos)` if the indices are a subtype of `Integer`, or as `E.e1:E.en` if the indices are of an enumeration type `E = enumeration(e1, ..., en)`, or as `false:true` if the indices are of type `Boolean`. If it is used to subscript several expressions, their ranges must be identical. The IDENT may also, inside a reduction-expression, array constructor expression, for-statement, or for-equation, occur freely outside of subscript positions, but only as a reference to the variable IDENT, and not for deducing ranges.

The IDENT may also be used as a subscript for an array in a component of an expandable connector but it is only seen as a reference to the variable IDENT and cannot be used for deducing ranges.

[Example:

```

Real x[4];
Real xsquared[:] = {x[i] * x[i] for i};
// Same as: {x[i] * x[i] for i in 1 : size(x, 1)}
Real xsquared2[size(x, 1)];
Real xsquared3[size(x, 1)];
equation
  for i loop // Same as: for i in 1 : size(x, 1) loop ...
    xsquared2[i] = x[i]^2;
  end for;
algorithm
  for i loop // Same as: for i in 1 : size(x, 1) loop ...
    xsquared3[i] := x[i]^2;
  end for;

```

```

type FourEnums = enumeration(one, two, three, four);
Real xe[FourEnums] = x;
Real xsquared3[FourEnums] = {xe[i] * xe[i] for i};
Real xsquared4[FourEnums] = {xe[i] * xe[i] for i in FourEnums};
Real xsquared5[FourEnums] = {x[i] * x[i] for i};

```

]

The size of an array – the iteration range – is evaluated on entry to the for-loop and the array size shall not change during the execution of the for-loop.

11.2.2.2 Types as Iteration Ranges

The iteration range can be specified as `Boolean` or as an enumeration type. This means iteration over the type from min to max, i.e. for `Boolean` it is the same as `false:true` and for an enumeration `E` it is the same as `E.min:E.max`. This can be used for **for** loops and reduction expressions.

[Example:

```

type FourEnums = enumeration(one, two, three, four);
Real xe[FourEnums];
Real xsquared1[FourEnums];

```

```

Real xsquared2[FourEnums] = {xe[i] * xe[i] for i in FourEnums};
equation
  for i in FourEnums loop
    xsquared1[i] = xe[i]^2;
  end for;

```

11.2.2.3 Nested For-Loops and Reduction Expressions with Multiple Iterators

The notation with several iterators is a shorthand notation for nested for-statements or for-equations (or reduction-expressions). For for-statements or for-equations it can be expanded into the usual form by replacing each ‘,’ by “loop for” and adding extra “end for”. For reduction-expressions it can be expanded into the usual form by replacing each ‘,’ by “) for” and prepending the reduction-expression with “functionName(”.

[Example:

```

Real x[4,3];
algorithm
  for j, i in 1:2 loop
    // The loop-variable j takes the values 1,2,3,4 (due to use)
    // The loop-variable i takes the values 1,2 (given range)
    x[j,i] := j+i;
  end for;

```

11.2.3 While-Statement

The while-statement has the following syntax:

```

while expression loop
  { statement ";" }
end while

```

The expression of a while-statement shall be a scalar Boolean expression. The while-statement corresponds to while-statements in programming languages, and is formally defined as follows:

1. The expression of the while-statement is evaluated.
2. If the expression of the while-statement is false, the execution continues after the while-statement.
3. If the expression of the while-statement is true, the entire body of the while-statement is executed (except if a break-statement, see section 11.2.4, or a return-statement, see section 11.2.5, is executed), and then execution proceeds at step 1.

11.2.4 Break-Statement

The break-statement breaks the execution of the innermost while- or for-loop enclosing the break-statement and continues execution after the while- or for-loop. It can only be used in a while- or for-loop in an algorithm section. It has the following syntax:

```

break;

```

[Example: (Note that this could alternatively use return).

```

function findValue "Returns position of val or 0 if not found"
  input Integer x[:];
  input Integer val;
  output Integer index;
algorithm
  index := size(x, 1);
  while index >= 1 loop
    if x[index] == val then

```

```

    break;
  else
    index := index - 1;
  end if;
end while;
end findValue;

```

11.2.5 Return-Statements

Can only be used inside functions, see section 12.1.2.

11.2.6 If-Statement

If-statements have the following syntax:

```

if expression then
  { statement ";" }
  { elseif expression then
  { statement ";" }
  }
  [ else
  { statement ";" }
  ]
end if;

```

The **expression** of an if- or elseif-clause must be scalar **Boolean** expression. One if-clause, and zero or more elseif-clauses, and an optional else-clause together form a list of branches. One or zero of the bodies of these if-, elseif- and else-clauses is selected, by evaluating the conditions of the if- and elseif-clauses sequentially until a condition that evaluates to true is found. If none of the conditions evaluate to true the body of the else-clause is selected (if an else-clause exists, otherwise no body is selected). In an algorithm section, the selected body is then executed. The bodies that are not selected have no effect on that model evaluation.

11.2.7 When-Statements

A when-statement has the following syntax:

```

when expression then
  { statement ";" }
  { elseif expression then
  { statement ";" } }
end when

```

The **expression** of a when-statement shall be a discrete-time **Boolean** scalar or vector expression. The algorithmic statements within a when-statement are activated when the scalar or any one of the elements of the vector-expression becomes true.

[Example: Algorithms are activated when x becomes > 2 :

```

when x > 2 then
  y1 := sin(x);
  y3 := 2*x + y1+y2;
end when;

```

The statements inside the when-statement are activated on the positive edge of any of the expressions $x > 2$, $\text{sample}(0, 2)$, or $x < 5$:

```

when {x > 2, sample(0,2), x < 5} then
  y1 := sin(x);
  y3 := 2*x + y1+y2;
end when;

```

For when-statements in algorithm sections the order is significant and it is advisable to have only one assignment within the when-statement and instead use several algorithm sections having when-statements with identical conditions, e.g.:

```

algorithm
  when x > 2 then
    y1 := sin(x);
  end when;
equation
  y2 = sin(y1);
algorithm
  when x > 2 then
    y3 := 2*x +y1+y2;
  end when;
    
```

Merging the when-statements can lead to less efficient code and different models with different behavior depending on the order of the assignment to y1 and y3 in the algorithm.]

11.2.7.1 Restrictions on Where a When-statement may occur

- A when-statement shall not be used within a function.
- When-statements shall not occur inside initial algorithms.
- When-statements cannot be nested.
- When-statements shall not occur inside while, if, and for-clauses in algorithms.

[Example: The following nested when-statement is invalid:

```

when x > 2 then
  when y1 > 3 then
    y2 := sin(x);
  end when;
end when;
    
```

]

11.2.7.2 Restrictions on Statements within When-Statements

[In contrast to when-equations, section 8.3.5.3, there are no additional restrictions within when-statements:

- *In algorithms, all assignment statements are already restricted to left-hand-side variables.*
- *If at least one element of an array appears on the left-hand-side of the assignment operator inside a when-statement, it is as if the entire array appears in the left-hand-side according to section 11.1.2. Thus, there is no need to restrict the indices to parameter expressions.*
- *For-clauses and if-clauses are not problematic inside when-statements in algorithms, since all left-hand-side variables inside when-statements are assigned to their pre-values before the start of the algorithm, according to section 11.1.2.*

]

11.2.7.3 Defining When-Statements by If-Statements

A when-statement:

```

algorithm
  when {x>1, ..., y>p} then
    ...
  elseif x > y.start then
    ...
  end when;
    
```

is similar to the following special if-statement, where `Boolean b1[N]`; and `Boolean b2`; are necessary because `edge` can only be applied to variables

```

Boolean b1[N](start={x.start>1, ...,
y.start>p});
Boolean b2(start=x.start>y.start);
algorithm
  b1:={x>1, ..., y>p};
  b2:=x>y.start;
  if edge(b1[1]) or edge(b1[2]) or ...
    edge(b1[N]) then
    ...
    elseif edge(b2) then
    ...
  end if;

```

with `edge(A) = A and not pre(A)` and the additional guarantee, that the statements within this special if-statement are only evaluated at event instants. The difference compared to the when-statements is that e.g. `pre` may only be used on continuous-time real variables inside the body of a when-clause and not inside these if-statements.

11.2.8 Special Statements

These special statements have the same form and semantics as the corresponding equations, apart from the general difference in semantics between equations and statements.

11.2.8.1 Assert Statement

See section 8.3.7. A failed `assert` stops the execution of the current algorithm.

11.2.8.2 Terminate Statement

See section 8.3.8. The `terminate` statement shall not be in functions. In an algorithm outside a function it does not stop the execution of the current algorithm.

Chapter 12

Functions

This chapter describes the Modelica function construct.

12.1 Function Declaration

A Modelica function is a specialized class (section 12.2) using the keyword `function`. The body of a Modelica function is an algorithm section that contains procedural algorithmic code to be executed when the function is called, or alternatively an external function specifier (section 12.9). Formal parameters are specified using the `input` keyword, whereas results are denoted using the `output` keyword. This makes the syntax of function definitions quite close to Modelica class definitions, but using the keyword `function` instead of `class`.

[The structure of a typical function declaration is sketched by the following schematic function example:

```
function  functionname
  input  TypeI1 in1;
  input  TypeI2 in2;
  input  TypeI3 in3 = default_expr1 "Comment" annotation(...);
  ...
  output TypeO1 out1;
  output TypeO2 out2 = default_expr2;
  ...
protected
  <local variables>
  ...
algorithm
  ...
  <statements>
  ...
end functionname;
```

]

Optional explicit default values can be associated with any input or output formal parameter through binding equations. Comment strings and annotations can be given for any formal parameter declaration, as usual in Modelica declarations.

[Explicit default values are shown for the third input parameter and the second output parameter in the example above.]

[All internal parts of a function are optional; i.e., the following is also a legal function:

```
function functionname
end functionname;
```

]

12.1.1 Ordering of Formal Parameters

The relative ordering between input formal parameter declarations is significant since that determines the matching between actual arguments and formal parameters at function calls with positional parameter passing. Likewise, the relative ordering between the declarations of the outputs is significant since that determines the matching with receiving variables at function calls of functions with multiple results. However, the declarations of the inputs and outputs can be intermixed as long as these internal orderings are preserved.

[*Mixing declarations in this way is not recommended, however, since it makes the code hard to read.*]

[*Example:*

```
function <functionname>
  output Type01 out1; // Intermixed declarations of inputs and outputs
  input TypeI1 in1; // not recommended since code becomes hard to read
  input TypeI2 in2;
  ...
  output Type02 out2;
  input TypeI3 in3;
  ...
end <functionname>;
```

]

12.1.2 Function return-statements

The return-statement terminates the current function call, see section 12.4. It can only be used in an algorithm section of a function. It has the following form:

```
return;
```

[*Example: (Note that this could alternatively use `break`.)*

```
function findValue "Returns position of val or 0 if not found"
  input Integer x[:];
  input Integer val;
  output Integer index;
algorithm
  for i in 1:size(x,1) loop
    if x[i] == val then
      index := i;
      return;
    end if;
  end for;
  index := 0;
  return;
end findValue;
```

]

12.1.3 Inheritance of Functions

It is allowed for a function to inherit and/or modify another function following the usual rules for inheritance of classes (chapter 7).

[*For example, it is possible to modify and extend a `function` class to add default values for input variables.*]

12.2 Function as a Specialized Class

The function concept in Modelica is a specialized class (section 4.6).

[The syntax and semantics of a function have many similarities to those of the `block` specialized class. A function has many of the properties of a general class, e.g. being able to inherit other functions, or to redeclare or modify elements of a function declaration.]

Modelica functions have the following restrictions compared to a general Modelica `class`:

- Each input formal parameter of the function must be prefixed by `input`, and each result formal parameter by `output`. All public variables are formal parameters.
- Input formal parameters are read-only after being bound to the actual arguments or default values, i.e., they shall not be assigned values in the body of the function.
- A function shall *not be used in connections*, shall not have *equations*, shall not have *initial algorithms*.
- A function can have at most *one algorithm* section or *one external function interface* (not both), which, if present, is the body of the function.
- A function may only contain components of the restricted classes `type`, `record`, `operator record`, and `function`; and it must not contain e.g. `model`, `block`, `operator` or `connector` components.
- The elements of a function shall not have prefixes `inner`, or `outer`.
- A function may have zero or one external function interface, which, if present, is the external definition of the function.
- For a function to be called in a simulation model, the function shall not be partial, and the output variables must be assigned inside the function either in binding equations or in an algorithm section, or have an external function interface as its body, or be defined as a function partial derivative. The output variables of a function should be computed.

[It is a quality of implementation how much analysis a tool performs in order to determine if the output variables are computed.]

A function *cannot contain* calls to the Modelica *built-in operators* `der`, `initial`, `terminal`, `sample`, `pre`, `edge`, `change`, `reinit`, `delay`, `cardinality`, `inStream`, `actualStream`, to the operators of the built-in package `Connections`, to the operators defined in chapter 16 and chapter 17, and is not allowed to contain when-statements.

- The dimension *sizes* not declared with `(:)` of each array result or array local variable (i.e., a non-input components) of a function must be either given by the input formal parameters, or given by constant or parameter expressions, or by expressions containing combinations of those (section 12.4.4).
- For initialization of local variables of a function see section 12.4.4).
- Components of a function will inside the function behave as though they had discrete-time variability.

Modelica functions have the following enhancements compared to a general Modelica `class`:

- Functions can be called, section 12.4.
 - The calls can use a mix of positional and named arguments, see section 12.4.1.
 - Instances of functions have a special meaning, see section 12.4.2.
 - The lookup of the `function` class to be called is extended, see section 5.3.2.
- A function can be *recursive*.
- A formal parameter or local variable may be initialized through a *binding* `(=)` of a default value in its declaration, see section 12.4.4. Using assignment `(:=)` is deprecated. If a non-input component in the function uses a record class that contain one or more binding equations they are viewed as initialization of those component of the record component.
- A function is dynamically instantiated when it is called rather than being statically instantiated by an instance declaration, which is the case for other kinds of classes.
- A function may have an external function interface specifier as its body.

- A function may have a `return` statement in its algorithm section body.
- A function allows dimension sizes declared with `(:)` to be resized for non-input array variables, see section 12.4.5.
- A function may be defined in a short function definition to be a function partial derivative.

12.3 Pure Modelica Functions

Modelica functions are normally *pure* which makes it easy for humans to reason about the code since they behave as mathematical functions, and possible for compilers to optimize.

- *Pure* Modelica functions always give the same output values or errors for the same input values and only the output values influence the simulation result, i.e. is seen as equivalent to a mathematical map from input values to output values. Some input values may map to errors. Pure functions are thus allowed to fail by calling `assert`, or `ModelicaError` in C code, or dividing by zero. Such errors will only be reported when and if the function is called. *Pure* Modelica functions are not assumed to be thread-safe.
- A Modelica function which does not have the *pure* function properties is *impure*.

The declaration of functions follow these rules:

- Functions defined in Modelica (non-external) are *normally* assumed to be pure (the exception is the deprecated case below), if they are impure they shall be marked with the `impure` keyword. They can be explicitly marked as `pure`.

[However, since functions as default are pure it is not recommended to explicitly declare them as pure.]

- External functions must be explicitly declared with `pure` or `impure`.
- A deprecated semantics is that external functions (and functions defined in Modelica directly or indirectly calling them) without `pure` or `impure` keyword are assumed to be impure, but without any restriction on calling them. Except for the function `Modelica.Utilities.Streams.print`, diagnostics must be given if called in a simulation model.

Calls of pure functions used inside expression may be skipped if the resulting expression will not depend on the possible returned value; ignoring the possibility of the function generating an error.

A call to a function with no declared outputs is assumed to have desired side-effects or assertion checks.

[A tool shall thus not remove such function calls, with exception of non-triggered assert calls. A pure function, used in an expression or used with a non-empty left hand side, need not be called if the output from the function call do not mathematically influence the simulation result, even if errors would be generated if it were called.]

[Comment 1: This property enables writing declarative specifications using Modelica. It also makes it possible for Modelica compilers to freely perform algebraic manipulation of expressions containing function calls while still preserving their semantics. For example, a tool may use common subexpression elimination to call a pure function just once, if it is called several times with identical input arguments. However, since functions may fail we can e.g. only move a common function call from inside a loop to outside the loop if the loop is run at least once.]

[Comment 2: The Modelica translator is responsible for maintaining this property for pure non-external functions. Regarding external functions, the external function implementor is responsible. Note that external functions can have side-effects as long as they do not influence the internal Modelica simulation state, e.g. caching variables for performance or printing trace output to a log file.]

With the prefix keyword `impure` it is stated that a Modelica function is *impure* and it is only allowed to call such a function from within:

- Another function marked with the prefix `impure`.
- A when-equation.

- A when-statement.
- `pure(impureFunctionCall(...))` – which allows calling impure functions in any pure context.
- Initial equations and initial algorithms.
- Binding equations for components declared as parameter – which is seen as syntactic sugar for having a parameter with `fixed=false` and the binding as an initial equation.

[Thus, evaluation of the same function call at a later time during simulation is not guaranteed to result in the same value as when the parameter was initialized, seemingly breaking the declaration equation.]

- Binding equations for external objects.

For initial equations, initial algorithms, and bindings it is an error if the function calls are part of systems of equations and thus have to be called multiple times.

[A tool is not allowed to perform any optimizations on function calls to an impure function, e.g., re-ordering calls from different statements in an algorithm or common subexpression elimination is not allowed.]

It is possible to mark a function formal parameter as `impure`. Only if the function formal parameter is marked `impure`, it is allowed to pass an `impure` function to it. A function having a formal function parameter marked `impure` must be marked `pure` or `impure`.

[Comment: The semantics are undefined if the function call of an impure function is part of an algebraic loop.]

[Example:

```
function evaluateLinear // pure function
  input Real a0;
  input Real a1;
  input Real x;
  output Real y;
algorithm
  y := a0 + a1*x;
end evaluateLinear;

impure function receiveRealSignal // impure function
  input HardwareDriverID id;
  output Real y;
external "C"
  y = receiveSignal(id);
end receiveRealSignal;
```

Examples of allowed optimizations of pure functions:

```
model M // Assume sin, cos, asin are pure functions with normal derivatives.
  input Real x[2];
  input Real w;
  Real y[2] = [cos(w), sin(w); -sin(w), cos(w)] * x;
  Real z[2] = der(y);
  Real a = 0 * asin(w);
end M;
```

A tool only needs to generate one call of the pure function `cos(w)` in the model `M` – a single call used for both the two elements of the matrix, as well as for the derivative of that matrix. A tool may also skip the possible error for `asin(w)` and assume that `a` is zero.

Examples of restrictions on optimizing pure functions:

```
Real x =
  if noEvent(abs(x)) < 1 then
    asin(x) // Cannot move asin(x) out of if-branch.
  else
    0;
```

```

algorithm
  assertCheck(p, T); // Must call function
algorithm
  if b then
    y := 2 * someOtherFunction(x);
  end if;
  y := y + asin(x);
  y := y + someOtherFunction(x);
  // Cannot evaluate someOtherFunction(x) before asin(x) – unless b is true
  // The reason is that asin(x) may fail and someOtherFunction may hang,
  // and it might be possible to recover from this error.

```

]

12.4 Function Call

Function classes and record constructors (section 12.6) and enumeration type conversions (section 4.8.5.3) can be called as described in this section.

12.4.1 Positional or Named Input Arguments of Functions

A function call has optional positional arguments followed by zero, one or more named arguments, such as

```
f(3.5, 5.76, arg3=5, arg6=8.3);
```

The formal syntax of a function call (simplified by removing reduction expression, section 10.3.4.1):

```

primary :
  component-reference function-call-args

function-call-args :
  "(" [ function-arguments ] ")"

function-arguments :
  function-argument [ "," function-arguments ]
  | named-arguments

named-arguments: named-argument [ "," named-arguments ]

named-argument: IDENT "=" function-argument

function-argument : function-partial-application | expression

```

The interpretation of a function call is as follows: First, a list of unfilled slots is created for all formal input parameters. If there are N positional arguments, they are placed in the first N slots, where the order of the parameters is given by the order of the component declarations in the function definition. Next, for each named argument `identifier = expression`, the `identifier` is used to determine the corresponding slot. The value of the argument is placed in the slot, filling it (it is an error if this slot is already filled). When all arguments have been processed, the slots that are still unfilled are filled with the corresponding default value of the function definition. The default values may depend on other inputs (these dependencies must be acyclical in the function) – the values for those other inputs will then be substituted into the default values (this process may be repeated if the default value for that input depend on another input). The default values for inputs shall not depend on non-input variables in the function. The list of filled slots is used as the argument list for the call (it is an error if any unfilled slots still remain).

Special purpose operators with function syntax defined in the specification shall not be called with named arguments, unless otherwise noted.

The type of each argument must agree with the type of the corresponding parameter, except where the standard type coercion, section 10.6.13, can be used to make the types agree. (See also section 12.4.6 on applying scalar functions to arrays.)

[Example: Assume a function `RealToString` is defined as follows to convert a `Real` number to a `String`:

```
function RealToString
  input Real number;
  input Real precision = 6 "number of significantdigits";
  input Real length = 0 "minimum length of field";
  output String string "number as string";
  ...
end RealToString;
```

Then the following applications are equivalent:

```
RealToString(2.0);
RealToString(2.0, 6, 0);
RealToString(2.0, 6);
RealToString(2.0, precision=6);
RealToString(2.0, length=0);
RealToString(2.0, 6, precision=6); // error: slot is used twice
```

]

12.4.2 Functional Input Arguments to Functions

A functional input argument to a function is an argument of function type. The declared type of such an input formal parameter in a function can be the type-specifier of a partial function that has no replaceable elements. It cannot be the type-specifier of a record or enumeration (i.e., record constructor functions and enumeration type conversions are not allowed in this context). Such an input formal parameter of function type can also have an optional functional default value.

[Example:

```
function quadrature "Integrate function y=integrand(x) from x1 to x2"
  input Real x1;
  input Real x2;
  input Integrand integrand; // Integrand is a partial function,
  see below
  // With default: input Integrand integrand =
  Modelica.Math.sin;
  output Real integral;
algorithm
  integral :=(x2-x1)*(integrand(x1) + integrand(x2))/2;
end quadrature;

partial function Integrand
  input Real u;
  output Real y;
end Integrand;
```

]

A functional argument can be provided in one of the following forms to be passed to a scalar formal parameter of function type in a function call:

- as a function type-specifier (`Parabola` example below),
- as a function partial application (section 12.4.2.1 below),
- as a function that is a component (i.e., a formal parameter of function type of the enclosing function),
- as a function partial application of a function that is a component (example in section 12.4.2.1 below).

In all cases the provided function must be *function type compatible* (section 6.6) to the corresponding formal parameter of function type.

[Example: A function as a positional input argument according to case (a):

```
function Parabola
  extends Integrand;
algorithm
  y := x*x;
end Parabola;
area = quadrature(0, 1, Parabola);
```

The `quadrature2` example below uses a function `integrand` that is a component as input argument according to case (c):

```
function quadrature2 "Integrate function y=integrand(x) from x1 to x2"
  input Real x1;
  input Real x2;
  input Integrand integrand; // Integrand is a partial function type
  output Real integral;
algorithm
  integral := quadrature(x1, (x1+x2)/2, integrand)+ quadrature((x1+x2)/2, x2,
    integrand);
end quadrature2;
```

]

12.4.2.1 Function Partial Application

A function partial application is similar to a function call with certain formal parameters bound to expressions, the specific rules are specified in this section and are not identical to the ones for function call in section 12.4.1. A function partial application returns a partially evaluated function that is also a function, with the remaining not bound formal parameters still present in the same order as in the original function declaration. A function partial application is specified by the `function` keyword followed by a function call to `func_name` giving named formal parameter associations for the formal parameters to be bound, e.g.:

```
function func_name(..., formal_parameter_name = expr, ...)
```

[Note that the keyword `function` in a function partial application differentiates the syntax from a normal function call where some parameters have been left out, and instead supplied via default values.]

The function created by the function partial application acts as the original function but with the bound formal input parameters(s) removed, i.e., they cannot be supplied arguments at function call. The binding occurs when the partially evaluated function is created. A partially evaluated function is *function compatible* (see section 6.6) to the same function where all bound arguments are removed.

[Thus, for checking function type compatibility, bound formal parameters are ignored.]

[Example: Function partial application as argument, positional argument passing, according to case (b) above:

```
model Test
  parameter Integer N;
  Real area;
algorithm
  area := 0;
  for i in 1:N loop
    area := area + quadrature(0, 1, function Sine(A=2, w=i*time));
  end for;
end Test;

function Sine "y = Sine(x,A,w)"
  extends Integrand;
  input Real A;
  input Real w;
algorithm
  y:=A*Modelica.Math.sin(w*x);
end Sine;
```


Call with function partial application as named input argument:

```
area := area + quadrature(0, 1, integrand = function Sine(A=2, w=i*time));
```

]

[Example: Function types are matching after removing the bound arguments *A* and *w* in a function partial application:

```
function Sine2 "y = Sine2(A,w,x)"
  input Real A;
  input Real w;
  input Real x; // Note: x is now last in argument list.
  output Real y;
algorithm
  y:=A*Modelica.Math.sin(w*x);
end Sine2;
area = quadrature(0, 1, integrand = function Sine2(A=2, w=3));
```

The partially evaluated *Sine2* has only one argument: *x* – and is thus type compatible with *Integrand*.]

[Example: Function partial application of a function that is a component, according to case (d) above:

```
partial function SurfaceIntegrand
  input Real x;
  input Real y;
  output Real z;
end SurfaceIntegrand;

function quadratureOnce
  input Real x;
  input Real y1;
  input Real y2;
  input SurfaceIntegrand integrand;
  output Real z;
algorithm
  z := quadrature(y1, y2, function integrand(y=x));
  // This is according to case (d) and needs to bind the 2nd argument
end quadratureOnce;

function surfaceQuadrature
  input Real x1;
  input Real x2;
  input Real y1;
  input Real y2;
  input SurfaceIntegrand integrand;
  output Real integral;
algorithm
  integral := quadrature(x1, x2,
    function quadratureOnce(y1=y1, y2=y2, integrand=integrand));
  // Case (b) and (c)
end surfaceQuadrature;
```

]

12.4.3 Output Formal Parameters of Functions

A function may have more than one output component, corresponding to multiple return values. The only way to use more than the first return value of such a function is to make the function call the right hand side of an equation or assignment. In this case, the left hand side of the equation or assignment shall contain a list of component references within parentheses:

```
(out1, out2, out3) = f(...);
```

The component references are associated with the output components according to their position in the list. Thus output component *i* is set equal to, or assigned to, component reference *i* in the list, where

the order of the output components is given by the order of the component declarations in the function definition. The type of each component reference in the list must agree with the type of the corresponding output component.

A function application may be used as expression whose value and type is given by the value and type of the first output component, if at least one return result is provided.

It is possible to omit left hand side component references and/or truncate the left hand side list in order to discard outputs from a function call.

[Optimizations to avoid computation of unused output results can be automatically deduced by an optimizing compiler.]

[Example: Function `eigen` to compute eigenvalues and optionally eigenvectors may be called in the following ways:

```

ev = eigen(A); // calculate eigenvalues
x = isStable(eigen(A)); // used in an expression
(ev, vr) = eigen(A) // calculate eigenvectors
(ev, vr, vl) = eigen(A) // and also left eigenvectors
(ev, , vl) = eigen(A) // no right eigenvectors
    
```

The function may be defined as:

```

function eigen "calculate eigenvalues and optionally eigenvectors"
  input Real A[:, size(A,1)];
  output Real eigenValues[size(A,1),2];
  output Real rightEigenvectors[size(A,1),size(A,1)];
  output Real leftEigenvectors [size(A,1),size(A,1)];
algorithm
  // The output variables are computed separately (and not, e.g., by one
  // call of a Fortran function) in order that an optimizing compiler can
  // remove
  // unnecessary computations, if one or more output arguments are missing
  // compute eigenvalues
  // compute right eigenvectors using the computed eigenvalues
  // compute left eigenvectors using the computed eigenvalues
end eigen;
    
```

]

The only permissible use of an expression in the form of a list of expressions in parentheses, is when it is used as the left hand side of an equation or assignment where the right hand side is an application of a function.

[Example: The following are illegal:

```

(x+1, 3.0, z/y) = f(1.0, 2.0); // Not a list of component references.
(x, y, z) + (u, v, w) // Not LHS of suitable eqn/assignment.
    
```

]

12.4.4 Initialization and Binding Equations of Components in Functions

Components in a function can be divided into three groups:

- Public components which are input formal parameters.
- Public components which are output formal parameters.
- Protected components which are local variables, parameters, or constants.

When a function is called components of a function do not have start-attributes. However, a binding equation (= `expression`) with an expression may be present for a component.

[Declaration assignments of the form `:= expression` are deprecated, but otherwise identical to binding equations.]

A binding equation for a non-input component initializes the component to this **expression** at the start of every function invocation (before executing the algorithm section or calling the external function). These bindings must be executed in an order where a variable is not used before its binding equations has been executed; it is an error if no such order exists (i.e. the binding must be acyclic).

Binding equations can only be used for components of a function. If no binding equation is given for a non-input component the variable is uninitialized (except for record components where modifiers may also initialize that component). It is an error to use (or return) an uninitialized variable in a function. Binding equations for input formal parameters are interpreted as default arguments, as described in section 12.4.1.

[It is recommended to check for use of uninitialized variables statically – if this is not possible a warning is recommended combined with a run-time check.]

[The properties of components in functions described in this section are also briefly described in section 12.2.]

12.4.5 Flexible Array Sizes and Resizing of Arrays in Functions

[Flexible setting of array dimension sizes of arrays in functions is also briefly described in section 12.2.]

A dimension size not specified with colon(:) for a non-input array component of a function must be given by the inputs or be constant.

[Example:

```
function joinThreeVectors
  input Real v1[:], v2[:], v3[:];
  output Real vres[size(v1,1)+size(v2,1)+size(v3,1)];
algorithm
  vres := cat(1, v1, v2, v3);
end joinThreeVectors;
```

]

- A non-input array component declared in a function with a dimension size specified by colon(:) and no binding equation, can change size according to these special rules: Prior to execution of the function algorithm the dimension size is zero.
- The entire array (without any subscripts) may be assigned with a corresponding array with arbitrary dimension size (the array variable is re-sized).

These rules also apply if the array component is an element of a record component in a function.

[Example: A function to collect the positive elements in a vector:

```
function collectPositive
  input Real x[:];
  output Real xpos[:];
algorithm
  for i in 1:size(x,1) loop
    if x[i]>0 then
      xpos:=cat(1, xpos, x[i:i]);
    end if;
  end for;
end collectPositive;
```

]

12.4.6 Scalar Functions Applied to Array Arguments

Functions with one scalar return value can be applied to arrays element-wise, e.g. if **A** is a vector of reals, then **sin(A)** is a vector where each element is the result of applying the function **sin** to the corresponding element in **A**. Only **function** classes that are transitively non-replaceable (section 6.3.1 and section 7.1.4) may be called vectorized.

Consider the expression $f(\mathbf{arg1}, \dots, \mathbf{argn})$, an application of the function f to the arguments $\mathbf{arg1}, \dots, \mathbf{argn}$. Potential vectorization of this call is defined as follows. For each passed argument, the type of the argument is checked against the type of the corresponding formal parameter of the function:

1. If the types match, nothing is done.
2. If the types do not match, and a type conversion can be applied, it is applied. Continue with step 1.
3. If the types do not match, and no type conversion is applicable, the passed argument type is checked to see if it is an n -dimensional array of the formal parameter type. If it is not, the function call is invalid. If it is, we call this a *foreach argument*.
4. For all foreach arguments, the number and sizes of dimensions must match. If they do not match, the function call is invalid.
5. If no foreach argument exists, the function is applied in the normal fashion, and the result has the type specified by the function definition.
6. The result of the function call expression is an n -dimensional array \mathbf{e} with the same dimension sizes as the foreach arguments. Each element $\mathbf{e}[i, \dots, j]$ is the result of applying f to arguments constructed from the original arguments in the following way:
 - If the argument is not a foreach argument, it is used as-is.
 - If the argument is a foreach argument, the element at index $[i, \dots, j]$ is used.

If more than one argument is an array, all of them have to be the same size, and they are traversed in parallel.

[Example:

```

sin({a, b, c}) = {sin(a), sin(b), sin(c)} // argument is a vector
sin([a, b, c]) = [sin(a), sin(b), sin(c)] // argument may be a matrix
atan2({a, b, c}, {d, e, f}) = {atan2(a, d), atan2(b, e), atan2(c, f)}
    
```

This works even if the function is declared to take an array as one of its arguments. If `pval` is defined as a function that takes one argument that is a `Real` vector and returns a `Real`, then it can be used with an actual argument which is a two-dimensional array (a vector of vectors). The result type in this case will be a vector of `Real`.

```

pval([1,2;3,4]) = [pval([1,2]); pval([3,4])]
sin([1,2;3,4]) = [sin({1,2}); sin({3,4})]
               = [sin(1), sin(2); sin(3), sin(4)]
    
```

```

function Add
  input Real e1, e2;
  output Real sum1;
algorithm
  sum1 := e1 + e2;
end Add;
    
```

`Add(1, [1,2,3])` adds one to each of the elements of the second argument giving the result `[2,3,4]`. However, it is illegal to write `1 + [1,2,3]`, because the rules for the built-in operators are more restrictive.]

12.4.7 Empty Function Calls

An *empty* function call is a call that does not return any results.

[An empty call is of limited use in Modelica since a function call without results does not contribute to the simulation, but it is useful to check assertions and in certain cases for desired side-effects, see section 12.3.]

An empty call can occur either as a kind of “null equation” or “null statement”.

[Example: The empty calls to `eigen()` are examples of a “null equation” and a “null statement”:

```

equation
  Modelica.Math.Matrices.eigen(A); // Empty function call as an equation
algorithm
  Modelica.Math.Matrices.eigen(A); // Empty function call as a statement

```

]

12.5 Built-in Functions

There are basically four groups of built-in functions in Modelica:

- Intrinsic mathematical and conversion functions, see section 3.7.1.
- Derivative and special operators with function syntax, see section 3.7.4.
- Event-related operators with function syntax, see section 3.7.5.
- Built-in array functions, see section 10.3.

Note that when the specification references a function having the name of a built-in function it references the built-in function, not a user-defined function having the same name.

12.6 Record Constructor Functions

Whenever a record is defined, a record constructor function with the same name and in the same scope as the record class is implicitly defined according to the following rules:

The declaration of the record is partially flattened including inheritance, modifications, redeclarations, and expansion of all names referring to declarations outside of the scope of the record to their fully qualified names.

[The partial flattening is performed in order to remove potentially conflicting import-clauses in the record constructor function due to flattening the inheritance tree.]

All record elements (i.e., components and local class definitions) of the partially flattened record declaration are used as declarations in the record constructor function with the following exceptions:

- Component declarations which do not allow a modification (such as `final parameter Real`) are declared as protected components in the record constructor function.
- Prefixes (`constant`, `parameter`, `final`, `discrete`, ...) of the remaining record components are removed.
- The prefix `input` is added to the public components of the record constructor function.

An instance of the record is declared as output parameter using a name not appearing in the record, together with a modification. In the modification, all input parameters are used to set the corresponding record variables.

A record constructor can only be called if the referenced record class is found in the global scope, and thus cannot be modified.

[This allows constructing an instance of a record, with an optional modification, at all places where a function call is allowed.]

Examples:

```

record Complex "Complex number"
  Real re "real part";
  Real im "imaginary part";
end Complex;

function add
  input Complex u, v;
  output Complex w(re=u.re + v.re, im=u.im+v.re);

```

```

end add;

Complex c1, c2;
equation
  c2 = add(c1, Complex(sin(time), cos(time)));

```

In the following example, a convenient data sheet library of components is built up:

```

package Motors
  record MotorData "Data sheet of a motor"
    parameter Real inertia;
    parameter Real nominalTorque;
    parameter Real maxTorque;
    parameter Real maxSpeed;
  end MotorData;

  model Motor "Motor model" // using the generic MotorData
    MotorData data;
    ...
  equation
    ...
  end Motor;

  record MotorI123 = MotorData( // data of a specific motor
    inertia = 0.001,
    nominalTorque = 10,
    maxTorque = 20,
    maxSpeed = 3600) "Data sheet of motor I123";
  record MotorI145 = MotorData( // data of another specific motor
    inertia = 0.0015,
    nominalTorque = 15,
    maxTorque = 22,
    maxSpeed = 3600) "Data sheet of motor I145";
end Motors

model Robot
  import Motors.*;
  Motor motor1(data = MotorI123()); // just refer to data sheet
  Motor motor2(data = MotorI123(inertia=0.0012));
  // data can still be modified (if no final declaration in record)
  Motor motor3(data = MotorI145());
  ...
end Robot;

```

Example showing most of the situations, which may occur for the implicit record constructor function creation. With the following record definitions:

```

package Demo;
  record Record1;
    parameter Real r0 = 0;
  end Record1;

  record Record2
    import Modelica.Math.*;
    extends Record1;
    final constant Real c1 = 2.0;
    constant Real c2;
    parameter Integer n1 = 5;
    parameter Integer n2;
    parameter Real r1 "comment";
    parameter Real r2 = sin(c1);
    final parameter Real r3 = cos(r2);
    Real r4;
    Real r5 = 5.0;
  end Record2;

```

```

    Real r6[n1];
    Real r7[n2];
end Record2;
end Demo;

```

The following record constructor functions are implicitly defined (the name of the output, given in *italic* below, is not defined; it should be chosen to not cause any conflict):

```

package Demo;
function Record1
    input Real r0 = 0;
    output Record1 result (r0 = r0);
end Record1;

function Record2
    input Real r0 = 0;
    input Real c2;
    input Integer n1 = 5;
    input Integer n2;
    input Real r1 "comment"; // the comment also copied from record
    input Real r2 = Modelica.Math.sin(c1);
    input Real r4;
    input Real r5 = 5.0;
    input Real r6[n1];
    input Real r7[n2];
    output Record2 result (r0=r0, c2=c2, n1=n1, n2=n2, r1=r1, r2=r2, r4=r4, r5=r5, r6=r6,
    r7=r7);
protected
    final constant Real c1 = 2.0; // referenced from r2
    final parameter Real r3 = Modelica.Math.cos(r2);
end Record2;
end Demo;

```

and can be applied in the following way

```

Demo.Record2 r1 = Demo.Record2(r0=1, c2=2, n1=2, n2=3, r1=1, r2=2, r4=5, r5=5,
    r6={1,2}, r7={1,2,3});
Demo.Record2 r2 = Demo.Record2(1,2,2,3,1,2,5,5,{1,2},{1,2,3});
parameter Demo.Record2 r3 = Demo.Record2(c2=2, n2=1, r1=1, r4=4, r6=1:5, r7={1})
;

```

The above example is only used to show the different variants appearing with prefixes, but it is not very meaningful, because it is simpler to just use a direct modifier.]

12.6.1 Casting to Record

A constructor of a record *R* can be used to cast an instance *m* of a [model](#), [block](#), [connector](#) class *M* to a value of type *R*, provided that for each component defined in *R* (that do not have a default value) there is also a public component defined in *M* with identical name and type. A nested record component of *R* is handled as follows, if the corresponding component of *M* is a [model/block/connector](#) a nested record constructor is called – otherwise the component is used directly; and the resulting call/component is used as argument to the record constructor *R*. If the corresponding component of *R* in *M* is a conditional component, it is an error. The instance *m* is given as single (un-named) argument to the record constructor of *R*. The interpretation is that *R(m)* is replaced by a record constructor of type *R* where all public components of *M* that are present in *R* are assigned to the corresponding components of *R*. The record cast can be used in vectorized form according to section 12.4.6.

[The problem if *R* would be a conditional component is that the corresponding binding would be illegal since it is not a connect-statement.]

[The record cast operation is uniquely distinguished from a record constructor call, because an argument of the record constructor cannot be a [model](#), [block](#) or [connector](#) instance.]

[Example:

```

connector Flange
  Real phi;
  flow Real tau;
end Flange;

model Model1
  Real m1;
  Boolean b1;
  Flange flange;
end Model1;

model Model2
  Real r1;
  Real r2;
  Integer i2
  Pin p1, p2;
  Model1 sub1;
  protected
  Integer i1;
  ...
end Model2;

record MyFlange
  Real tau;
end MyFlange;

record MyRecord1
  Boolean b1;
  MyFlange flange;
end MyRecord1;

record MyRecord2
  Real r1;
  Integer i2;
  MyRecord1 sub1;
end MyRecord2;

model Model
  Model2 s1;
  Model2 s2[2];
  MyRecord2 rec1 = MyRecord2(s1);
  MyRecord2 rec2[2] = MyRecord2(s2);
  ...
end Model;
// Model is conceptually mapped to
model ModelExpanded
  Model2 s1;
  Model2 s2[2];
  MyRecord2 rec1 = MyRecord2(r1=s1.r1, i2=s1.i2,
  sub1 = MyRecord1(b1=s1.sub1.b1,
  flange = MyFlange(tau=s1.sub1.flange.tau));
  MyRecord2 rec2[2] = {MyRecord2(r1=s2[1].r1, i2=s2[1].i2,
  sub1 = MyRecord1(b1=s2[1].sub1.b1,
  flange = MyFlange(tau=s1[1].sub1.flange.tau)),
  MyRecord2(r1=s2[2].r1, i2=s2[2].i2,
  sub1 = MyRecord1(b1=s2[2].sub1.b1,
  flange = MyFlange(tau=s2[2].sub1.flange.tau))};
  ...
end ModelExpanded;
    
```

]

12.7 Declaring Derivatives of Functions

Derivatives of functions can be declared explicitly using the `derivative` annotation, see section 12.7.1, whereas a function can be defined as a partial derivative of another function using the `der`-operator in a short function definition, see section 12.7.2.

12.7.1 Using the Derivative Annotation

A function declaration can have an annotation `derivative` specifying the derivative function or preferably, for a function written in Modelica, use the `smoothOrder` annotation to indicate that the tool can construct the derivative function automatically, section 18.3. The `derivative` annotation can influence simulation time and accuracy and can be applied to both functions written in Modelica and to external functions. A `derivative` annotation can state that it is only valid under certain restrictions on the input arguments. These restrictions are defined using the following optional attributes: `order` (only a restriction if `order>1`, the default for `order` is 1), `noDerivative`, and `zeroDerivative`. The given derivative-function can only be used to compute the derivative of a function call if these restrictions are satisfied. There may be multiple restrictions on the derivative, in which case they must all be satisfied. The restrictions also imply that some derivatives of some inputs are excluded from the call of the derivative (since they are not necessary). A function may supply multiple derivative functions subject to different restrictions, the first one that can be used (i.e. satisfying the restrictions) will be used for each call.

[This means that the most restrictive derivatives should be written first.]

[Example:

```
function foo0 annotation(derivative=foo1);
end foo0;

function foo1 annotation(derivative(order=2)=foo2);
end foo1;

function foo2 end foo2;
```

]

The inputs to the derivative function of `order` 1 are constructed as follows:

- First are all inputs to the original function, and after all them we will in order append one derivative for each input containing reals. These common inputs must have the same name, type, and declaration order for the function and its derivative.
- The outputs are constructed by starting with an empty list and then in order appending one derivative for each output containing reals. The outputs must have the same type and declaration order for the function and its derivative.

If the Modelica function call is a n th derivative ($n \geq 1$), i.e. this function call has been derived from an $(n-1)$ th derivative by differentiation inside the tool, an `annotation(order=n+1)=...`, specifies the $(n+1)$ th derivative, and the $(n+1)$ th derivative call is constructed as follows:

- The input arguments are appended with the $(n+1)$ th derivative, which are constructed in order from the n th `order` derivatives.
- The output arguments are similar to the output argument for the n th derivative, but each output is one higher in derivative order. The outputs must have the same type and declaration order for the function and its derivative.

[The restriction that only the result of differentiation can use higher-order derivatives ensures that the derivatives `x`, `der_x`, ... are in fact derivatives of each other. Without that restriction we would have both `der(x)` and `x_der` as inputs (or perform advanced tests to verify that they are the same).]

[Example: Given the declarations

```
function foo0
...

```

```

input Real x;
input Boolean linear;
input ...;
output Real y;
...
annotation(derivative=foo1);
end foo0;

function foo1
...
input Real x;
input Boolean linear;
input ...;
input Real der_x;
...
output Real der_y;
...
annotation(derivative(order=2)=foo2);
end foo1;

function foo2
...
input Real x;
input Boolean linear;
input ...;
input Real der_x;
...;
input Real der_2_x;
...
output Real der_2_y;
...

```

the equation

$$(\dots, y(t), \dots) = \text{foo0}(\dots, x(t), b, \dots)$$

implies that:

$$\begin{aligned}
 (\dots, \frac{dy(t)}{dt}, \dots) &= \text{foo1}(\dots, x(t), b, \dots, \dots, \frac{dx(t)}{dt}, \dots) \\
 (\dots, \frac{d^2y(t)}{dt^2}, \dots) &= \text{foo2}(\dots, x(t), b, \dots, \frac{dx(t)}{dt}, \dots, \dots, \frac{d^2x(t)}{dt^2}, \dots)
 \end{aligned}$$

]

An input or output to the function may be any simple type (**Real**, **Boolean**, **Integer**, **String** and enumeration types) or a record. For a record containing Reals the corresponding derivative uses a derivative record, that only contain the real-predefined types and sub-records containing reals (handled recursively) from the original record. When using `smoothOrder`, then the derivative record is automatically constructed. The function must have at least one input containing reals. The output list of the derivative function shall not be empty.

[Example: Here is one example use case with records mixing **Real** and non-**Real** as inputs and outputs:

```

record ThermodynamicState "Thermodynamic state"
  SpecificEnthalpy h "Specific enthalpy";
  AbsolutePressure p "Pressure";
  Integer phase(min=1, max=2, start=1);
end ThermodynamicState;

record ThermoDynamicState_der "Derivative"
  SpecificEnthalpyDerivative h "Specific enthalpy derivative";
  PressureDerivative p "Pressure derivative";
  // Integer input is skipped

```

```

end ThermodynamicState_der;

function density
  input ThermodynamicState state "Thermodynamic state";
  output Density d "Density";
algorithm
  ...
  annotation(derivative=density_der);
end density;

function density_der
  input ThermodynamicState state "Thermodynamic state";
  input ThermodynamicState_der state_der;
  output DensityDerivative d "Density derivative";
algorithm
  ...
end density_der;

function setState_ph
  input Pressure p;
  input SpecificEnthalpy h;
  input Integer phase = 0;
  output ThermodynamicState state;
algorithm
  ...
  annotation(derivative = setState_ph_der);
end setState_ph;

function setState_ph_der
  input Pressure p;
  input SpecificEnthalpy h;
  input Integer phase;
  input PressureDerivative p_der;
  input SpecificEnthalpyDerivative h_der;
  output ThermodynamicState_der state_der;
algorithm
  ...
end setState_ph_der;

ThermodynamicState state1 = setState_ph(p=..., h=..., phase=...);
Density rho1=density(state1);
DensityDerivative d_rho1=der (rho1);
Density rho2=density(setState_ph(p=..., h=..., phase=...));
DensityDerivative d_rho2=der (rho2);

```

- zeroDerivative = inputVar1 { , zeroDerivative = inputVar2 }

The derivative function is only valid if `inputVar1` (and `inputVar2` etc.) are independent of the variables the function call is differentiated with respect to (i.e. that the derivative of `inputVar1` is zero). The derivative of `inputVar1` (and `inputVar2` etc.) are excluded from the argument list of the derivative-function. If the derivative-function also specifies a derivative the common variables should have consistent `zeroDerivative`.

[Assume that function `f` takes a matrix and a scalar. Since the matrix argument is usually a parameter expression it is then useful to define the function as follows (the additional derivative = `f_general_der` is optional and can be used when the derivative of the matrix or offset is non-zero). Note that `f_der` must have `zeroDerivative` for both `y` and `offset`, but `f_general_der` shall not have `zeroDerivative` for either of them (it may `zeroDerivative` for `x_der`, `y_der`, or `offset_der`).

```

function f "Simple table lookup"
  input Real x;
  input Real y[:, 2];

```

```

    input Real offset;
    output Real z;
  algorithm
    ...
    annotation(derivative(zeroDerivative=y, zeroDerivative=offset)= f_der,
                derivative=f_general_der);
  end f;

function f_der "Derivative of simple table lookup"
  input Real x;
  input Real y[:, 2];
  input Real offset;
  input Real x_der;
  output Real z_der;
algorithm
  ...
  annotation(derivative(zeroDerivative=y, zeroDerivative=offset, order=2) =
              f_der2);
end f_der;

function f_der2 "Second derivative of simple table lookup"
  input Real x;
  input Real y[:, 2];
  input Real offset;
  input Real x_der;
  input Real x_der2;
  output Real z_der2;
algorithm
  ...
end f_der2;

function f_general_der "Derivative of table lookup taking
into account varying tables"
  input Real x;
  input Real y[:, 2];
  input Real offset;
  input Real x_der;
  input Real y_der[:, 2];
  input Real offset_der;
  output Real z_der;
algorithm
  ...
  //annotation(derivative(order=2) = f_general_der2);
end f_general_der;

```

- noDerivative = inputVar1

The derivative of inputVar1 is excluded from the argument list of the derivative-function. This relies on assumptions on the arguments to the function; and the function should document these assumptions (it is not always straightforward to verify them). In many cases even the undifferentiated function will only behave correctly under these assumptions.

The inputs excluded using zeroDerivative or noDerivative may be of any type (including types not containing reals).

[Assume that function fg is defined as a composition $f(x, g(x))$. When differentiating f it is useful to give the derivative under the assumption that the second argument is defined in this way:

```

function fg
  input Real x;
  output Real z;
algorithm
  z := f(x, g(x));

```

```

end fg;

function f
  input Real x;
  input Real y;
  output Real z;
algorithm
  ...
  annotation(derivative(noDerivative = y) = f_der);
end f;

function f_der
  input Real x;
  input Real y;
  input Real x_der;
  output Real z_der;
algorithm
  ...
end f_der;
    
```

This is useful if g represents the major computational effort of fg .]

12.7.2 Partial Derivatives of Functions

A class defined as:

```
IDENT "=" der "(" name "," IDENT { "," IDENT } ")" comment
```

is the partial derivative of a function, and may only be used as declarations of functions.

The semantics is that a function (and only a function) can be specified in this form, defining that it is the partial derivative of the function to the right of the equal sign (looked up in the same way as a short class definition, and the looked up name must be a function), and partially differentiated with respect to each IDENT in order (starting from the first one). The IDENT must be `Real` inputs to the function.

The comment allows a user to comment the function (in the info-layer and as one-line description, and as icon).

[Example: The specific enthalpy can be computed from a Gibbs-function as follows:

```

function Gibbs
  input Real p, T;
  output Real g;
algorithm
  ...
end Gibbs;
function Gibbs_T = der(Gibbs, T);
function specificEnthalpy
  input Real p, T;
  output Real h;
algorithm
  h := Gibbs(p, T) - T * Gibbs_T(p, T);
end specificEnthalpy;
    
```

]

12.8 Declaring Inverses of Functions

Every function with one output formal parameter may have one or more **inverse** annotations to define inverses of this function:

```

function f1
  input A1 u1;
  ...
    
```

```

input T1 uk;
...
input Am um = am;
...
input An un;
output T2 y;
algorithm
...
  annotation(inverse(uk = f2(..., y, ...), ui = f3(..., y, ...), ...));
end f1;
    
```

The meaning is that function f_2 is one inverse to function f_1 where the previous output y is now an input and the previous input u_k is now an output. More than one inverse can be defined within the same `inverse` annotation. Several inverses are separated by commas.

[The inverse requires that for all valid values of the input arguments of $f_2(\dots, y, \dots)$ and u_k being calculated as $u_k := f_2(\dots, y, \dots)$ implies the equality $y = f_1(\dots, u_k, \dots)$ up to a certain precision.]

Function f_1 can have any number and types of formal parameters with and without default value. The restriction is that the *number of unknown variables* (see section 4.7) in the output formal parameter of both f_1 and f_2 must be the same and that f_2 should have a union of output and formal parameters that is the same or a sub-set of that union for f_1 , but the order of the formal parameters may be permuted.

[Example: Same union of variables:

```

function h_pTX
  input Real p "pressure";
  input Real T "temperature";
  input Real X[:] "mass fractions";
  output Real h "specific enthalpy";
algorithm
...
  annotation(inverse(T = T_phX(p,h,X)));
end h_pTX;

function T_phX
  input Real p "pressure";
  input Real h "specific enthalpy";
  input Real X[:] "mass fractions";
  output Real T "temperature";
algorithm
...
end T_phX;
    
```

]

The sub-set case is useful if f_1 computes the inverse of f_2 within a region, or up to a certain tolerance. Then, f_1 may specify f_2 as inverse with fewer arguments, skipping the arguments for tolerance and/or the region.

[Example:

```

function inv_sine
  input Real x;
  input Real angleOrig;
  output Real angle;
  // Finds sine(angle) = x with angle closest to angleOrig.
algorithm
...
  annotation(inverse(x = sine(angle)));
end inv_sine;

function sine
    
```

```

input Real angle;
output Real x;
algorithm
  x := sin(angle);
  // Note: No inverse.
end sine;

```

]

12.9 External Function Interface

Here, the word *function* is used to refer to an arbitrary external routine, whether or not the routine has a return value or returns its result via output parameters (or both). The Modelica external function call interface provides the following:

- Support for external functions written in C (specifically C89) and FORTRAN 77. Other languages, e.g. C++ and Fortran 90, may be supported in the future, and provided the function is link-compatible with C89 or FORTRAN 77 it can be written in any language.
- Mapping of argument types from Modelica to the target language and back.
- Natural type conversion rules in the sense that there is a mapping from Modelica to standard libraries of the target language.
- Handling arbitrary parameter order for the external function.
- Passing arrays to and from external functions where the dimension sizes are passed as explicit integer parameters.
- Handling of external function parameters which are used both for input and output, by passing an output that has a binding equation to the external function.

[Binding equations are executed prior to calling the external function.]

The format of an external function declaration is as follows.

```

function IDENT description-string
  { component-clause ";" }
  [ protected { component-clause ";" } ]
external [ language-specification ]
  [ external-function-call ]
  [ annotation ] ";"
  [ annotation ";" ]
end IDENT;

```

Components in the public part of an external function declaration shall be declared either as **input** or **output**.

[This is just as for any other function. The components in the protected part allow local variables for temporary storage to be declared.]

The **language-specification** must currently be one of "builtin" (deprecated), "C", "C..." (for one of the specific C standards like C89, C99, and C11 – specifying that it relies on the C standard library of that version) or "FORTRAN 77". Unless the external language is specified, it is assumed to be "C".

[The intended use of e.g. C99 is to detect if the user tries to link with a C99-function using a C89 compiler.]

The deprecated "builtin" specification is only used for the elementary mathematical functions described in section 3.7.3. The external function call mechanism for "builtin" functions is implementation-defined.

[Typically, for functions from the standard C library, the prototype of the function is provided but no Library annotation. Currently, there are no other builtin functions defined in Modelica.]

[Example:

```

package Modelica
  package Math
    function sin
      input Real x;
      output Real y;
      external "builtin"
        y = sin(x);
      end sin;
    end Math;
  end Modelica;

model UserModel
  parameter Real p = Modelica.Math.sin(2);
end UserModel;
    
```

]

The `external-function-call` specification allows functions whose prototypes do not match the default assumptions as defined below to be called. It also gives the name used to call the external function. If the external call is not given explicitly, this name is assumed to be the same as the Modelica name.

The only permissible kinds of expressions in the argument list are component references, scalar constants, and the function `size` applied to an array and a constant dimension number. The annotations are used to pass additional information to the compiler when necessary.

A component reference to a component that is part of an input or output is treated the same way as a top-level input or output in the external call.

If the function has `annotation(Include="includeDirective")`, section 12.9.4 it is assumed that it contains an actual prototype and no prototype shall be automatically generated. In that case the input argument pointers shall be const pointers (indicating that they do not modify the inputs), and non-const pointers are deprecated. The optional `external-function-call` is still used for determining the name of the function, and order of arguments, as described below.

12.9.1 Argument type Mapping

The arguments of the external function are declared in the same order as in the Modelica declaration, unless specified otherwise in an explicit external function call. Protected variables (i.e. temporaries) are passed in the same way as outputs, whereas constants and size-expression are passed as inputs.

12.9.1.1 Simple Types

Arguments of *simple* types are by default mapped as follows for C:

<i>Modelica</i>	<i>C</i>	
	<i>Input</i>	<i>Output</i>
Real	double	double *
Integer	int	int *
Boolean	int	int *
String	const char *	const char **
Enumeration type	int	int *

An exception is made when the argument is of the form `size(..., ...)`. In this case the corresponding C type is `size_t`.

Strings are NUL-terminated (i.e., terminated by `'\0'`) to facilitate calling of C functions. When returning a non-literal string, see section 12.9.6.2 for details on memory allocation.

Boolean values are mapped to C such that `false` in Modelica is 0 in C and `true` in Modelica is 1 in C. If the returned value from C is 0 it is treated as `false` in Modelica; otherwise as `true`.

[It is recommended that the C function should interpret any non-zero value as true.]

Arguments of simple types are by default mapped as follows for FORTRAN 77:

<i>Modelica</i>	<i>FORTRAN 77</i>	
	<i>Input</i>	<i>Output</i>
Real	DOUBLE PRECISION	DOUBLE PRECISION
Integer	INTEGER	INTEGER
Boolean	LOGICAL	LOGICAL
String	<i>Special</i>	<i>Not available</i>
Enumeration type	INTEGER	INTEGER

Sending string literals to FORTRAN 77 subroutines/functions is supported for Lapack/Blas-routines, and the strings are NUL-terminated for compatibility with C. Returning strings from FORTRAN 77 subroutines/functions is currently not supported.

Enumeration types used as arguments are mapped to type int when calling an external C function, and to type INTEGER when calling an external FORTRAN function. The i :th enumeration literal is mapped to integer value i , starting at 1.

Return values are mapped to enumeration types analogously: integer value 1 is mapped to the first enumeration literal, 2 to the second, etc. Returning a value which does not map to an existing enumeration literal for the specified enumeration type is an error.

12.9.1.2 Arrays

Unless an explicit function call is present in the external-clause, an array is passed by its address followed by n arguments of type `size_t` with the corresponding array dimension sizes, where n is the number of dimensions.

[The type `size_t` is a C unsigned integer type.]

Arrays are by default stored in row-major order when calling C functions and in column-major order when calling FORTRAN 77 functions. These defaults can be overridden by the `arrayLayout` annotation. See the example below.

The table below shows the mapping of an array argument in the absence of an explicit external function call when calling a C function. The type T is allowed to be any of the simple types which can be passed to C as defined in section 12.9.1.1 or a record type as defined in section 12.9.1.3 and it is mapped to the type T' as defined in these sections for input arguments.

<i>Modelica</i>	<i>C</i>
	<i>Input and output</i>
T[dim_1]	T' *, <code>size_t dim₁</code>
T[dim_1, dim_2]	T' *, <code>size_t dim₁, size_t dim₂</code>
T[dim_1, \dots, dim_n]	T' *, <code>size_t dim₁, \dots, size_t dim_n</code>

The method used to pass array arguments to FORTRAN 77 functions in the absence of an explicit external function call is similar to the one defined above for C: first the address of the array, then the dimension sizes as integers. See the table below. The type T is allowed to be any of the simple types which can be passed to FORTRAN 77 as defined in section 12.9.1.1 and it is mapped to the type T' as defined in that section.

<i>Modelica</i>	<i>FORTRAN 77</i>
	<i>Input and output</i>
T[dim_1]	T', INTEGER dim_1
T[dim_1, dim_2]	T', INTEGER dim_1 , INTEGER dim_2
T[dim_1, \dots, dim_n]	T', INTEGER $dim_1, \dots, INTEGER dim_n$

[Example: The following two examples illustrate the default mapping of array arguments to external C and FORTRAN 77 functions.]

```
function foo
  input Real a[:, :, :];
  output Real x;
  external;
end foo;
```

The corresponding C prototype is as follows:

```
double foo(double *, size_t, size_t, size_t);
```

If the external function is written in FORTRAN 77, i.e.:

```
function foo
  input Real a[:, :, :];
  output Real x;
  external "FORTRAN 77";
end foo;
```

the default assumptions correspond to a FORTRAN 77 function defined as follows:

```
FUNCTION foo(a, d1, d2, d3)
  DOUBLE PRECISION(d1, d2, d3) a
  INTEGER          d1
  INTEGER          d2
  INTEGER          d3
  DOUBLE PRECISION          foo
  ...
END
```

]

When an explicit call to the external function is present, the array and the sizes of its dimensions must be passed explicitly.

[Example: This example shows how arrays can be passed explicitly to an external FORTRAN 77 function when the default assumptions are unsuitable.

```
function foo
  input Real x[:];
  input Real y[size(x,1),:];
  input Integer i;
  output Real u1[size(y,1)];
  output Integer u2[size(y,2)];
  external "FORTRAN 77"
  myfoo(x, y, size(x,1), size(y,2), u1, i, u2);
end foo;
```

The corresponding FORTRAN 77 subroutine would be declared as follows:

```
SUBROUTINE myfoo(x, y, n, m, u1, i, u2)
  DOUBLE PRECISION(n) x
  DOUBLE PRECISION(n,m) y
  INTEGER n
  INTEGER m
  DOUBLE PRECISION(n) u1
  INTEGER i
  DOUBLE PRECISION(m) u2
  ...
END
```

This example shows how to pass an array in column major order to a C function.

```
function fie
  input Real[:, :] a;
  output Real b;
  external;
  annotation(arrayLayout = "columnMajor");
end fie;
```

This corresponds to the following C prototype:

```
double fie(double *, size_t, size_t);
```

]

12.9.1.3 Records

Mapping of record types is only supported for C. A Modelica record class that contains simple types, other record elements, is mapped as follows:

- The record class is represented by a struct in C.
- Each element of the Modelica record is mapped to its corresponding C representation.
- The elements of the Modelica record class are declared in the same order in the C struct.
- Arrays cannot be mapped.

Records are passed by reference (i.e. a pointer to the record is being passed).

[*Example:*

```
record R
  Real x;
  Real z;
end R;
```

is mapped to:

```
struct R {
  double x;
  double z;
};
```

]

12.9.2 Return Type Mapping

If there is a single output parameter and no explicit call of the external function, or if there is an explicit external call in the form of an equation, in which case the LHS must be one of the output parameters, the external routine is assumed to be a value-returning function. Mapping of the return type of functions is performed as indicated in the table below. Storage for arrays as return values is allocated by the calling routine, so the dimensions of the returned array are fixed at call time. Otherwise the external function is assumed not to return anything; i.e., it is really a procedure or, in C, a void-function.

[*In the case of an external function not returning anything, argument type mapping according to section 12.9.1.1 is performed in the absence of any explicit external function call.*]

Return types are by default mapped as follows for C and FORTRAN 77:

<i>Modelica</i>	<i>C</i>	<i>FORTRAN 77</i>
Real	double	DOUBLE PRECISION
Integer	int	INTEGER
Boolean	int	LOGICAL
String	const char*	<i>Not allowed</i>
T[<i>dim</i> ₁ , ..., <i>dim</i> _{<i>n</i>}]	<i>Not allowed</i>	<i>Not allowed</i>
Enumeration type	int	INTEGER
Record	See section 12.9.1.3	<i>Not allowed</i>

The element type T of an array can be any simple type as defined in section 12.9.1.1 or, for C, a record type is returned as a value of the record type defined in section 12.9.1.3.

12.9.3 Aliasing

Any potential aliasing in the external function is the responsibility of the tool and not the user. An external function is not allowed to internally change the inputs (even if they are restored before the end of the function).

[*Example:*

```
function foo
  input Real x;
  input Real y;
  output Real z = x;
  external "FORTRAN 77"
    myfoo(x, y, z);
end foo;
```

The following Modelica function:

```
function f
  input Real a;
  output Real b;
algorithm
  b := foo(a, a);
  b := foo(b, 2 * b);
end f;
```

can on most systems be transformed into the following C function:

```
double f(double a) {
  extern void myfoo_(double*, double*, double*);
  double b, temp1, temp2;

  myfoo_(&a, &a, &b);
  temp1 = 2 * b;
  temp2 = b;
  myfoo_(&b, &temp1, &temp2);

  return temp2;
}
```

The reason for not allowing the external function to change the inputs is to ensure that inputs can be stored in static memory and to avoid superfluous copying (especially of matrices). If the routine does not satisfy the requirements the interface must copy the input argument to a temporary. This is rare but occurs e.g. in `dormlq` in some Lapack implementations. In those special cases the writer of the external interface have to copy the input to a temporary. If the first input was changed internally in `myfoo` the designer of the interface would have to change the interface function `foo` to:

```
function foo
  input Real x;
  protected Real xtemp = x; // Temporary used because myfoo changes its input
  public input Real y;
  output Real z;
  external "FORTRAN 77"
    myfoo(xtemp, y, z);
end foo;
```

Note that we discuss input arguments for Fortran-routines even though FORTRAN 77 does not formally have input arguments and forbid aliasing between any pair of arguments to a function (Section 15.9.3.6 of X3J3/90.4). For the few (if any) FORTRAN 77 compilers that strictly follow the standard and are unable to handle aliasing between input variables the tool must transform the first call of `foo` into:

```
temp1 = a; /* Temporary to avoid aliasing */
myfoo_(&a, &temp1, &b);
```

The use of the function `foo` in Modelica is uninfluenced by these considerations.]

12.9.4 Annotations for External Libraries and Include Files

The following annotations are useful in the context of calling external functions from Modelica, and they should occur on the external-clause and no other standard annotations should occur on the external-

clause. They can all specify either a scalar value or an array of values as indicated below for the `Library` annotation:

- The `annotation(Library="libraryName")`, used by the linker to include the library file where the compiled external function is available.
- The `annotation(Library=("libraryName1", "libraryName2"))`, used by the linker to include the library files where the compiled external function is available and additional libraries used to implement it. For shared libraries it is recommended to include all non-system libraries in this list.
- The `annotation(Include="includeDirective")`, used to include source files needed for calling the external function in the code generated by the Modelica compiler. The included code should be valid C89 code.

[Examples of files that can be included are header files or source files that contain the functions referenced in the external function declaration.]

- The `annotation(IncludeDirectory="modelica://ModelicaLibraryName/Resources/Include")`, used to specify a location for header files. The preceding one is the default and need not be specified; but another location could be specified by using an URI name for the include directory, see section 13.5.
- The `annotation(LibraryDirectory="modelica://ModelicaLibraryName/Resources/Library")`, used to specify a location for library files. The preceding one is the default and need not be specified; but another location could be specified by using an URI name for the library directory, see section 13.5. Different versions of one object library can be provided (e.g. for Windows and for Linux) by providing a *platform* directory below the `LibraryDirectory`. If no platform directory is present, the object library must be present in the `LibraryDirectory`. The following *platform* names are standardized:
 - "win32" (Microsoft Windows 32 bit)
 - "win64" (Microsoft Windows 64 bit)
 - "linux32" (Linux Intel 32 bit)
 - "linux64" (Linux Intel 64 bit)

The win32 or win64 directories may contain gcc47, vs2010, vs2012 for specific versions of these compilers and these are used instead of the general win32 or win64 directories, and similarly for other platforms.

The library on Windows may refer to a lib-file (static library), both a lib- and dll-file (in this case the lib-file is an import-library), or just a dll-file. It shall not refer to an obj-file.

If the directory for the specific compiler version is missing the platform specific directory is used.

[A tool may give diagnostics if the directory corresponding to the selected compiler version is missing. The directories may use symbolic links or use a text-file as described below: e.g. a text-file vs2008 containing the text ../win32/vs2005 (or vs2005) suggesting that it is compatible with vs2005.]

The `ModelicaLibraryName` used for `IncludeDirectory` and `LibraryDirectory` indicates the top-level class where the annotation is found in the Modelica source code.

[Example: Use of external functions and of object libraries:

```

package ExternalFunctions
  model Example
    Real x(start = 1.0), y(start = 2.0);
  equation
    der(x) = -ExternalFunc1(x);
    der(y) = -ExternalFunc2(y);
  end Example;

  model OtherExample
    Real x(start = 1.0);
  equation
    der(x) = -ExternalFunc3(x);
  end OtherExample;

```

```

function ExternalFunc1
  input Real x;
  output Real y;
  external "C"
    y = ExternalFunc1_ext(x)
    annotation(Library = "ExternalLib1",
              Include = "#include \"ExternalFunc1.h\"");
end ExternalFunc1;

function ExternalFunc2
  input Real x;
  output Real y;
  external "C"
    annotation(Library = "ExternalLib2",
              Include = "#include \"ExternalFunc2.h\"");
end ExternalFunc2;

function ExternalFunc3
  input Real x;
  output Real y;
  external "C"
    annotation(Include = "#include \"ExternalFunc3.c\"");
end ExternalFunc3;
end ExternalFunctions;

package MyExternalFunctions
  extends ExternalFunctions;
end MyExternalFunctions;
    
```

Directory structure:

```

ExternalFunctions
  package.mo // contains the Modelica code from above
  Resources
    Include // contains the include files
      ExternalFunc1.h // C header file
      ExternalFunc2.h // C header file
      ExternalFunc3.c // C source file
    Library // contains the object libraries for different
  platforms
    win32
      ExternalLib1.lib // static link library for VisualStudio
      ExternalLib2.lib // statically linking the dynamic link library
      ExternalLib2.dll // dynamic link library (with manifest)
    linux32
      libExternalLib1.a // static link library
      libExternalLib2.so // shared library
  MyExternalFunctions
    package.mo
    
```

Note that calling `MyExternalFunctions.ExternalFunc1` will use header and library files from `ExternalFunction`, the `ExternalFunctions`. Example will not use `ExternalFunc3.c`, and one library file may contain multiple functions.

Header file for the function in the dynamic link / shared library `ExternalLib2` so that the desired functions are defined to be exported for Microsoft VisualStudio and for GNU C compiler (note, for Linux it is recommended to use the compiler option `-fPIC` to build shared libraries or object libraries that are later transformed to a shared library):

```

// File ExternalFunc2.h
#ifdef __cplusplus
extern "C" {
#endif
    
```

```

#ifdef _MSC_VER
#ifdef EXTERNAL_FUNCTION_EXPORT
# define EXTLIB2_EXPORT __declspec( dllexport )
#else
# define EXTLIB2_EXPORT __declspec( dllimport )
#endif
#elif __GNUC__ >= 4
/* In gnuc, all symbols are by default exported. It is still often
useful,
to not export all symbols but only the needed ones */
# define EXTLIB2_EXPORT __attribute__ ((visibility("default")))
#else
# define EXTLIB2_EXPORT
#endif

EXTLIB2_EXPORT void ExternalFunc2(<function arguments>);

#ifdef __cplusplus
}
#endif

```

]

The `Library` name and the `LibraryDirectory` name in the function annotation are mapped to a linkage directive in a compiler-dependent way thereby selecting the object library suited for the respective computer platform.

12.9.5 Examples

12.9.5.1 Input Parameters, Function Value

[Example: Here all parameters to the external function are input parameters. One function value is returned. If the external language is not specified, the default is "C", as below.

```

function foo
  input Real x;
  input Integer y;
  output Real w;
  external;
end foo;

```

This corresponds to the following C prototype:

```
double foo(double, int);
```

Example call in Modelica:

```
z = foo(2.4, 3);
```

Translated call in C:

```
z = foo(2.4, 3);
```

]

12.9.5.2 Arbitrary Placement of Output Parameters, No External Function Value

[Example: In the following example, the external function call is given explicitly which allows passing the arguments in a different order than in the Modelica version.

```

function foo
  input Real x;
  input Integer y;
  output Real u1;

```

```

output Integer u2;
external "C"
  myfoo(x, u1, y, u2);
end foo;
    
```

This corresponds to the following C prototype:

```
void myfoo(double, double *, int, int *);
```

Example call in Modelica:

```
(z1,i2) = foo(2.4, 3);
```

Translated call in C:

```
myfoo(2.4, \&z1, 3, \&i2);
```

]

12.9.5.3 External Function with Both Function Value and Output Variable

[Example: The following external function returns two results: one function value and one output parameter value. Both are mapped to Modelica output parameters.

```

function foo
  input Real x;
  input Integer y;
  output Real funcvalue;
  output Integer out1;
external "C"
  funcvalue = myfoo(x, y, out1);
end foo;
    
```

This corresponds to the following C prototype:

```
double myfoo(double, int, int *);
```

Example call in Modelica:

```
(z1,i2) = foo(2.4, 3);
```

Translated call in C:

```
z1 = myfoo(2.4, 3, \&i2);
```

]

12.9.6 Utility Functions

This section describes the utility functions declared in `ModelicaUtilities.h`, which can be called in external Modelica functions written in C.

12.9.6.1 Utility Functions for Reporting Errors

The functions listed below produce a message in different ways.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>ModelicaMessage(string)</code>	Message with fixed string	
<code>ModelicaWarning(string)</code>	Warning with fixed string	Function 12.1
<code>ModelicaError(string)</code>	Error with fixed string	
<code>ModelicaFormatMessage(format, ...)</code>	Message with <code>printf</code> style formatting	
<code>ModelicaFormatWarning(format, ...)</code>	Warning with <code>printf</code> style formatting	Function 12.2
<code>ModelicaFormatError(format, ...)</code>	Error with <code>printf</code> style formatting	
<code>ModelicaVFormatMessage(format, ap)</code>	Message with <code>vprintf</code> style formatting	
<code>ModelicaVFormatWarning(format, ap)</code>	Warning with <code>vprintf</code> style formatting	Function 12.3
<code>ModelicaVFormatError(format, ap)</code>	Error with <code>vprintf</code> style formatting	

The *Message*-functions only produce the message, but the *Warning*- and *Error*-functions combine this with error handling as follows.

The *Warning*-functions view the message as a warning and can skip duplicated messages similarly as an `assert` with `level = AssertionLevel.Warning` in the Modelica code.

The *Error*-functions never return to the calling function, but handle the error similarly to an `assert` with `level = AssertionLevel.Error` in the Modelica code.

Function 12.1 `ModelicaMessage`, `ModelicaWarning`, `ModelicaError`

```
void ModelicaMessage(const char* string)
void ModelicaWarning(const char* string)
void ModelicaError(const char* string)
```

Output the fixed message string (no format control).

Function 12.2 `ModelicaFormatMessage`, `ModelicaFormatWarning`, `ModelicaFormatError`

```
void ModelicaFormatMessage(const char* format, ...)
void ModelicaFormatWarning(const char* format, ...)
void ModelicaFormatError(const char* format, ...)
```

Output the message under the same format control as the C function `printf`.

Function 12.3 `ModelicaVFormatMessage`, `ModelicaVFormatWarning`, `ModelicaVFormatError`

```
void ModelicaVFormatMessage(const char* format, va_list ap)
void ModelicaVFormatWarning(const char* format, va_list ap)
void ModelicaVFormatError(const char* format, va_list ap)
```

Output the message under the same format control as the C function `vprintf`.

12.9.6.2 Utility Functions for Allocating Strings

The functions listed below are related to string allocation.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>ModelicaAllocateString(len)</code>	Allocate or error	Function 12.4
<code>ModelicaAllocateStringWithErrorReturn(len)</code>	Allocate or null	Function 12.5
<code>ModelicaDuplicateString(str)</code>	Duplicate or error	Function 12.6
<code>ModelicaDuplicateStringWithErrorReturn(str)</code>	Duplicate or null	Function 12.7

Function 12.4 `ModelicaAllocateString`

```
char* ModelicaAllocateString(size_t len)
```

Allocate memory for a writeable non-literal string which is used as a return argument of an external Modelica function. It allocates `len + 1` characters and the last one is set to NUL. If an error occurs, this function does not return, but calls `ModelicaError`.

Function 12.5 `ModelicaAllocateStringWithErrorReturn`

```
char* ModelicaAllocateStringWithErrorReturn(size_t len)
```

Same as `ModelicaAllocateString`, except that in case of error, the function returns 0. This allows the external function to close files and free other open resources in case of error. After cleaning up resources, use `ModelicaError` or `ModelicaFormatError` to signal the error.

Function 12.6 `ModelicaDuplicateString`

```
char* ModelicaDuplicateString(const char* str)
```

Returns a writeable duplicate of the NUL-terminated string `str`. If an error occurs, this function does not return, but calls `ModelicaError`.

Function 12.7 `ModelicaDuplicateStringWithErrorReturn`

```
char* ModelicaDuplicateStringWithErrorReturn(const char* str)
```

Same as `ModelicaDuplicateString`, except that in case of error, the function returns 0. This allows the external function to close files and free other open resources in case of error. After cleaning up resources, use `ModelicaError` or `ModelicaFormatError` to signal the error.

The valid return values for an external function returning a `String` are:

- A literal `String`.
- A string given as `String` input to the external function.
- A string pointer returned by one the functions in the table above.

Thus if an external function wants to create a non-literal string it must be allocated with one of the functions in this section, e.g., `ModelicaAllocateString`. After return of the external function, the Modelica environment is responsible for the memory allocated with `ModelicaAllocateString` (e.g., to free this memory, when appropriate). It is not allowed to access memory that was allocated with `ModelicaAllocateString` in a previous call of this external function.

[Memory that is not passed to the Modelica simulation environment, such as memory that is freed before leaving the function, or in an `ExternalObject`, see section 12.9.7, should be allocated with the standard C mechanisms, like `calloc`.]

[The reason why one should avoid, for instance, `malloc` for string allocation is that a Modelica simulation environment may have its own allocation scheme, e.g., a special stack for local variables of a function.]

12.9.7 External Objects

External functions may need to store their internal memory between function calls. Within Modelica this memory is defined as instance of the predefined class `ExternalObject` according to the following rules:

- There is a predefined partial class `ExternalObject`.

[Since the class is partial, it is not possible to define an instance of this class.]

- An external object class shall be directly extended from `ExternalObject`, shall have exactly two function definitions, called `constructor` and `destructor`, and shall not contain other elements. The functions `constructor` and `destructor` shall not be replaceable.
- The `constructor` function is called exactly once before the first use of the object. For each completely constructed object, the `destructor` is called exactly once, after the last use of the object, even if an error occurs. The `constructor` shall have exactly one output argument in which the constructed instance derived from `ExternalObject` is returned. The `destructor` shall have no output arguments and the only input argument of the `destructor` shall be of the type derived from `ExternalObject`. It is not legal to call explicitly the `constructor` and `destructor` functions. The `constructor` shall initialize the object, and must not require any other calls to be made for the initialization to be complete (e.g., from an initial algorithm or initial equation). The `destructor` shall delete the object, and must not require any other calls to be made for the deletion to be complete (e.g., from a `when terminal()` clause). The `constructor` shall not assume that pointers sent to the external object will remain valid for the life-time of the external object. An exception is that if the pointer to another external object is given as argument to the `constructor`, that pointer will remain valid as long as the other external object lives.

[External objects may be a protected component (or part of one) in a function. The `constructor` is in that case called at the start of the function call, and the `destructor` when the function returns, or when recovering from errors in the function.]

[External objects may be an input (or part of an input) to a function, in that case the `destructor` is not called (since the external object is active before and after the function call). Normally this is an external function, but it could be a non-external function as well (e.g. calling external functions one or more times). The function input shall not have a default value using the `constructor`.]

- Classes derived from `ExternalObject` can neither be used in an `extends`-clause nor in a short class definition.

- Only the constructor may return external objects and an external object can only be bound in component declarations and neither modified later nor assigned to.

[It follows that a function cannot return a component containing an external object, since only the constructor may return an external object and the constructor exactly returns the external object.]

- External functions may be defined which operate on the internal memory of an `ExternalObject`. An `ExternalObject` used as input argument or return value of an external C function is mapped to the C type `void*`.

[Example: A user-defined table may be defined in the following way as an `ExternalObject` (the table is read in a user-defined format from file and has memory for the last used table interval):

```
class MyTable
  extends ExternalObject;
  function constructor
    input String fileName = "";
    input String tableName = "";
    output MyTable table;
  external "C"
    table = initMyTable(fileName, tableName);
  end constructor;

  function destructor "Release storage of table"
    input MyTable table;
  external "C"
    closeMyTable(table);
  end destructor;
end MyTable;
```

and used in the following way:

```
model test "Define a new table and interpolate in it"
  MyTable table=MyTable(fileName = "testTables.txt",
    tableName="table1"); // call initMyTable
  Real y;
equation
  y = interpolateMyTable(table, time);
end test;
```

This requires to provide the following Modelica function:

```
function interpolateMyTable "Interpolate in table"
  input MyTable table;
  input Real u;
  output Real y;
  external "C"
    y = interpolateMyTable(table, u);
  end interpolateTable;
```

The external C functions may be defined in the following way:

```
typedef struct { /* User-defined datastructure of the table */
  double* array; /* nrow*ncolumn vector */
  int nrow; /* number of rows */
  int ncol; /* number of columns */
  int type; /* interpolation type */
  int lastIndex; /* last row index for search */
} MyTable;

void* initMyTable(const char* fileName, const char* tableName) {
  MyTable* table = malloc(sizeof(MyTable));
  if ( table == NULL ) ModelicaError("Not enough memory");
  // read table from file and store all data in *table
  return (void*) table;
```

```
};  
  
void closeMyTable(void* object) { /* Release table storage */  
  MyTable* table = (MyTable*) object;  
  if ( object == NULL ) return;  
  free(table->array);  
  free(table);  
}  
  
double interpolateMyTable(void* object, double u) {  
  MyTable* table = (MyTable*) object;  
  double y;  
  // Interpolate using "table" data (compute y)  
  return y;  
};
```

]

Chapter 13

Packages

Packages in Modelica may contain definitions of constants and classes including all kinds of specialized classes, functions, and subpackages. By the term subpackage we mean that the package is declared inside another package, no inheritance relationship is implied. Parameters and variables cannot be declared in a package. The definitions in a package should typically be related in some way, which is the main reason they are placed in a particular package. Packages are useful for a number of reasons:

- Definitions that are related to some particular topic are typically grouped into a package. This makes those definitions easier to find and the code more understandable.
- Packages provide encapsulation and coarse-grained structuring that reduces the complexity of large systems. An important example is the use of packages for construction of (hierarchical) class libraries.
- Name conflicts between definitions in different packages are eliminated since the package name is implicitly prefixed to names of definitions declared in a package.
- Information hiding and encapsulation can be supported to some extent by declaring `protected` classes, types, and other definitions that are available only inside the package and therefore inaccessible to outside code.
- Modelica defines a method for locating a package by providing a standard mapping of package names to storage places, typically file or directory locations in the file system.

13.1 Package as Specialized Class

The package concept is a specialized class (section 4.6), using the keyword `package`.

13.2 Importing Definitions from a Package

The import-clause makes public classes and other public definitions declared in some package available for use by shorter names in a class or a package. It is the only way of referring to definitions declared in some other package for use inside an encapsulated package or class.

[Import-clauses in a package or class fill the following two needs:

- *Making definitions from other packages available for use (by shorter names) in a package or class.*
- *Explicit declaration of usage dependences on other packages.*

]

An `import`-clause can occur in one of the following syntactic forms:

`import definitionname;` (qualified import of top-level definition)

`import packagename.definitionname;` (qualified import)

```

import packagename.{def1, def2, ..., defn}; (multiple definition import)
import packagename.*; (unqualified import)
import shortname = definitionname; (renaming import of top-level definition)
import shortname = packagename.definitionname; (renaming import)

```

Here *packagename* is the fully qualified name of the imported package including possible dot notation and *definitionname* is the name of an element in a package. The multiple definition import is equivalent to multiple single definition imports with corresponding *packagename* and definition names.

13.2.1 Lookup of Imported Names

This section only defines how the imported name is looked up in the import-clause. For lookup in general – including how import-clauses are used – see section 5.3.

Lookup of the name of an imported package or class deviates from the normal lexical lookup. For example, consider `A.B.C` in the clauses `import A.B.C;`, `import D = A.B.C;`, or `import A.B.C.*;`. Here, lookup starts with the lexical lookup of the first part of the name (`A`) at the top-level.

Qualified import-clauses may only refer to packages or elements of packages, i.e., in `import A.B.C;` or `import D = A.B.C;`, `A.B` must be a package. Unqualified import-clauses may only import from packages, i.e., in `import A.B.*;`, `A.B` must be a package.

[In `import A;` the class `A` can be any class which is an element of the unnamed top-level package.]

[For example, if the package `ComplexNumbers` would have been declared as a subpackage inside the package `Modelica.Math`, its fully qualified name would be `Modelica.Math.ComplexNumbers`. *definitionname* is the simple name without dot notation of a single definition that is imported. A *shortname* is a simple name without dot notation that can be used to refer to the package after import instead of the presumably much longer *packagename*.

The forms of `import` are exemplified below assuming that we want to access the addition operation of the hypothetical package `Modelica.Math.ComplexNumbers`:

```

import Modelica.Math.ComplexNumbers;           // Accessed by ComplexNumbers.Add
import Modelica.Math.ComplexNumbers.Add;       // Accessed by Add
import Modelica.Math.ComplexNumbers.{Add,Sub}; // Accessed by Add and Sub
import Modelica.Math.ComplexNumbers.*;        // Accessed by Add
import Co = Modelica.Math.ComplexNumbers;      // Accessed by Co.Add

```

]

13.2.2 Summary of Rules for Import Clauses

The following rules apply to import-clauses:

- Import-clauses are *not* inherited.
- Import-clauses are not named elements of a class or package. This means that import-clauses cannot be changed by modifiers or redeclarations.
- The *order* of import-clauses does not matter.
- One can only import *from* packages, not from other kinds of classes. Both packages and classes can be imported *into* i.e., they may contain import-clauses.
- An imported package or definition should always be referred to by its fully qualified name in the import-clause.
- Multiple qualified import-clauses shall not have the same import name.

13.3 The Modelica Library Path – MODELICAPATH

The top-level scope implicitly contains a number of classes stored externally. If a top-level name is not found at global scope, a Modelica translator shall look up additional classes in an ordered list of library roots, called MODELICAPATH.

[The implementation of MODELICAPATH is tool dependent. In order that a user can work in parallel with different Modelica tools, it is advisable to not have this list as environment variable, but as a setting in the respective tool. Since MODELICAPATH is tool dependent, it is not specified in which way the list of library roots is stored. Typically, on a Windows system MODELICAPATH is a string with path names separated by ‘;’ whereas on a Linux system it is a string with path names separated by a ‘.’.]

In addition a tool may define an internal list of libraries, since it is in general not advisable for a program installation to modify global environment variables. The version information for a library (as defined in section 18.8) may also be used during this search to search for a specific version of the library (e.g. if Modelica library version 2.2 is needed and the first directory in MODELICAPATH contain Modelica library version 2.1, whereas the second directory contains Modelica version 2.2, then Modelica library version 2.2 is loaded from the second directory.).

[The first part of the path A.B.C (i.e., A) is located by searching the ordered list of roots in MODELICAPATH. If no root contains A the lookup fails. If A has been found in one of the roots, the rest of the path is located in A; if that fails, the entire lookup fails without searching for A in any of the remaining roots in MODELICAPATH.]

13.3.1 Example of Searching MODELICAPATH

If during lookup a top-level name is not found in the unnamed top-level scope, the search continues in the package hierarchies stored in these directories.

[Example: Figure 13.1 below shows an example MODELICAPATH = "C:\library;C:\lib1;C:\lib2", with three directories containing the roots of the package hierarchies Modelica, MyLib, and ComplexNumbers. The first two are represented as the subdirectories C:\library\MyLib and C:\lib1\MyLib, whereas the third is stored as the file C:\lib2\ComplexNumbers.mo.




Figure 13.1: Roots of package hierarchies, e.g., Modelica, MyLib, and ComplexNumbers in MODELICAPATH = "C:\library;C:\lib1;C:\lib2".

Assume that we want to access the package MyLib.Pack2 in figure 13.1 above, e.g. through an `import`-clause `import MyLib.Pack2;`. During lookup we first try to find a package MyLib corresponding to the first part of the import name. It is not found in the top-level scope since it has not previously been loaded into the environment.

Since the name was not found in the top-level scope the search continues in the directories in the MODELICAPATH in the specified order. For the search to succeed, there must be a subdirectory MyLib or a file MyLib.mo in one of the directories mentioned in the MODELICAPATH. If there is no such subdirectory or file, the lookup fails. If MyLib is found in one of the directories, the rest of the name, in this case Pack2, is located in MyLib. If that fails, the entire lookup fails without continuing the search in possibly remaining directories.

In this example the name matches the subdirectory named MyLib in the second directory C:\lib1 mentioned in the MODELICAPATH. This subdirectory must have a file `package.mo` containing a definition of the

package MyLib, according to the Modelica rules on how to map a package hierarchy to the file system. The subpackage Pack2 is stored in its own subdirectory or file in the subdirectory MyLib. In this case the search succeeds and the package MyLib.Pack2 is loaded into the environment.]

13.4 Mapping Package/Class Structures to a Hierarchical File System

Packages/classes may be represented in the hierarchical structure of the operating system (the file system). For classes with version information see also section 18.8.3. The nature of such an external entity falls into one of the following two groups:

- Directory in the file system.
- File in the file system.

Each Modelica file in the file-system is stored in UTF-8 format (defined by The Unicode Consortium; <http://www.unicode.org>). A deprecated feature is that the file may start with the UTF-8 encoded BOM (byte order mark; 0xef 0xbb 0xbf); this is treated as white-space in the grammar. Since the use of BOM is deprecated, tools can ignore any BOM when reading, and it is recommended to never write it.

[Tools may also store classes in data-base systems, but that is not standardized.]

13.4.1 Mapping a Package/Class Hierarchy into a Directory Hierarchy (Structured Entity)

A directory shall contain a node, the file `package.mo`. The node shall contain a stored-definition that defines a class `A` with a name matching the name of the structured entity.

[The node typically contains documentation and graphical information for a package, but may also contain additional elements of the class A.]

A directory may also contain one or more sub-entities (directories or files). The sub-entities are mapped as elements of the class defined by their enclosing structured entity. Two sub-entities shall not define classes with identical names

[Example: If directory A contains the three files `package.mo`, `B.mo` and `C.mo`, the classes defined are A, A.B, and A.C.]

[Example: A directory shall not contain both the sub-directory A and the file `A.mo`.]

In order to preserve the order of sub-entities it is advisable to create a file `package.order` where each line contains the name of one class or constant (using its Modelica `IDENT` form). If a `package.order` is present when reading a structured entity the classes and constants are added in this order; if the contents does not exactly match the classes and constants in the package, the resulting order is tool specific and a warning may be given. Classes and constants that are stored in `package.mo` are also present in `package.order` but their relative order should be identical to the one in `package.mo` (this ensures that the relative order between classes and constants stored in different ways is preserved).

13.4.2 Mapping a Package/Class Hierarchy into a Single File (Nonstructured Entity)

When mapping a package or class-hierarchy to a file (e.g. the file `A.mo`), that file shall only define a single class `A` with a name matching the name of the nonstructured entity. In a file hierarchy the files shall have the extension `.mo`.

A `.mo` file defining more than one class cannot be part of the mapping to file-structure and it is an error if it is loaded from the `MODELICAPATH`.

13.4.3 The within Clause

A within-clause has the following syntax:


```
within [ packageprefixname ] ";"
```

A non-top-level entity shall begin with a within-clause which for the class defined in the entity specifies the location in the Modelica class hierarchy. A top-level class may contain a within-clause with no name. For a sub-entity of an enclosing structured entity, the within-clause shall designate the class of the enclosing entity; and this class must exist and must not have been defined using a short class definition.

[*Example: The subpackage `Rotational` declared within `Modelica.Mechanics` has the fully qualified name `Modelica.Mechanics.Rotational`, which is formed by concatenating the `packageprefixname` with the short name of the package. The declaration of `Rotational` could be given as below:*

```
within Modelica.Mechanics;
package Rotational // Modelica.Mechanics.Rotational
...
]
```

13.5 External resources

In order to reference external resources from documentation (such as links and images in html-text) and/or to reference images in the `Bitmap` annotation (see section 18.6.5.6). URIs should be used, for example `file:///` and the URI scheme `modelica://` which can be used to retrieve resources associated with a package. According to the URI specification scheme names are case-insensitive, but the lower-case form should be used, that is `Modelica://` is allowed but `modelica://` is the recommended form.

The Modelica-scheme has the ability to reference a hierarchical structure of resources associated with packages. The same structure is used for all kind of resource references, independent of use (external file, image in documentation, bitmap in icon layer, and link to external file in the documentation), and regardless of the storage mechanism.

Any Modelica-scheme URI containing a slash after the package-name is interpreted as a reference to a resource. The *authority* portion of the URI is interpreted as a fully qualified package name and the *path* portion of the URI is interpreted as the path (relative to the package) of the resource. Each storage scheme can define its own interpretation of the path (but care should be taken when converting from one storage scheme or when restructuring packages that resource references resolve to the same resource). Any storage scheme should be constrained such that a resource with a given path should be unique for any package name that precedes it. The first part of the path shall not be the name of a class in the package given by the authority.

When Modelica packages are stored hierarchically in a file-system (i.e. package `A` in a directory `A` containing `package.mo`) the resource `modelica://A/Resources/C.jpg` should be stored in the file `A/Resources/C.jpg`, it is not recommend to use `modelica://A.B/C.jpg` for referencing resources; it could be stored in the file `A/B/C.jpg` – which is counter-intuitive if `A.B` is stored together with `A`. When Modelica packages are stored in other formats a similar mapping should be defined, such that a resource with a given path should be unique for any package name that precedes it. The first part of the path shall not be the name of a class in the package given by the authority. As above for `Modelica 3.2.1/package.mo` i.e. resources starting from `Modelica 3.2.1`, and `modelica://Modelica.Mechanics/C.jpg` is `Modelica 3.2.1/Mechanics/C.jpg` – regardless of whether `Modelica.Mechanics` is stored in `Modelica 3.2.1/package.mo`, `Modelica 3.2.1/Mechanics/package.mo`, or `Modelica 3.2.1/Mechanics.mo`.

For a Modelica-package stored as a single file, `A.mo`, the resource `modelica://A/C.jpg` refers to a file `C.jpg` stored in the same directory as `A.mo`, but using resources in this variant is not recommended since multiple packages will share resources.

In case the name of the class contains quoted identifiers, the single-quote ‘`’` and any reserved characters (‘`:`’, ‘`/`’, ‘`?`’, ‘`\#`’, ‘`[`’, ‘`]`’, ‘`@`’, ‘`!`’, ‘`\$`’, ‘`\&`’, ‘`(`’, ‘`)`’, ‘`*`’, ‘`+`’, ‘`,`’, ‘`;`’, ‘`=`’) should be percent-encoded as normal in URIs.

[*Example: Consider a top-level package `Modelica` and a class `Mechanics` inside it, a reference such as `modelica://Modelica.Mechanics/C.jpg` is legal, while `modelica://Modelica/Mechanics/C.jpg` is illegal. The reference `modelica://Modelica.Mechanics/C.jpg` must also refer to a different resource than `modelica://Modelica/C.jpg`.]*

Chapter 14

Overloaded Operators

A Modelica `operator record` can overload the behavior for operations such as constructing, adding, multiplying etc.

The overloading is defined in such a way that ambiguities are not allowed and give an error. Furthermore, it is sufficient to define overloading for scalars. Overloaded array operations are automatically deduced from the overloaded scalar operations.

14.1 Overview of overloaded operators

In an `operator record` the definition of operations are done using the specialized class `operator` (a restricted class similar to `package`, see section 4.6) followed by the name of the operation. Each `operator` class is comprised of functions implementing different variants of the operation for the `operator record` class in which the definition resides.

- Overloaded constructors, see section 14.3:
'constructor', '0'
- Overloaded string conversions, see section 14.4:
'String'
- Overloaded binary operations, see section 14.5:
'+', '-' (subtraction), '*', '/', '^',
'==', '<=', '>', '<', '>=', '<=', 'and', 'or'
- Overloaded unary operations, see section 14.6:
'-' (negation), 'not'

The functions defined in the operator-class must take at least one component of the record class as input, except for the constructor-functions which instead must return one component of the record class. All of the functions shall return exactly one output.

The functions can be either called as defined in this section, or they can be called directly using the hierarchical name. The operator or operator function must be encapsulated; this allows direct calls of the functions and prohibits the functions from using the elements of operator record class.

The `operator record` may also contain additional functions, and declarations of components of the record. It is not legal to extend from an `operator record`, except as a short class definition modifying the default attributes for the component elements directly inside the operator record.

If an operator record was derived by a short class definition, the overloaded operators of this operator record are the operators that are defined in its base class, for subtyping see chapter 6.

The precedence and associativity of the overloaded operators is identical to the one defined in table 3.1 in section 3.2.

[*Note, the operator overloading as defined in this section is only a short hand notation for function calls.*]

14.2 Matching Function

All functions defined inside the `operator` class must return one output (based on the restriction above), and may include functions with optional arguments, i.e. functions of the form

```
function f
  input A1 u1;
  ...
  input Am um := Am;
  ...
  input An un;
  output B y;
algorithm
  ...
end f;
```

The vector P indicates whether argument m of f has a default value (`true` for default value, `false` otherwise). A call $f(A_1, a_2, \dots, a_k, b_1 = w_1, \dots, b_p = w_p)$ with distinct names b_j is a valid match for the function f , provided (treating `Integer` and `Real` as the same type)

- $A_i = \text{typeOf}(A_i)$ for $1 \leq i \leq k$,
- the names $b_j = u_{Q_j}$, $Q_j > k$, $A_{Q_j} = \text{typeOf}(w_i)$ for $1 \leq j \leq p$, and
- if the union of $\{i: 1 \leq i \leq k\}$, $\{Q_j: 1 \leq j \leq p\}$, and $\{m: P_m \text{ true and } 1 \leq m \leq n\}$ is the set $\{i: 1 \leq i \leq n\}$.

[This corresponds to the normal treatment of function calls with named arguments, requiring that all inputs have some value given by a positional argument, named argument, or a default value (and that positional and named arguments do not overlap). Note, that this only defines a valid call, but does not explicitly define the set of domains.]

14.3 Overloaded Constructors

Let C denote an operator record class and consider an expression $C(A_1, a_2, \dots, a_k, b_1=w_1, \dots, b_p=w_p)$.

1. If there exists a unique function f in C .`'constructor'` such that $(A_1, a_2, \dots, a_k, b_1=w_1, \dots, b_p=w_p)$ is a valid match for the function f , then $C(A_1, a_2, \dots, a_k, b_1=w_1, \dots, b_p=w_p)$ is resolved to C .`'constructor'`. $f(A_1, a_2, \dots, a_k, b_1=w_1, \dots, b_p=w_p)$.
2. If there is no operator C .`'constructor'` the automatically generated record constructor is called.
3. Otherwise the expression is erroneous.

Restrictions:

- The operator C .`'constructor'` shall only contain functions that declare one output component, which shall be of the operator record class C .
- For an operator recordclass there shall not exist any potential call that lead to multiple matches in item 1 above.

[How to verify this is not specified.]

- For a pair of operator record classes C and D and components c and d of these classes both of C .`'constructor'` (d) and D .`'constructor'` (c) shall not both be legal.

[Hence, one of the two definitions must be removed.]

[By the last restriction the following problem for binary operators is avoided:

Assume there are two operator record classes C and D that both have a constructor from `Real`. If we want to extend $c + c$ and $d + d$ to support mixed operations, one variant would be to define $c + d$ and $d + c$; but then $c + 2$ becomes ambiguous (since it is not clear which instance should be converted to). Without mixed operations expressions such as $c + d$ are only ambiguous if both conversion from C to D and back from D to C are both available, and this possibility is not allowed by the restriction above.]

Additionally there is an operator '0' defining the zero-value which can also be used to construct an element. The operator '0' for an operator record C can contain only one function, having zero inputs and one output of type C (the called function is therefore unambiguous). It should return the identity element of addition, and is used for generating flow-equations for connect-equations and zero elements for matrix multiplication.

14.4 Overloaded String Conversions

Consider an expression `String(A1, a2, ..., ak, b1=w1, ..., bp=wp)`, $k \geq 1$ where A_1 is an element of class A.

1. If A is a predefined type, i.e., `Boolean`, `Integer`, `Real`, `String` or an enumeration, or a type derived from them, then the corresponding built-in operation is performed.
2. If A is an operator record class and there exists a unique function f in `A.'String'` such that `A.'String'.f(A1, a2, ..., ak, b1=w1, ..., bp=wp)` is a valid match for f , then `String(A1, a2, ..., ak, b1=w1, ..., bp=wp)` is evaluated to `A.'String'.f(A1, a2, ..., ak, b1=w1, ..., bp=wp)`.
3. Otherwise the expression is erroneous.

Restrictions:

- The operator `A.'String'` shall only contain functions that declare one output component, which shall be of the `String` type, and the first input argument shall be of the operator record class A.
- For an operator record class there shall not exist any call that lead to multiple matches in (2) above.

[How to verify this is not specified.]

14.5 Overloaded Binary Operations

Let X denote a binary operator and consider an expression `a X b` where `a` is an instance or array of instances of class A and `b` is an instance or array of instances of class B.

1. If A and B are predefined types of such, then the corresponding built-in operation is performed.
2. Otherwise, if there exists *exactly one* function f in the union of `A.X` and `B.X` such that $f(\mathbf{a}, \mathbf{b})$ is a valid match for the function f , then `a X b` is evaluated using this function. It is an error, if multiple functions match. If A is not an operator record class, `A.X` is seen as the empty set, and similarly for B.

[Having a union of the operators ensures that if A and B are the same, each function only appears once.]

3. Otherwise, consider the set given by f in `A.X` and an operator record class C (different from B) with a constructor, g , such that `C.'constructor'.g(b)` is a valid match, and $f(\mathbf{a}, \mathbf{C.'constructor'.g(b)})$ is a valid match; and another set given by f in `B.X` and an operator record class D (different from A) with a constructor, h , such that `D.'constructor'.h(a)` is a valid match and $f(\mathbf{D.'constructor'.h(a)}, \mathbf{b})$ is a valid match. If the sum of the sizes of these sets is one this gives the unique match. If the sum of the sizes is larger than one it is an error.

[Informally, this means: If there is no direct match of `a X b`, then it is tried to find a direct match by automatic type casts of `a` or `b`, by converting either `a` or `b` to the needed type using an appropriate constructor function from one of the operator record classes used as arguments of the overloaded op functions. Example using the `Complex`-definition below:

```
Real a;
Complex b;
Complex c = a * b; // interpreted as:
// Complex.'*'.multiply(Complex.'constructor'.fromReal(a), b);
```

]

4. Otherwise, if **a** or **b** is an array expression, then the expression is conceptually evaluated according to the rules of section 10.6 with the following exceptions concerning section 10.6.4:

(a) *vector* * *vector* should be left undefined.

[The scalar product of table 10.9 does not generalize to the expected linear and conjugate linear scalar product of complex numbers.]

(b) *vector* * *matrix* should be left undefined.

[The corresponding definition of table 10.9 does not generalize to complex numbers in the expected way.]

(c) If the inner dimension for *matrix* * *vector* or *matrix* * *matrix* is zero, this uses the overloaded '0' operator of the result array element type. If the operator '0' is not defined for that class it is an error if the inner dimension is zero.

[For array multiplication it is assumed that the scalar elements form a non-commutative ring that does not necessarily have a multiplicative identity.]

5. Otherwise the expression is erroneous.

For an element-wise operator, **a** .op **b**, items 1, 4, and 5 are used; e.g. the operator .+ will always be defined in terms of '+'.

Restrictions:

- A function is allowed for a binary operator if and only if it has at least two inputs; at least one of which is of the operator record class, and the first two inputs shall not have default values, and all inputs after the first two must have default values.
- For an operator record class there shall not exist any (potential) call that lead to multiple matches in (2) above.

14.6 Overloaded Unary Operations

Let *X* denote a unary operator and consider an expression *X* **a** where **a** is an instance or array of instances of class **A**. Then *X* **a** is evaluated in the following way.

1. If **A** is a predefined type, then the corresponding built-in operation is performed.
2. If **A** is an operator record class and there exists a unique function *f* in **A**.*X* such that **A**.*X*.*f*(**a**) is a valid match, then *X* **a** is evaluated to **A**.*X*.*f*(**a**). It is an error, if there are multiple valid matches.
3. Otherwise, if **a** is an array expression, then the expression is conceptually evaluated according to the rules of section 10.6.
4. Otherwise the expression is erroneous.

Restrictions:

- A function is allowed for a unary operator if and only if it has least one input; and the first input is of the record type (or suitable arrays of such) and does not have a default value, and all inputs after the first one must have default values.
- For an operator record class there shall not exist any (potential) call that lead to multiple matches in (2) above.
- A binary and/or unary operator-class may only contain functions that are allowed for this binary and/or unary operator-class; and in case of '-' it is the union of these sets, since it may define both a unary (negation) and binary (subtraction) operator.

14.7 Example of Overloading for Complex Numbers

[Example: The rules in the previous subsections are demonstrated at hand of a record class to work conveniently with complex numbers:

```

operator record Complex "Record defining a Complex number"
  Real re "Real part of complex number";
  Real im "Imaginary part of complex number";
  encapsulated operator 'constructor'
    import Complex;
    function fromReal
      input Real re;
      input Real im := 0;
      output Complex result(re=re, im=im);
    algorithm
      annotation(Inline=true);
    end fromReal;
  end 'constructor';

  encapsulated operator function '+' // short hand notation, see section 4.6
    import Complex;
    input Complex c1;
    input Complex c2;
    output Complex result "= c1 + c2";
  algorithm
    result := Complex(c1.re + c2.re, c1.im + c2.im);
    annotation(Inline=true);
  end '+';

  encapsulated operator '-'
    import Complex;
    function negate
      input Complex c;
      output Complex result "= - c";
    algorithm
      result := Complex(-c.re, -c.im);
      annotation(Inline=true);
    end negate;

    function subtract
      input Complex c1;
      input Complex c2;
      output Complex result "= c1 - c2";
    algorithm
      result := Complex(c1.re - c2.re, c1.im - c2.im);
      annotation(Inline=true);
    end subtract;
  end '-';

  encapsulated operator function '*'
    import Complex;
    input Complex c1;
    input Complex c2;
    output Complex result "= c1 * c2";
  algorithm
    result := Complex(c1.re*c2.re - c1.im*c2.im, c1.re*c2.im + c1.im*c2.re);
    annotation(Inline=true);
  end '*';

  encapsulated operator function '/'
    import Complex; input Complex c1;
    input Complex c2;
    output Complex result "= c1 / c2";
  algorithm
    result := Complex(( c1.re*c2.re + c1.im*c2.im)/(c2.re^2 +
      c2.im^2),
      (-c1.re*c2.im + c1.im*c2.re)/(c2.re^2 + c2.im^2));
    annotation(Inline=true);
  end '/';

```

```

end '/';

encapsulated operator function '=='
  import Complex;
  input Complex c1;
  input Complex c2;
  output Boolean result "= c1 == c2";
algorithm
  result := c1.re == c2.re and c1.im == c2.im;
  annotation(Inline=true);
end '==';

encapsulated operator function 'String'
  import Complex;
  input Complex c;
  input String name := "j" "Name of variable representing sqrt(-1) in the
string";
  input Integer significantDigits=6 "Number of significant digits to be shown
";
  output String s;
algorithm
  s := String(c.re, significantDigits=significantDigits);
  if c.im <> 0 then
    s := if c.im > 0 then s + " + "
  else s + " - ";
    s := s + String(abs(c.im), significantDigits=significantDigits) + name;
  end if;
end 'String';

encapsulated function j
  import Complex;
  output Complex c;
algorithm
  c := Complex(0,1);
  annotation(Inline=true);
end j;

encapsulated operator function '0'
  import Complex;
  output Complex c;
algorithm
  c := Complex(0,0);
  annotation(Inline=true);
end '0';
end Complex;

function eigenValues
  input Real A [:::];
  output Complex ev[size(A, 1)];
protected
  Integer nx=size(A, 1);
  Real eval[nx,2];
  Integer i;
algorithm
  eval := Modelica.Math.Matrices.eigenValues(A);
  for i in 1:nx loop
    ev[i] := Complex(eval[i, 1], eval[i, 2]);
  end for;
end eigenValues;

// Usage of Complex number above:
Complex j = Complex.j();
Complex c1 = 2 + 3*j;

```

```

Complex c2 = 3 + 4*j;
Complex c3 = c1 + c2;
Complex c4[:] = eigenValues([1,2; -3,4]);
algorithm
  Modelica.Utilities.Streams.print("c4 = " + String(c4));
  // results in output:
  // c4 = {2.5 + 1.93649j, 2.5 - 1.93649j}
    
```

How overloaded operators can be symbolically processed. Example:

```

Real a;
Complex b;
Complex c = a + b;
    
```

Due to inlining of functions, the equation for c is transformed to:

```

c = Complex.'+'.add(Complex.'constructor'.fromReal(a), b);
  = Complex.'+'.add(Complex(re=a,im=0), b)
  = Complex(re=a+b.re, im=b.im);
    
```

or

```

c.re = a + b.re;
c.im = b.im;
    
```

These equations can be symbolically processed as other equations.

Complex can be used in a connector:

```

operator record ComplexVoltage = Complex(re(unit="V"),im(unit="V"));
operator record ComplexCurrent = Complex(re(unit="A"),im(unit="A"));

connector ComplexPin
  ComplexVoltage v;
  flow ComplexCurrent i;
end ComplexPin;

ComplexPin p1,p2,p3;
equation
  connect(p1,p2);
  connect(p1,p3);
    
```

The two connect-equations result in the following connection equations:

```

p1.v = p2.v;
p1.v = p3.v;
p1.i + p2.i + p3.i = Complex.'0'();
// Complex.'+'(p1.i, Complex.'+'(p2.i, p3.i)) = Complex.'0'();
    
```

The restrictions on extends are intended to avoid combining two variants inheriting from the same operator record, but with possibly different operations; thus `ComplexVoltage` and `ComplexCurrent` still use the operations from `Complex`. The restriction that it is not legal to extend from any of its enclosing scopes implies that:

```

package A
  extends Icon; // Ok
  operator record B ... end B;
end A;

package A2
  extends A(...); // Not legal
end A2;

package A3 = A(...); // Not legal
    
```

]

Chapter 15

Stream Connectors

The two basic variable types in a connector – *potential* (or *across*) variable and *flow* (or *through*) variable – are not sufficient to describe in a numerically sound way the bi-directional flow of matter with convective transport of specific quantities, such as specific enthalpy and chemical composition. The values of these specific quantities are determined from the upstream side of the flow, i.e., they depend on the flow direction. When using across and through variables, the corresponding models would include nonlinear systems of equations with `Boolean` unknowns for the flow directions and singularities around zero flow. Such equation systems cannot be solved reliably in general. The model formulations can be simplified when formulating two different balance equations for the two possible flow directions. This is not possible with across and through variables though.

This fundamental problem is addressed in Modelica by introducing a third type of connector variable, called stream variable, declared with the prefix `stream`. A stream variable describes a quantity that is carried by a flow variable, i.e., a purely convective transport phenomenon. The value of the stream variable is the specific property inside the component close to the boundary, assuming that matter flows out of the component into the connection point. In other words, it is the value the carried quantity would have if the fluid was flowing out of the connector, irrespective of the actual flow direction.

The rationale of the definition and typical use cases are described in appendix D.

15.1 Definition of Stream Connectors

If at least one variable in a connector has the `stream` prefix, the connector is called *stream connector* and the corresponding variable is called *stream variable*. The following definitions hold:

- The `stream` prefix can only be used in a connector declaration.
- A stream connector must have exactly one scalar variable with the `flow` prefix.
[The idea is that all stream variables of a connector are associated with this flow variable.]
- For every outside connector (see section 9.1.2), one equation is generated for every variable with the `stream` prefix (to describe the propagation of the stream variable along a model hierarchy). For the exact definition, see the end of section 15.2.
- For inside connectors (see section 9.1.2), variables with the `stream` prefix do not lead to connection equations.
- Connection equations with stream variables are generated in a model when using `inStream` or `actualStream`, see section 15.2 and section 15.3.

[Example:

```
connector FluidPort
  replaceable package Medium =
    Modelica.Media.Interfaces.PartialMedium;
  Medium.AbsolutePressure p "Pressure in connection point";
  flow Medium.MassFlowRate m_flow "> 0, if flow into component";
```

```

stream Medium.SpecificEnthalpy h_outflow "h close to port if m_flow < 0";
stream Medium.MassFraction X_outflow[Medium.nX] "X close to port if m_flow <
0";
end FluidPort;
    
```

`FluidPort` is a stream connector, because some connector variables have the `stream` prefix. The `Medium` definition and the stream variables are associated with the only flow variable (`m_flow`) that defines a fluid stream. The `Medium` and the stream variables are transported with this flow variable. The stream variables `h_outflow` and `X_outflow` are the stream properties inside the component close to the boundary, when fluid flows out of the component into the connection point. The stream properties for the other flow direction can be inquired with the built-in `inStream`. The value of the stream variable corresponding to the actual flow direction can be inquired through the built-in `actualStream`, see section 15.3.]

15.2 Stream Operator inStream and Connection Equations

In combination with the stream variables of a connector, `inStream` is designed to describe in a numerically reliable way the bi-directional transport of specific quantities carried by a flow of matter.

`inStream(v)` is only allowed on stream variables `v` and is informally the value the stream variable has, assuming that the flow is from the connection point into the component. This value is computed from the stream connection equations of the flow variables and of the stream variables.

For the following definition it is assumed that N inside connectors $m_j.c$ ($j = 1, 2, \dots, N$) and M outside connectors c_k ($k = 1, 2, \dots, M$) belonging to the same connection set (see definition in section 9.1.2) are connected together and a stream variable `h_outflow` is associated with a flow variable `m_flow` in connector `c`.

```

connector FluidPort
...
flow Real m_flow "Flow of matter; m_flow > 0 if flow into component";
stream Real h_outflow "Specific variable in component if m_flow < 0"
end FluidPort

model FluidSystem
...
FluidComponent m1, m2, ..., mN;
FluidPort c1, c2, ..., cM;
equation
connect(m1.c, m2.c);
connect(m1.c, m3.c);
...
connect(m1.c, mN.c);
connect(m1.c, c1);
connect(m1.c, c2);
...
connect(m1.c, cM);
...
end FluidSystem;
    
```




Figure 15.1: Exemplary FluidSystem with $N = 3$ and $M = 2$.

[The connection set represents an infinitesimally small control volume, for which the stream connection

equations are equivalent to the conservation equations for mass and energy.]

With these prerequisites, the semantics of the expression `inStream(mi.c.h_outflow)` is given implicitly by defining an additional variable `h_mix_ini`, and by adding to the model the conservation equations for mass and energy corresponding to the infinitesimally small volume spanning the connection set. The connection equation for the flow variables has already been added to the system according to the connection semantics of flow variables defined in section 9.2.

```
// Standard connection equation for flow variables
0 = sum(mj.c.m_flow for j in 1:N) + sum(-ck.m_flow for k in 1:M);
```

Whenever `inStream` is applied to a stream variable of an inside connector, the balance equation of the transported property must be added under the assumption of flow going into the connector

```
// Implicit definition of inStream applied to inside connector i
0 =
  sum(mj.c.m_flow *
    (if mj.c.m_flow > 0 or j==i then h_mix_ini else mj.c.h_outflow)
    for j in 1:N) +
  sum(-ck.m_flow *
    (if -ck.m_flow > 0 then h_mix_ini else inStream(ck.h_outflow)
    for k in 1:M);
inStream(mi.c.h_outflow) = h_mix_ini;
```

Note that the result of `inStream(mi.c.h_outflow)` is different for each port i , because the assumption of flow entering the port is different for each of them.

Additional equations need to be generated for the stream variables of outside connectors.

```
// Additional connection equations for outside connectors
for q in 1:M loop
  0 =
    sum(mj.c.m_flow *
      (if mj.c.m_flow > 0 then h_mix_outq else mj.c.h_outflow)
      for j in 1:N) +
    sum(-ck.m_flow *
      (if -ck.m_flow > 0 or k==q then h_mix_outq else inStream(ck.h_outflow))
      for k in 1:M);
  cq.h_outflow = h_mix_outq;
end for;
```

Neglecting zero flow conditions, the solution of the above-defined stream connection equations for `inStream` values of inside connectors and outflow stream variables of outside connectors is (for a derivation, see appendix D):

```
inStream(mi.c.h_outflow) :=
  (sum(max(-mj.c.m_flow,0)*mj.c.h_outflow for j in cat(1, 1:i-1, i+1:N) +
    sum(max( ck.m_flow,0)*inStream(ck.h_outflow) for k in 1:M))
  /
  (sum(max(-mj.c.m_flow,0) for j in cat(1, 1:i-1, i+1:N) +
    sum(max( ck.m_flow ,0) for k in 1:M)));

// Additional equations to be generated for outside connectors q
for q in 1:M loop
  cq.h_outflow :=
    (sum(max(-mj.c.m_flow,0)*mj.c.h_outflow for j in 1:N) +
    sum(max( ck.m_flow,0)*inStream(ck.h_outflow) for k in cat(1, 1:q-1, q+1:M))
    /
    (sum(max(-mj.c.m_flow,0) for j in 1:N) +
    sum(max( ck.m_flow ,0) for k in cat(1, 1:q-1, q+1:M)));
end for;
```

[Note, that `inStream(ck.h_outflow)` is computed from the connection set that is present one hierarchical level above. At this higher level `ck.h_outflow` is no longer an outside connector, but an inside connector and then the formula from above for inside connectors can be used to compute it.]

If the argument of `inStream` is an array, the implicit equation system holds elementwise, i.e., `inStream` is vectorizable.

The stream connection equations have singularities and/or multiple solutions if one or more of the flow variables become zero. When all the flows are zero, a singularity is always present, so it is necessary to approximate the solution in an open neighbourhood of that point.

[For example, assume that $m_j.c.m_flow = c_k.m_flow = 0$, then all equations above are identically fulfilled and `inStream` can have any value.]

However, specific optimizations may be applied to avoid the regularization if the flow through one port is zero or non-negative, see appendix D. It is required that `inStream` is appropriately approximated when regularization is needed and the approximation must fulfill the following requirements:

1. `inStream(mi.c.h_outflow)` and `inStream(ck.h_outflow)` must be unique with respect to all values of the flow and stream variables in the connection set, and must have a continuous dependency on them.
2. Every solution of the implicit equation system above must fulfill the equation system identically (upto the usual numerical accuracy), provided the absolute value of every flow variable in the connection set is greater than a small value ($|m_i.c.m_flow| > \text{eps}$ for $1 \leq i \leq N$ and $|c_i.m_flow| > \text{eps}$ for $1 \leq i \leq M$).

[Based on the above requirements, the following implementation is recommended:

- $N = 1, M = 0$:

```
inStream(m1.c.h_outflow) = m1.c.h_outflow;
```

- $N = 2, M = 0$:

```
inStream(m1.c.h_outflow) = m2.c.h_outflow;
inStream(m2.c.h_outflow) = m1.c.h_outflow;
```

- $N = 1, M = 1$:

```
inStream(m1.c.h_outflow) = inStream(c1.h_outflow);
// Additional equation to be generated
c1.h_outflow = m1.c.h_outflow;
```

- $N = 0, M = 2$:

```
// Additional equation to be generated
c1.h_outflow = inStream(c2.h_outflow);
c2.h_outflow = inStream(c1.h_outflow);
```

- All other cases:

```
if m_j.c.m_flow.min >= 0 for all j = 1:N with j <> i and
   c_k.m_flow.max <= 0 for all k = 1:M
then
   inStream(m_i.c.h_outflow) = m_i.c.h_outflow;
else
   si = sum(max(-m_j.c.m_flow,0) for j in cat(1,1:i-1, i+1:N) +
            sum(max(c_k.m_flow,0) for k in 1:M);
   inStream(m_i.c.h_outflow) =
      (sum(positiveMax(-m_j.c.m_flow,s_i)*m_j.c.h_outflow)
       + sum(positiveMax(c_k.m_flow,s_i)*inStream(c_k.h_outflow)))/
      (sum(positiveMax(-m_j.c.m_flow,s_i))
       + sum(positiveMax(c_k.m_flow,s_i)))
       for j in 1:N and i <> j and m_j.c.m_flow.min < 0,
       for k in 1:M and c_k.m_flow.max > 0
// Additional equations to be generated
for q in 1:M loop
   if m_j.c.m_flow.min >= 0 for all j = 1:N and
      c_k.m_flow.max <= 0 for all k = 1:M and k <> q
   then
```

```

cq.h_outflow = 0;
else
sq = (sum(max(-mj.c.m_flow,0) for j in 1:N) +
        sum(max(ck.m_flow,0) for k in cat(1,1:q-1, q+1:M)));
cq.h_outflow = (sum(positiveMax(-mj.c.m_flow,sq)*mj.c.h_outflow) +
                sum(positiveMax(ck.m_flow,sq)* inStream(ck.h_outflow)))/
                (sum(positiveMax(-mj.c.m_flow,sq)) +
                 sum(positiveMax(ck.m_flow,sq)))
for j in 1:N and mj.c.m_flow.min < 0,
    for k in 1:M and k <> q and ck.m_flow.max > 0
end for;

```

The operator `positiveMax(-mj.c.m_flow, si)` should be such that:

- `positiveMax(-mj.c.m_flow, si) = -mj.c.m_flow` if $-m_j.c.m_flow > eps1_j \geq 0$, where $eps1_j$ are small flows, compared to typical problem-specific values,
- all denominators should be greater than $eps2 > 0$, where $eps2$ is also a small flow, compared to typical problem-specific values.

Trivial implementation of `positiveMax` guarantees continuity of `inStream`:

```
positiveMax(-mj.c.m_flow, si) = max(-mj.c.m_flow, eps1); // so si is not needed
```

More sophisticated implementation, with smooth approximation, applied only when all flows are small:

```

// Define a "small number" eps (nominal(v) is the nominal value of v, see
// section 4.8.6)
eps := relativeTolerance*min(nominal(mj.c.m_flow));

// Define a smooth curve, such that alpha(si >= eps)=1 and alpha(si < 0)=0
alpha := smooth(1, if si > eps then 1
                 else if si > 0 then (si/eps)^2*(3-2*si/eps)
                 else 0);

// Define function positiveMax(v,si) as a linear combination of max(v,0)
// and of eps along alpha
positiveMax((-mj.c.m_flow,si) := alpha*max(-mj.c.m_flow,0) + (1-alpha)*eps;

```

The derivation of this implementation is discussed in appendix D. Note that in the cases $N = 1$, $M = 0$ (unconnected port, physically corresponding to a plugged-up flange), and $N = 2$, $M = 0$ (one-to-one connection), the result of `inStream` is trivial and no non-linear equations are left in the model, despite the fact that the original definition equations are nonlinear.

The following properties hold for this implementation:

- `inStream` is continuous (and differentiable), provided that $m_j.c.h_outflow$, $m_j.c.m_flow$, $c_k.h_outflow$, and $c_k.m_flow$ are continuous and differentiable.
- A division by zero can no longer occur (since $\text{sum}(\text{positiveMax}(-m_j.c.m_flow, s_i)) \geq eps2 > 0$), so the result is always well-defined.
- The balance equations are exactly fulfilled if the denominator is not close to zero (since the exact formula is used, if $\text{sum}(\text{positiveMax}(-m_j.c.m_flow, s_i)) > eps$).
- If all flows are zero, $\text{inStream}(m_i.c.h_outflow) = \text{sum}(m_j.c.h_outflow \text{ for } j \neq i \text{ and } m_j.c.m_flow.min < 0) / N_p$, i.e., it is the mean value of all the N_p variables $m_j.c.h_outflow$, such that $j \neq i$ and $m_j.c.m_flow.min < 0$. This is a meaningful approximation, considering the physical diffusion effects that are relevant at small flow rates in a small connection volume (thermal conduction for enthalpy, mass diffusion for mass fractions).

The value of `relativeTolerance` should be larger than the relative tolerance of the nonlinear solver used to solve the implicit algebraic equations.

As a final remark, further symbolic simplifications could be carried out by taking into account equations that affect the flows in the connection set (i.e., equivalent to $m_j.c.m_flow = 0$, which then implies

$m_j.c.m_flow.min \geq 0$). This is interesting, e.g., in the case of a valve when the stem position is set identically to closed by its controller.]

15.3 Stream Operator actualStream

`actualStream` is provided for convenience, in order to return the actual value of the stream variable, depending on the actual flow direction. The only argument of this built-in operator needs to be a reference to a stream variable. The operator is vectorizable, in the case of vector arguments. For the following definition it is assumed that an (inside or outside) connector `c` contains a stream variable `h_outflow` which is associated with a flow variable `m_flow` in the same connector `c`:

```
actualStream(c.h_outflow) = if c.m_flow > 0 then inStream(c.h_outflow) else c.h_outflow;
```

[`actualStream` is typically used in two contexts:

```
der(U) = c.m_flow * actualStream(c.h_outflow); // (1) energy balance equation
h_c = actualStream(c.h); // (2) monitoring the enthalpy
at port c
```

In the case of equation (1), although `actualStream` is discontinuous, the product with the flow variable is not, because `actualStream` is discontinuous when the flow is zero by construction. Therefore, a tool might infer that the expression is `smooth(0, ...)` automatically, and decide whether or not to generate an event. If a user wants to avoid events entirely, he/she may enclose the right-hand side of (1) with `noEvent`.

Equations like (2) might be used for monitoring purposes (e.g. plots), in order to inspect what the actual enthalpy of the fluid flowing through a port is. In this case, the user will probably want to see the change due to flow reversal at the exact instant, so an event should be generated. If the user doesn't bother, then he/she should enclose the right-hand side of (2) with `noEvent`. Since the output of `actualStream` will be discontinuous, it should not be used by itself to model physical behaviour (e.g., to compute densities used in momentum balances) – `inStream` should be used for this purpose. `actualStream` should be used to model physical behaviour only when multiplied by the corresponding flow variable (like in the above energy balance equation), because this removes the discontinuity.]

Chapter 16

Synchronous Language Elements

This chapter defines synchronous behavior suited for implementation of control systems. The synchronous behavior relies on an additional kind of discrete-time variables and equations, as well as an additional kind of when-clause. The benefits of synchronous behavior is that it allows a model to define large sampled data systems in a safe way, so that the translator can provide good diagnostics in case of a modeling error.

The following small example shows the most important elements:




Figure 16.1: A continuous plant and a sampled data controller connected together with sample and (zero-order) hold elements.

- A periodic clock is defined with `Clock(3)`. The argument of `Clock` defines the sampling interval (for details see section 16.3).
- Clocked variables (such as `yd`, `xd`, `ud`) are associated uniquely with a clock and can only be directly accessed when the associated clock is active. Since all variables in a clocked equation must belong to the same clock, clocking errors can be detected at compile time. If variables from different clocks shall be used in an equation, explicit cast operators must be used, such as `sample` to convert from continuous-time to clocked discrete-time or `hold` to convert from clocked discrete-time to continuous-time.

- A continuous-time variable is sampled at a clock tick with `sample`. The operator returns the value of the continuous-time variable when the clock is active.
- When no argument is defined for `clock`, the clock is deduced by clock inference.
- For a `when`-clause with an associated clock, all equations inside the `when`-clause are clocked with the given clock. All equations on an associated clock are treated together and in the same way regardless of whether they are inside a `when`-clause or not. This means that automatic sampling and hold of variables inside the `when`-clause does not apply (explicit sampling and hold is required) and that general equations can be used in such when-clauses (this is not allowed for `when`-clauses with `Boolean` conditions, that require a variable reference on the left-hand side of an equation).
- The `when`-clause in the controller could also be removed and the controller could just be defined by the equations:

```
/* Discrete controller */
E * xd = A * previous(xd) + B * yd;
  ud = C * previous(xd) + D * yd;
```

- `previous(xd)` returns the value of `xd` at the previous clock tick. At the first sample instant, the start value of `xd` is returned.
- A discrete-time signal (such as `ud`) is converted to a continuous-time signal with `hold`.
- If a variable belongs to a particular clock, then all other equations where this variable is used, with the exception of as argument to certain special operators, belong also to this clock, as well as all variables that are used in these equations. This property is used for *clock inference* and allows defining an associated clock only at a few places (above only in the sampler, whereas in the discrete controller and the hold the sampling period is inferred).
- The approach in this chapter is based on the clock calculus and inference system proposed by Colaço and Pouzet (2003) and implemented in Lucid Synchrone version 2 and 3 (Pouzet 2006). However, the Modelica approach also uses multi-rate periodic clocks based on rational arithmetic introduced by Forget, Boniol, Lesens, and Pagetti (2008), as an extension of the Lucid Synchrone semantics. These approaches belong to the class of synchronous languages (Benveniste, Caspi, Edwards, Halbwachs, Le Guernic, and Simone 2003).

16.1 Rationale for Clocked Semantics

[Periodically sampled control systems could also be defined with standard when-clauses, see section 8.3.5, and the `sample` operator, see section 3.7.5. For example:

```
when sample(0, 3) then
  xd = A * pre(xd) + B * y;
  u  = C * pre(xd) + D * y;
end when;
```

Equations in a when-clause with a `Boolean` condition have the property that (a) variables on the left hand side of the equal sign are assigned a value when the when-condition becomes true and otherwise hold their value, (b) variables not assigned in the when-clause are directly accessed (= automatic `sample` semantics), and (c) the variables assigned in the when-clause can be directly accessed outside of the when-clause (= automatic `hold` semantics).

Using standard when-clauses works well for individual simple sampled blocks, but the synchronous approach using clocks and clocked equations provide the following benefits (especially for large sampled systems):

1. *Possibility to detect inconsistent sampling rate, since clock partitioning (see section 16.7), replaces the automatic `sample` and `hold` semantics. Examples:*
 - a. *If when-clauses in different blocks should belong to the same controller part, but by accident different when-conditions are given, then this is accepted (no error is detected).*
 - b. *If a sampled data library such as the `Modelica.LinearSystems2.Controller` library is used, at every block the sampling of the block has to be defined as integer multiple of a base sampling*

rate. If several blocks should belong to the same controller part, and different integer multiples are given, then the translator has to accept this (no error is detected).

Note: Clocked systems can mix different sampling rates in well-defined ways when needed.

2. Fewer initial conditions are needed, as only a subset of clocked variables need initial conditions – the clocked state variables (see section 16.4). For a standard when-clause all variables assigned in a when-clause must have an initial value because they might be used, before they are assigned a value the first time. As a result, all these variables are “discrete-time states” although in reality only a subset of them need an initial value.
3. More general equations can be used, compared to standard when-clauses that require a restricted form of equations where the left hand side has to be a variable, in order to identify the variables that are assigned in the when-clause. This restriction can be circumvented for standard when-clauses, but is absent for clocked equations and make it more convenient to define nonlinear control algorithms.
4. Clocked equations allow clock inference, meaning that the sampling need only be given once for a sub-system. For a standard when-clause the condition (sampling) must be explicitly propagated to all blocks, which is tedious and error prone for large systems.
5. Possible to use general continuous-time models in synchronous models (e.g. some advanced controllers use an inverse model of a plant in the feedforward path of the controller, see Thümmel, Looye, Kurze, Otter, and Bals (2005)). This powerful feature of Modelica to use a nonlinear plant model in a controller would require to export the continuous-time model with an embedded integration method and then import it in an environment where the rest of the controller is defined. With clocked equations, clocked controllers with continuous-time models can be directly defined in Modelica.
6. Clocked equations are straightforward to optimize because they are evaluated exactly once at each an event instant. In contrast a standard when-clause with `sample` conceptually requires several evaluations of the model (in some cases tools can optimize this to avoid unneeded evaluations). The problem for the standard when-clause is that after `v` is changed, `pre(v)` shall be updated and the model re-evaluated, since the equations could depend on `pre(v)`. For clocked equations this iteration can be omitted since `previous(v)` can only occur in the clocked equations that are only run the first event iterations.
7. Clocked subsystems using arithmetic blocks are straightforward to optimize. When a standard math-block (e.g. addition) is part of a clocked sub-system it is automatically clocked and only evaluated when the clocked equations trigger. For standard when-clauses one either needs a separate sampled math-block for each operation, or it will conceptually be evaluated all the time. However, tools may perform a similar optimization for standard when-clauses and it is only relevant in large sampled systems.

]

16.2 Definitions

In this section various terms are defined.

16.2.1 Clocks and Clocked Variables

In section 3.8.3 the term *discrete-time* Modelica expression and in section 3.8.4 the term *continuous-time* Modelica expression is defined. In this chapter, two additional kinds of discrete-time expressions/variables are defined that are associated to clocks and are therefore called *clocked discrete-time* expressions. The different kinds of discrete-time variables in Modelica are defined below.

Definition 16.1. Piecewise-constant variable. (See section 3.8.3.) Variables $m(t)$ of base type `Real`, `Integer`, `Boolean`, enumeration, and `String` that are *constant* inside each interval $t_i \leq t < t_{i+1}$ (i.e., piecewise constant continuous-time variables). In other words, $m(t)$ changes value only at events: $m(t) = m(t_i)$, for $t_i \leq t < t_{i+1}$. Such variables depend continuously on time and they are discrete-time variables. See figure 16.2. □




Figure 16.2: A piecewise-constant variable.

Definition 16.2. Clock variable. Clock variables $c(t_i)$ are of base type `Clock`. A clock is either defined by a constructor (such as `Clock(3)`) that defines when the clock ticks (is active) at a particular time instant, or it is defined with clock operators relatively to other clocks, see section 16.5.1. See figure 16.3. \square

[Example: Clock variables:

```
Clock c1 = Clock(...);
Clock c2 = c1;
Clock c3 = subSample(c2, 4);
```

]




Figure 16.3: A clock variable. The value of a clock variable is not defined – the plot marks only indicate when the clock is active.

Definition 16.3. Clocked variable. The elements of clocked variables $r(t_i)$ are of base type `Real`, `Integer`, `Boolean`, enumeration, `String` that are associated uniquely with a clock $c(t_i)$. A clocked variable can only be directly accessed at the event instant where the associated clock is active. A constant and a parameter can always be used at a place where a clocked variable is required.

At time instants where the associated clock is not active, the value of a clocked variable can be inquired by using an explicit cast operator, see below. In such a case `hold` semantics is used, in other words the value of the clocked variable from the last event instant is used. See figure 16.4. \square




Figure 16.4: A clocked variable. The `hold` extrapolation of the value at the last event instant is illustrated with dashed green lines.

16.2.2 Base-Clock and Sub-Clock Partitions

There are two kinds of *clock partitions*:

Definition 16.4. Base-clock partition. A base-clock partition identifies a set of equations and a set of variables which must be executed together in one task. Different base-clock partitions can be associated to separate tasks for asynchronous execution. \square

Definition 16.5. Sub-clock partition. A sub-clock partition identifies a subset of equations and a subset of variables of a base-clock partition which are partially synchronized with other sub-clock partitions of the same base-clock partition, i.e., synchronized when the ticks of the respective clocks are simultaneous. \square

16.2.3 Argument Restrictions (Component Expression)

The built-in operators (with function syntax) defined in the following sections have partially restrictions on their input arguments that are not present for Modelica functions. To define the restrictions, the following term is used.

Definition 16.6. Component expression. A component expression is a **component-reference** which is a valid expression, i.e., not referring to models or blocks with equations. In detail, it is an instance of a (a) base type, (b) derived type, (c) record, (d) an array of such an instance (a-c), (e) one or more elements of such an array (d) defined by index expressions which are parameter expressions (see below), or (f) an element of records.

[The essential features are that one or several values are associated with the instance, that start values can be defined on these values, and that no equations are associated with the instance. A component expression can be constant or can vary with time.] \square

In the following sections, when defining an operator with function calling syntax, there are some common restrictions being used for the input arguments (operands). For example, an input argument to the operator may be required to be a component expression (definition 16.6) or parameter expression (section 3.8). To emphasize that there are no such restrictions, an input argument may be said to be just an *expression*.

[The reason for restricting an input argument to be a component expression is that the start value of the input argument is returned before the first tick of the clock of the input argument and this is not possible for a general expression.]

The reason for restricting an input argument to be a parameter expression is that the value of the input argument needs to be evaluated during translation, in order that clock analysis can be performed during translation.]

[Example: The input argument to `previous` is restricted to be a component expression.]

```

Real u1;
Real u2[4];
Complex c;
Resistor R;
...
y1 = previous(u1);      // fine
y2 = previous(u2);     // fine
y3 = previous(u2[2]);  // fine
y4 = previous(c.im);   // fine
y5 = previous(2 * u);  // error (general expression, not component expression)
y6 = previous(R);     // error (component, not component expression)

```

]

[Example: The named argument `factor` of `subSample` is restricted to be a parameter expression.]

```

Real u;
parameter Real p=3;
...

```

```

y1 = subSample(u, factor = 3);           // fine (literal)
y2 = subSample(u, factor = 2 * p - 3); // fine (parameter expression)
y3 = subSample(u, factor = 3 * u);     // error (general expression)
    
```

]

Note that the operators defined in this chapter do not automatically vectorize, but some operate on arrays in a similar way.

16.3 Clock Constructors

The overloaded constructors listed below are available to generate clocks, and it is possible to call them with the specified named arguments, or with positional arguments (according to the order shown in the details after the table).

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>Clock()</code>	Inferred clock	Operator 16.1
<code>Clock(intervalCounter, resolution)</code>	Rational interval clock	Operator 16.2
<code>Clock(interval)</code>	Real interval clock	Operator 16.3
<code>Clock(condition, startInterval)</code>	Event clock	Operator 16.4
<code>Clock(c, solverMethod)</code>	Solver clock	Operator 16.5

Operator 16.1 Clock

`Clock()`

Inferred clock. The operator returns a clock that is inferred.

[*Example:*

```

when Clock() then // equations are on the same clock
  x = A * previous(x) + B * u;
  Modelica.Utilities.Streams.print
    ("clock ticks at = " + String(sample(time)));
end when;
    
```

Note, in most cases, the operator is not needed and equations could be written without a when-clause (but not in the example above, since the print statement is otherwise not associated to a clock). This style is useful if a modeler would clearly like to mark the equations that must belong to one clock (although a tool could figure this out as well, if the when-clause is not present).]

Operator 16.2 Clock

`Clock(intervalCounter=intervalCounter, resolution=resolution)`

Rational interval clock. The first input argument, *intervalCounter*, is a clocked component expression (definition 16.6) or a parameter expression of type `Integer` with `min = 0`. The optional second argument *resolution* (defaults to 1) is a parameter expression of type `Integer` with `min = 1` and `unit = "Hz"`. If *intervalCounter* is a parameter expression with value zero, the period of the clock is derived by clock inference, see section 16.7.5. The result is of base type `Clock` that ticks when `time` becomes t_{start} , $t_{\text{start}} + interval_1$, $t_{\text{start}} + interval_1 + interval_2$, ... The clock starts at the start of the simulation t_{start} or when the controller is switched on. At the start of the simulation, `previous(intervalCounter) = intervalCounter.start` and the clocks ticks the first time. At the first clock tick *intervalCounter* must be computed and the second clock tick is then triggered at $interval_1 = intervalCounter/resolution$. At the second clock tick at time $t_{\text{start}} + interval_1$, a new value for *intervalCounter* must be computed and the next clock tick is scheduled at $interval_2 = intervalCounter/resolution$, and so on. If interval is a parameter expression, the clock defines a periodic clock.

[*The given interval and time shift can be modified by using the `subSample`, `superSample`, `shiftSample` and `backSample` operators on the returned clock, see section 16.5.2.]*

[*Example:*

```

// first clock tick: previous(nextInterval) = 2
Integer nextInterval(start = 2);
Real y1(start = 0);
Real y2(start = 0);
equation
  when Clock(2, 1000) then
    // periodic clock that ticks at 0, 0.002, 0.004, ...
    y1 = previous(y1) + 1;
  end when;

  when Clock(nextInterval, 1000) then
    // interval clock that ticks at 0, 0.003, 0.007, 0.012, ...
    nextInterval = previous(nextInterval) + 1;
    y2 = previous(y2) + 1;
  end when;

```

]

Note that operator `interval(c)` of `Clock c = Clock(nextInterval, resolution)` returns: `previous(intervalCounter) / resolution` (in seconds)

Operator 16.3 Clock

`Clock(interval=interval)`

Real interval clock. The input argument, *interval*, is a clocked component expression (definition 16.6) or a parameter expression. The *interval* must be strictly positive ($interval > 0$) of type `Real` with `unit = "s"`. The result is of base type `Clock` that ticks when `time` becomes t_{start} , $t_{start} + interval_1$, $t_{start} + interval_1 + interval_2$, ... The clock starts at the start of the simulation t_{start} or when the controller is switched on. Here the next clock tick is scheduled at $interval_1 = previous(interval) = interval.start$. At the second clock tick at time $t_{start} + interval_1$, the next clock tick is scheduled at $interval_2 = previous(interval)$, and so on. If *interval* is a parameter expression, the clock defines a periodic clock.

[*Note, the clock is defined with `previous(interval)`. Therefore, for sorting the input argument is treated as known. The given interval and time shift can be modified by using the `subSample`, `superSample`, `shiftSample` and `backSample` operators on the returned clock, see section 16.5.2. There are restrictions where this operator can be used, see `Clock` expressions below.*]

Operator 16.4 Clock

`Clock(condition=condition, startInterval=startInterval)`

Event clock. The first input argument, *condition*, is a continuous-time expression of type `Boolean`. The optional *startInterval* argument (defaults to 0) is the value returned by `interval()` at the first tick of the clock, see section 16.9. The result is of base type `Clock` that ticks when `edge(condition)` becomes `true`.

[*This clock is used to trigger a clocked partition due to a state event, that is a zero-crossing of a `Real` variable, in a continuous-time partition or due to a hardware interrupt that is modeled as `Boolean` in the simulation model.*]

[*Example:*

```

Clock c = Clock(angle > 0, 0.1); // before first tick of c:
                                // interval(c) = 0.1

```

]

[*The implicitly given interval and time shift can be modified by using the `subsample`, `superSample`, `shiftSample` and `backSample` operators on the returned clock, see section 16.5.2, provided the base interval is not smaller than the implicitly given interval.*]

Operator 16.5 Clock

`Clock(c=c, solverMethod=solverMethod)`

Solver clock. The first input argument, *c*, is a clock and the operator returns this clock. The returned clock is associated with the second input argument *solverMethod* of type `String`. The meaning of *solverMethod* is defined in section 16.8.2. If *solverMethod* is the empty `String`, then this `Clock` construct does not associate an integrator with the returned clock.

[*Example:*

```

Clock c1 = Clock(1, 10);           // 100 ms, no solver
Clock c2 = Clock(c1, "ImplicitTrapezoid"); // 100 ms, ImplicitTrapezoid
      solver
Clock c3 = Clock(c2, "");         // 100 ms, no solver
]

```

Besides inferred clocks and solver clocks, one of the following mutually exclusive associations of clocks are possible in one base partition:

1. One or more rational interval clocks, provided they are consistent with each other, see section 16.7.5.

[*Example:* Assume `y = subSample(u)`, and `Clock(1, 10)` is associated with `u` and `Clock(2, 10)` is associated with `y`, then this is correct, but it would be an error if `y` is associated with a `Clock(1, 3)`.]

2. Exactly one real interval clock.

[*Example:* Assume `Clock c = Clock(2.5)`, then variables in the same base partition can be associated multiple times with `c` but not multiple times with `Clock(2.5)`.]

3. Exactly one event clock.

4. A default clock, if neither a real interval, nor a rational interval nor an event clock is associated with a base partition. In this case the default clock is associated with the fastest sub-clock partition.

[*Typically, a tool will use `Clock(1.0)` as a default clock and will raise a warning, that it selected a default clock.*]

Clock variables can be used in a restricted form of expressions. Generally, every expression switching between clock variables must have parameter variability (in order that clock analysis can be performed when translating a model). Thus subscripts on clock variables and conditions of if-then-else switching between clock variables must be parameter expressions, and there are similar restrictions for sub-clock conversion operators section 16.5.2. Otherwise, the following expressions are allowed:

- Declaring arrays of clocks.

[*Example:* `Clock c1[3] = {Clock(1), Clock(2), Clock(3)}`]

- Array constructors of clocks: `{}`, `[]`, `cat`.

- Array access of clocks.

[*Example:* `sample(u, c1[2])`]

- Equality of clocks.

[*Example:* `c1 = c2`]

- If-expressions of clocks in equations.

[*Example:*

```

Clock c2 =
  if f > 0 then
    subSample(c1, f)
  elseif f < 0 then
    superSample(c1, f)
  else
    c1;
]

```

- Clock variables can be declared in models, blocks, connectors, and records. A clock variable can be declared with the prefixes `input`, `output`, `inner`, `outer`, but *not* with the prefixes `flow`, `stream`, `discrete`, `parameter`, or `constant`.

[Example:

```
connector ClockInput = input Clock;
```

]

16.4 Clocked State Variables

The previous value of a clocked variable can be accessed with the `previous` operator, listed below.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>previous(u)</code>	Previous value of clocked state variable	Operator 16.6

A variable to which `previous` has been applied is called a *clocked state variable*.

Operator 16.6 previous

```
previous(u)
```

The input argument u is a component expression (definition 16.6) or a parameter expression. The return argument has the same type as the input argument. Input and return arguments are on the same clock. At the first tick of the clock of u or after a reset transition (see section 17.3.2), the start value of u is returned, see section 16.9. At subsequent activations of the clock of u , the value of u from the previous clock activation is returned.

16.5 Partitioning Operators

A set of *clock conversion operators* together act as boundaries between different clock partitions.

16.5.1 Base-clock conversion operators

The operators listed below convert between a continuous-time and a clocked-time representation and vice versa.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>sample(u, clock)</code>	Sample continuous-time expression	Operator 16.7
<code>hold(u)</code>	Zeroth order hold of clocked-time variable	Operator 16.7

Operator 16.7 sample

```
sample(u, clock)
```

Input argument u is a continuous-time expression according to section 3.8.4. The optional input argument $clock$ is of type `Clock`, and can in a call be given as a named argument (with the name $clock$), or as positional argument. The operator returns a clocked variable that has $clock$ as associated clock and has the value of the left limit of u when $clock$ is active (that is the value of u just before the event of c is triggered). If argument $clock$ is not provided, it is inferred, see section 16.7.5.

[Since the operator returns the left limit of u , it introduces an infinitesimal small delay between the continuous-time and the clocked partition. This corresponds to the reality, where a sampled data system cannot act infinitely fast and even for a very idealized simulation, an infinitesimal small delay is present. The consequences for the sorting are discussed below.

Input argument u can be a general expression, because the argument is continuous-time and therefore has always a value. It can also be a constant, a parameter or a piecewise constant expression.

Note that `sample` is an overloaded function: If `sample` has two positional input arguments and the second argument is of type `Real`, it is the operator from section 3.7.5. If `sample` has one

input argument, or it has two input arguments and the second argument is of type `Clock`, it is the base-clock conversion operator from this section.]

Operator 16.8 hold

`hold(u)`

Input argument u is a clocked component expression (definition 16.6) or a parameter expression. The operator returns a piecewise constant signal of the same type as u . When the clock of u ticks, the operator returns u and otherwise returns the value of u from the last clock activation. Before the first clock activation of u , the operator returns the start value of u , see section 16.9.

[Since the input argument is not defined before the first tick of the clock of u , the restriction is present, that it must be a component expression (or a parameter expression), in order that the initial value of u can be used in such a case.]

[Example: Assume there is the following model:

```

Real y(start = 1), yc;
equation
  der(y) + y = 2;
  yc = sample(y, Clock(0.1));
initial equation
  der(y) = 0;

```

The value of `yc` at the first clock tick is `yc = 2` (and not `yc = 1`). The reason is that the continuous-time model `der(y) + y = 2` is first initialized and after initialization y has the value 2. At the first clock tick at time = 0, the left limit of y is 2 and therefore `yc = 2`.]

16.5.1.1 Sorting of a simulation model

[Since `sample(u)` returns the left limit of u , and the left limit of u is a known value, all inputs to a base-clock partition are treated as known during sorting. Since a periodic and interval clock can tick at most once at a time instant, and since the left limit of a variable does not change during event iteration (i.e., re-evaluating a base-clock partition associated with a condition clock always gives the same result because the `sample(u)` inputs do not change and therefore need not to be re-evaluated) all base-clock partitions, see section 16.7.3, need not to be sorted with respect to each other. Instead, at an event instant, active base-clock partitions can be evaluated first (and once) in any order. Afterwards, the continuous-time partition is evaluated.

Event iteration takes place only over the continuous-time partition. In such a scenario, accessing the left limit of u in `sample(u)` just means to pick the latest available value of u when the partition is entered, storing it in a local variable of the partition and only using this local copy during evaluation of the equations in this partition.]

16.5.2 Sub-clock conversion operators

The operators listed below convert between synchronous clocks.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>subSample(u, factor)</code>	Clock that is slower by a factor	Operator 16.9
<code>superSample(u, factor)</code>	Clock that is faster by a factor	Operator 16.10
<code>shiftSample(u, shiftCounter, resolution)</code>	Clock with time-shifted ticks	Operator 16.11
<code>backSample(u, backCounter, resolution)</code>	Inverse of <code>shiftSample</code>	Operator 16.12
<code>noClock(u)</code>	Clock that is always inferred	Operator 16.13

These operators have the following properties:

- The input argument u is a clocked expression or an expression of type `Clock`. (The operators can operate on all types of clocks.) If u is a clocked expression, the operator returns a clocked variable that has the same type as the expression. If u is an expression of type `Clock`, the operator returns a `Clock` – except for `noClock` where it is an error.
- The optional input arguments `factor` (defaults to 0, with `min = 0`), and `resolution` (defaults to 1, with `min = 1`) are parameter expressions of type `Integer`.

- Calls of the operators can use named arguments for the multi-letter arguments (i.e. not for u) with the given names, or positional arguments.

[Named arguments can make the calls easier to understand.]

- The input arguments `shiftCounter` and `backCounter` are parameter expressions of type `Integer` with `min = 0`.

Operator 16.9 subSample

```
subSample(u, factor=factor)
```

The clock of $y = \text{subSample}(u, \text{factor})$ is factor times slower than the clock of u . At every factor ticks of the clock of u , the operator returns the value of u . The first activation of the clock of y coincides with the first activation of the clock of u , and then every activation of the clock of y coincides with the every factor -th activation of the clock of u . If argument factor is not provided or is equal to zero, it is inferred, see section 16.7.5.

Operator 16.10 superSample

```
superSample(u, factor=factor)
```

The clock of $y = \text{superSample}(u, \text{factor})$ is factor times faster than the clock of u . At every tick of the clock of y , the operator returns the value of u from the last tick of the clock of u . The first activation of the clock of y coincides with the first activation of the clock of u , and then the interval between activations of the clock of u is split equidistantly into factor activations, such that the activation $1 + k \cdot \text{factor}$ of y coincides with the $1 + k$ activation of u .

[Thus $\text{subSample}(\text{superSample}(u, \text{factor}), \text{factor}) = u$.]

If argument factor is not provided or is equal to zero, it is inferred, see section 16.7.5. If an event clock is associated to a base-clock partition, all its sub-clock partitions must have resulting clocks that are sub-sampled with an `Integer` factor with respect to this base clock.

[Example:

```
Clock u = Clock(x > 0);
Clock y1 = subSample(u, 4);
Clock y2 = superSample(y1, 2); // fine; y2 = subSample(u, 2)
Clock y3 = superSample(u, 2); // error
Clock y4 = superSample(y1, 5); // error
```

]

Operator 16.11 shiftSample

```
shiftSample(u, shiftCounter=k, resolution=resolution)
```

The operator $c = \text{shiftSample}(u, k, \text{resolution})$ splits the interval between ticks of u into resolution equidistant intervals i . The clock c then ticks k intervals i after each tick of u .

It leads to

```
shiftSample(u, k, resolution) =
  subSample(shiftSample(superSample(u, resolution), k), resolution)
```

[Note, due to the restriction of `superSample` on event clocks, `shiftSample` can only shift the number of ticks of the event clock, but cannot introduce new ticks. Example:

```
// Rational interval clock
Clock u = Clock(3, 10); // ticks: 0, 3/10, 6/10, ..
Clock y1 = shiftSample(u, 1, 3); // ticks: 1/10, 4/10,
...
// Event clock
Clock u = Clock(sin(2 * pi * time) > 0, startInterval = 0.0);
// ticks: 0.0, 1.0, 2.0, 3.0, ...
Clock y1 = shiftSample(u, 2); // ticks: 2.0, 3.0, ...
Clock y2 = shiftSample(u, 2, 3); // error (resolution must be 1)
```

]

Operator 16.12 backSample

```
backSample(u, backCounter=cnt, resolution=res)
```

The input argument u is either a component expression (definition 16.6) or an expression of type `Clock`. This is an inverse of `shiftSample` such that `Clock y = backSample(u, cnt, res)` implicitly defines a clock y such that `shiftSample(y, cnt, res)` activates at the same times as u . It is an error if the clock of y starts before the base clock of u .

At every tick of the clock of y , the operator returns the value of u from the last tick of the clock of u . If u is a clocked component expression, the operator returns the start value of u , see section 16.9, before the first tick of the clock of u .

[Example:

```
// Rational interval clock 1
Clock u = Clock(3, 10);           // ticks: 0, 3/10, 6/10, ..
Clock y1 = shiftSample(u, 3);     // ticks: 9/10, 12/10, ..
Clock y2 = backSample(y1, 2);     // ticks: 3/10, 6/10,
...
Clock y3 = backSample(y1, 4);     // error (ticks before u)
Clock y4 = shiftSample(u, 2, 3);  // ticks: 2/10, 5/10,
...
Clock y5 = backSample(y4, 1, 3);  // ticks: 1/10, 4/10,
...
// Event clock
Clock u = Clock(sin(2 * pi * time) > 0, startInterval = xx)
// ticks: 0, 1.0, 2.0, 3.0, ....
Clock y1 = shiftSample(u, 3);     // ticks: 3.0, 4.0, ...
Clock y2 = backSample(y1, 2);     // ticks: 1.0, 2.0, ...
```

]

Operator 16.13 noClock

```
noClock(u)
```

The clock of $y = \text{noClock}(u)$ is always inferred, and u must be part of the same base-clock as y . At every tick of the clock of y , the operator returns the value of u from the last tick of the clock of u . If `noClock(u)` is called before the first tick of the clock of u , the start value of u is returned.

[Clarification of `backSample`:

Let a and b be positive integers with $a < b$, and

```
yb = backSample(u, a, b)
ys = shiftSample(u, b - a, b)
```

Then when ys exists, also yb exists and $ys = yb$.

The variable yb exists for the above parameterization with $a < b$ one clock tick before ys . Therefore, `backSample` is basically a `shiftSample` with a different parameterization and the clock of `backSample.y` ticks before the clock of u . Before the clock of u ticks, $yb = u.start$.]

[Clarification of `noClock` operator:

Note, that `noClock(u)` is not equivalent to `sample(hold(u))`. Consider the following model:

```
model NoClockVsSampleHold
  Clock clk1 = Clock(0.1);
  Clock clk2 = subSample(clk1, 2);
  Real x(start = 0), y(start = 0), z(start = 0);
equation
  when clk1 then
    x = previous(x) + 0.1;
  end when;
```

```

when clk2 then
  y = noClock(x);           // most recent value of x
  z = sample(hold(x));     // left limit of x (infinitesimally delayed)!
end when;
end NoClockVsSampleHold;

```

Due to the infinitesimal delay of `sample`, `z` will not show the current value of `x` as `clk2` ticks, but will show its previous value (left limit). However, `y` will show the current value, since it has no infinitesimal delay.]

Note that it is not legal to compute the derivative of the `sample`, `subSample`, `superSample`, `backSample`, `shiftSample`, and `noClock` operators.

16.6 Clocked When-Clause

In addition to the previously discussed conditional when-clause, a *clocked* when-clause is introduced:

```

when clock-expression then
  clocked-equation
  ...
end when;

```

The clocked when-clause cannot be nested and does not have any `elsewhen` part. It cannot be used inside an algorithm. General equations are allowed in a clocked when-clause.

For a clocked when-clause, all equations inside the when-clause are clocked with the same clock given by the *clock-expression*.

16.7 Clock Partitioning

This section defines how clock-partitions and clocks associated with equations are inferred.

[Typically clock partitioning is performed before sorting the equations. The benefit is that clocking and symbolic transformation errors are separated.]

Every clocked variable is uniquely associated with exactly one clock.

After model flattening, every equation in an equation section, every expression and every algorithm section is either continuous-time, or it is uniquely associated with exactly one clock. In the latter case it is called a *clocked equation*, a *clocked expression* or *clocked algorithm* section respectively. The associated clock is either explicitly defined by a when-clause, see section 16.5.2, or it is implicitly defined by the requirement that a clocked equation, a clocked expression and a clocked algorithm section must have the same clock as the variables used in them with exception of the expressions used as first arguments in the conversion operators of section 16.5. *Clock inference* means to infer the clock of a variable, an equation, an expression or an algorithm section if the clock is not explicitly defined and is deduced from the required properties in the previous two paragraphs.

All variables in an expression without clock conversion operators must have the same clock to infer the clocks for each variable and expression. The clock inference works both forward and backwards regarding the data flow and is also being able to handle algebraic loops. The clock inference method uses the set of variable incidences of the equations, i.e., what variables that appear in each equation.

Note that incidences of the first argument of clock conversion operators of section 16.5 are handled specially.

16.7.1 Flattening of Model

The clock partitioning is conceptually performed after model flattening, i.e., redeclarations have been elaborated, arrays of model components expanded into scalar model components, and overloading resolved. Furthermore, function calls to inline functions have been inlined.

[This is called conceptually, because a tool might do this more efficiently in a different way, provided the result is the same as if everything is flattened. For example, array and matrix equations and records don't need to be expanded if they have the same clock.]

Furthermore, each non-trivial expression (non-literal, non-constant, non-parameter, non-variable), $expr_i$, appearing as first argument of a clock conversion operator (except `hold` and `backSample`) is recursively replaced by a unique variable, v_i , and the equation $v_i = expr_i$ is added to the equation set.

16.7.2 Connected Components of the Equations and Variables Graph

Consider the set E of equations and the set V of unknown variables (not constants and parameters) in a flattened model, i.e., $M = \langle E, V \rangle$. The partitioning is described in terms of an undirected graph $\langle N, F \rangle$ with the nodes N being the set of equations and variables, $N = E \cup V$. The set $\text{incidence}(e)$ for an equation e in E is a subset of V , in general, the unknowns which lexically appear in e . There is an edge in F of the graph between an equation, e , and a variable, v , if $v \in \text{incidence}(e)$:

$$F = \{(e, v) : e \in E, v \in \text{incidence}(e)\}$$

A set of clock partitions is the *connected components* (Wikipedia, *Connected components*) of this graph with appropriate definition of the incidence operator.

16.7.3 Base-clock Partitioning

The goal is to identify all clocked equations and variables that should be executed together in the same task, as well as to identify the continuous-time partition.

The base-clock partitioning is performed with base-clock inference which uses the following incidence definition:

$\text{incidence}(e) =$
 the *unknown* variables, as well as variables x in `der(x)`, `pre(x)`, and `previous(x)`, which lexically appear in e
 except as first argument of base-clock conversion operators: `sample` and `hold` and `Clock(condition=..., startInterval=...)`.

The resulting set of connected components, is the partitioning of the equations and variables, $B_i = \langle E_i, V_i \rangle$, according to base-clocks and continuous-time partitions.

The base clock partitions are identified as *clocked* or as *continuous-time partitions* according to the following properties:

A variable u in `sample(u)`, a variable y in `y = hold(ud)`, and a variable b in `Clock(b, startInterval=...)` where the `Boolean` b is in a continuous-time partition.

Correspondingly, variables u and y in `y = sample(uc)`, `y = subSample(u)`, `y = superSample(u)`, `y = shiftSample(u)`, `y = backSample(u)`, `y = previous(u)`, are in a clocked partition. Equations in a clocked when clause are also in a clocked partition. Other partitions where none of the variables in the partition are associated with any of the operators above have an unspecified partition kind and are considered continuous-time partitions.

All continuous-time partitions are collected together and form *the continuous-time partition*.

[Example:

```

// Controller 1
ud1 = sample(y, c1);
0 = f1(yd1, ud1, previous(yd1));

// Controller 2
ud2 = superSample(yd1, 2);
0 = f2(yd2, ud2);

// Continuous-time system
u = hold(yd2);
```

```

0 = f3(der(x1), x1, u);
0 = f4(der(x2), x2, x1);
0 = f5(der(x3), x3);
0 = f6(y, x1, u);
    
```

After base clock partitioning, the following partitions are identified:

```

// Base partition 1 — clocked partition
ud1 = sample(y, c1);           // incidence(e) = {ud1}
0 = f1(yd1, ud1, previous(ud1)); // incidence(e) = {yd1, ud1}
ud2 = superSample(yd1, 2);     // incidence(e) = {ud2, yd1}
0 = f2(yd2, ud2);             // incidence(e) = {yd2, ud2}

// Base partition 2 — continuous-time partition
u = hold(yd2);                // incidence(e) = {u}
0 = f3(der(x1), x1, u);       // incidence(e) = {x1, u}
0 = f4(der(x2), x2, x1);     // incidence(e) = {x2, x1}
0 = f6(y, x1, u);            // incidence(e) = {y, x1, u}

// Identified as separate partition, but belonging to partition 2
0 = f5(der(x3), x3);         // incidence(e) = {x3}
    
```

]

16.7.4 Sub-clock Partitioning

For each clocked partition B_i , identified in section 16.7.3, the sub-clock partitioning is performed with sub-clock inference which uses the following incidence definition:

$\text{incidence}(e) =$

the *unknown* variables, as well as variables x in $\text{der}(x)$, $\text{pre}(x)$, and $\text{previous}(x)$, which lexically appear in e
 except as first argument of sub-clock conversion operators: `subSample`, `superSample`, `shiftSample`, `backSample`, `noClock`, and `Clock` with first argument of `Boolean` type.

The resulting set of connected components, is the partitioning of the equations and variables, $S_{ij} = \langle E_{ij}, V_{ij} \rangle$, according to sub-clocks.

The resulting sets of equations and variables shall be possible to solve separately, meaning that systems of equations cannot involve different sub-clocks.

It can be noted that:

$$\begin{aligned}
 E_{ij} \cap E_{kl} &= \emptyset, \forall i \neq k, j \neq l \\
 V_{ij} \cap V_{kl} &= \emptyset, \forall i \neq k, j \neq l \\
 V &= \bigcup V_{ij} \\
 E &= \bigcup E_{ij}
 \end{aligned}$$

[Example: After sub-clock partitioning of the example from section 16.7.3, the following partitions are identified:

```

// Base partition 1 (clocked partition)
// Sub-clock partition 1.1
ud1 = sample(y, c1);           // incidence(e) = {ud1}
0 = f1(yd1, ud1, previous(yd1)); // incidence(e) = {yd1, ud1}

// Sub-Clock partition 1.2
ud2 = superSample(yd1, 2);     // incidence(e) = {ud2}
0 = f2(yd2, ud2);             // incidence(e) = {yd2, ud2}

// Base partition 2 (no sub-clock partitioning, since continuous-time)
u = hold(yd2);
0 = f3(der(x1), x1, u);
    
```

```

0 = f4(der(x2), x2, x1);
0 = f5(der(x3), x3);
0 = f6(y, x1, u);

```

]

16.7.5 Sub-clock Inferencing

For each base-clock partition, the base interval needs to be determined and for each sub-clock partition, the sub-sampling factors and shift need to be determined. The sub-clock partition intervals are constrained by `subSample` and `superSample` factors which might be known (or parameter expression) or unspecified, as well as by `shiftSample`, `shiftCounter` and `resolution`, or `backSample`, `backCounter` and `resolution`. This constraint set is used to solve for all intervals and sub-sampling factors and shift of the sub-clock partitions. The model is erroneous if no solution exist.

[It must be possible to determine that the constraint set is valid at compile time. However, in certain cases, it could be possible to defer providing actual numbers until run-time.]

It is required that accumulated sub- and supersampling factors in the range of 1 to 2^{63} can be handled.

[64 bit internal representation of numerator and denominator with sign can be used and gives minimum resolution 1.08×10^{-19} seconds and maximum range 9.22×10^{18} seconds = 2.92×10^{11} years.]

16.8 Continuous-Time Equations in Clocked Partitions

[The goal is that every continuous-time Modelica model can be utilized in a sampled data control system. This is achieved by solving the continuous-time equations with a defined integration method between clock ticks. With this feature, it is for example possible to invert the nonlinear dynamic model of a plant, see Thümmel, Looye, Kurze, Otter, and Bals (2005), and use it in a feedforward path of an advanced control system that is associated with a clock.]

This feature also allows defining multi-rate systems: Different parts of the continuous-time model are associated to different clocks and are solved with different integration methods between clock ticks, e.g., a very fast sub-system with an implicit solver with a small step-size and a slow sub-system with an explicit solver with a large step-size.]

With the language elements defined in this section, continuous-time equations can be used in clocked partitions. Hereby, the continuous-time equations are solved with the defined integration method between clock ticks.

From the view of the continuous-time partition, the clock ticks are not interpreted as events, but as step-sizes of the integrator that the integrator must exactly hit. Hence, no event handling is triggered at clock ticks (provided an explicit event is not triggered from the model at this time instant).

[The interpretation of the clock ticks is the same assumption as for manually discretized controllers, such as the z-transform.]

[It is not defined, how events are handled that appear when solving the continuous-time partition. For example, a tool could handle events exactly in the same way as for a usual simulation. Alternatively, relations might be interpreted literally, so that events are no longer triggered (in order that the time for an integration step is always the same, as needed for hard real-time requirements).]

From the view of the clocked partition, the continuous-time partition is discretized and the discretized continuous-time variables have only a value at a clock tick. Therefore, such a partition is handled in the same way as any other clocked partition. Especially, operators such as `sample`, `hold`, `subSample` must be used to communicate signals of the discretized continuous-time partition with other partitions. Hereby, a discretized continuous-time partition is seen as a clocked partition.

16.8.1 Clocked Discrete-Time and Clocked Discretized Continuous-Time Partition

Additionally to the variability of expressions defined in section 3.8, an orthogonal concept *clocked variability* is defined in this section. If not explicitly stated otherwise, an expression with a variability such

as *continuous-time* or *discrete-time* means that the expression is inside a partition that is not associated to a clock. If an expression is present in a partition that is not a continuous-time partition, it is a *clocked expression* and has *clocked variability*.

After sub-clock inferencing, see section 16.7.5, every partition that is associated to a clock has to be categorized as *clocked discrete-time* or *clocked discretized continuous-time* partition.

If a clocked partition contains no operator `der`, `delay`, `spatialDistribution`, no event related operators from section 3.7.5 (with exception of `noEvent`), and no `when`-clause with a `Boolean` condition, it is a *clocked discrete-time* partition.

[That is, the clocked discrete-time partition is a standard sampled data system that is described by difference equations.]

If a clocked partition is not a *clocked discrete-time* partition, it is a *clocked discretized continuous-time* partition. Such a partition has to be solved with a *solver method* of section 16.8.2. When `previous(x)` is used on a continuous-time state variable `x`, then `previous(x)` uses the start value of `x` as value for the first clock tick.

In a clocked discrete-time partition all event generating mechanisms do no longer apply. Especially neither relations, nor any of the built-in operators of section 3.7.2 (event triggering mathematical functions) will trigger an event.

16.8.2 Solver Methods

The integration method associated with a clocked discretized continuous-time partition is defined with a string. A predefined type `ModelicaServices.Types.SolverMethod` defines the methods supported by the respective tool by using the `choices` annotation.

[The `ModelicaServices` package contains tool specific definitions. A string is used instead of an enumeration, since different tools might have different values and then the integer mapping of an enumeration is misleading since the same value might characterize different integrators.]

The following names of solver methods are standardized:

```

type SolverMethod = String annotation(choices(
  choice="External" "Solver specified externally",
  choice="ExplicitEuler" "Explicit Euler method (order 1)",
  choice="ExplicitMidPoint2" "Explicit mid point rule (order 2)",
  choice="ExplicitRungeKutta4" "Explicit Runge-Kutta method (order 4)",
  choice="ImplicitEuler" "Implicit Euler method (order 1)",
  choice="ImplicitTrapezoid" "Implicit trapezoid rule (order 2)"
)) "Type of integration method to solve differential equations in a clocked " +
  "discretized continuous-time partition."

```

If a tool supports one of the integrators of `SolverMethod`, it must use the solver method name of above.

[A tool may also support other integrators. Typically, a tool supports at least methods `"External"` and `"ExplicitEuler"`. If a tool does not support the integration method defined in a model, typically a warning message is printed and the method is changed to `"External"`.]

If the solver method is `"External"`, then the partition associated with this method is integrated by the simulation environment for an interval of length of `interval()` using a solution method defined in the simulation environment.

[An example of such a solution method could be to have a table of the clocks that are associated with discretized continuous-time partitions and a method selection per clock. In such a case, the solution method might be a variable step solver with step-size control that integrates between two clock ticks. The simulation environment might also combine all partitions associated with method `"External"`, as well as all continuous-time partitions, and integrate them together with the solver selected by the simulation environment.]

If the solver method is *not* `"External"`, then the partition is integrated using the given method with the step-size `interval()`.

[For a periodic clock, the integration is thus performed with fixed step size.]

The solvers are defined with respect to the underlying ordinary differential equation in state space form to which the continuous-time partition can be transformed, at least conceptually (t is time, $u_c(t)$ is the continuous-time `Real` vector of input variables, $u_d(t)$ is the discrete-time `Real/Integer/Boolean/String` vector of input variables, $x(t)$ is the continuous-time real vector of states, and $y(t)$ is the continuous-time or discrete-time `Real/Integer/Boolean/String` vector of algebraic and/or output variables):

$$\begin{aligned}\dot{x} &= f(x, u, t) \\ y &= g(x, u, t)\end{aligned}$$

A solver method is applied to a subclock partition. Such a partition has explicit inputs u marked by `sample(u)`, `subSample(u)`, `superSample(u)`, `shiftSample(u)` and/or `backSample(u)`. Furthermore, the outputs y of such a partition are marked by `hold(y)`, `subSample(y)`, `superSample(y)`, `shiftSample(y)`, and/or `backSample(y)`. The arguments of these operators are to be used as input signals u and output signals y in the conceptual ordinary differential equation above, and in the discretization formulae below, respectively.

The solver methods (with exception of "External") are defined by integrating from clock tick t_{i-1} to clock tick t_i and computing the desired variables at t_i , with $h = t_i - t_{i-1} = \text{interval}(u)$ and $x_i = x(t_i)$ (for all methods: $y_i = g(x_i, u_{c,i}, u_{d,i}, t_i)$):

SolverMethod	Solution method
"ExplicitEuler"	$x_i := x_{i-1} + h \cdot \dot{x}_{i-1}$ $\dot{x}_i := f(x_i, u_{c,i}, u_{d,i}, t_i)$
"ExplicitMidPoint2"	$x_i := x_{i-1} + h \cdot f(x_{i-1} + \frac{1}{2} \cdot h \cdot \dot{x}_{i-1}, \frac{u_{c,i-1} + u_{c,i}}{2}, u_{d,i-1}, t_{i-1} + \frac{1}{2} \cdot h)$ $\dot{x}_i := f(x_i, u_{c,i}, u_{d,i}, t_i)$
"ExplicitRungeKutta4"	$k_1 := h \cdot \dot{x}_{i-1}$ $k_2 := h \cdot f(x_{i-1} + \frac{1}{2}k_1, \frac{u_{c,i-1} + u_{c,i}}{2}, u_{d,i-1}, t_{i-1} + \frac{1}{2} \cdot h)$ $k_3 := h \cdot f(x_{i-1} + \frac{1}{2}k_2, \frac{u_{c,i-1} + u_{c,i}}{2}, u_{d,i-1}, t_{i-1} + \frac{1}{2} \cdot h)$ $k_4 := h \cdot f(x_{i-1} + k_3, u_{c,i}, u_{d,i}, t_i)$ $x_i := x_{i-1} + \frac{1}{6} \cdot (k_1 + 2 \cdot k_2 + 2 \cdot k_3 + k_4)$ $\dot{x}_i := f(x_i, u_{c,i}, u_{d,i}, t_i)$
"ImplicitEuler"	Equation system with unknowns: x_i, \dot{x}_i $x_i = x_{i-1} + h \cdot \dot{x}_i$ $\dot{x}_i = f(x_i, u_{c,i}, u_{d,i}, t_i)$
"ImplicitTrapezoid"	Equation system with unknowns: x_i, \dot{x}_i $x_i = x_{i-1} + \frac{1}{2}h \cdot (\dot{x}_i + \dot{x}_{i-1})$ $\dot{x}_i = f(x_i, u_{c,i}, u_{d,i}, t_i)$

The initial conditions will be used at the first tick of the clock, and the first integration step will go from the first to the second tick of the clock.

[Example: Assume the differential equation

```
input Real u;
Real x(start = 1, fixed = true);
equation
  der(x) = -x + u;
```

shall be transformed to a clocked discretized continuous-time partition with the "ExplicitEuler" method. The following model is a manual implementation:

```
input Real u;
parameter Real x_start = 1;
Real x(start = x_start); // previous(x) = x_start at first clock tick
Real der_x(start = 0); // previous(der_x) = 0 at first clock tick
protected
  Boolean first(start = true);
equation
```



```

when Clock() then
  first = false;
  if previous(first) then
    // first clock tick (initialize system)
    x = previous (x);
  else
    // second and further clock tick
    x = previous(x) + interval() * previous(der_x);
  end if;
  der_x = -x + u;
end when;

```

]

[For the implicit integration methods the efficiency can be enhanced by utilizing the discretization formula during the symbolic transformation of the equations. For example, linear differential equations are then mapped to linear and not non-linear algebraic equation systems, and also the structure of the equations can be utilized. For details see *Elmqvist, Otter, and Cellier (1995)*. It might be necessary to associate additional data for an implicit integration method, e.g. the relative tolerance to solve the non-linear algebraic equation systems, or the maximum number of iterations in case of hard realtime requirements. This data is tool specific and is typically either defined with a vendor annotation or is given in the simulation environment.]

16.8.3 Associating a Solver to a Partition

A `SolverMethod` can be associated to a clock with the overloaded `Clock` constructor `Clock(c, solverMethod = ...)`, see section 16.3. If a clock is associated with a clocked partition and a `SolverMethod` is associated with this clock, then the partition is integrated with it.

[Example:

```

// Continuous PI controller in a clocked partition
vd = sample(x2, Clock(Clock(1, 10), solverMethod="ImplicitEuler"));
e = ref - vd;
der(xd) = e / Ti;
u = k * (e + xd);

// Physical model
f = hold(u);
der(x1) = x2;
m * der(x2) = f;

```

]

16.8.4 Inferencing of solverMethod

If a `solverMethod` is not explicitly associated with a partition, it is inferred with a similar mechanism as for sub-clock inferencing, see section 16.7.5.

First, one set is constructed for each sub-clock partition, containing just this sub-clock partition. These sets are then merged as follows: For each set without a specified `solverMethod`, the set is merged with sets connected to it (these may contain a `solverMethod`), and this is repeated until it is not possible to merge more sets. The sets connected in this way should be part of the same base-clock partition and connected through a sub-clock conversion operator (`subSample`, `superSample`, `shiftSample`, `backSample`, or `noClock`).

- It is an error if this set contains multiple different values for `solverMethod`.
- If the set contains continuous time-equations:
 - It is an error if this set contains no `solverMethod`.
 - Otherwise, the specified `solverMethod` is used.
- If the set does not contain continuous time-equations, there is no need for a `solverMethod`.

[Example:

```

model InferenceTest
  Real x(start = 3) "Explicitly using ExplicitEuler";
  Real y "Explicitly using ImplicitEuler method";
  Real z "Inferred to use ExplicitEuler";
equation
  der(x) = -x + sample(1, Clock(Clock(1, 10), solverMethod="ExplicitEuler"));
  der(y) = subSample(x, 2) +
           sample(1, Clock(Clock(2, 10), solverMethod="ImplicitEuler"));
  der(z) = subSample(x, 2) + 1;
end InferenceTest;

model IllegalInference
  Real x(start = 3) "Explicitly using ExplicitEuler";
  Real y "Explicitly using ImplicitEuler method";
  Real z;
equation
  der(x) = -x + sample(1, Clock(Clock(1, 10), solverMethod="ExplicitEuler"));
  der(y) = subSample(x, 2) +
           sample(1, Clock(Clock(2, 10), solverMethod="ImplicitEuler"));
  der(z) = subSample(x, 4) + 1 + subSample(y);
end IllegalInference;

```

Here *z* is a continuous-time equation connected directly to both *x* and *y* partitions that have different solverMethod.]

16.9 Initialization of Clocked Partitions

The standard scheme for initialization of Modelica models does not apply for clocked discrete-time partitions. Instead, initialization is performed in the following way:

- Clocked discrete-time variables cannot be used in initial equation or initial algorithm sections.
- Attribute `fixed` cannot be applied to clocked discrete-time variables. The attribute `fixed` is true for variables to which `previous` is applied, otherwise false.

16.10 Other Operators

A few additional utility operators are listed below.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>firstTick(u)</code>	Test for first clock tick	Operator 16.14
<code>interval(u)</code>	Interval between previous and present tick	Operator 16.15

It is an error if these operators are called in the continuous-time partition.

Operator 16.14 `firstTick`

```
firstTick(u)
```

This operator returns true at the first tick of the clock of the expression, in which this operator is called. The operator returns false at all subsequent ticks of the clock. The optional argument *u* is only used for clock inference, see section 16.7.

Operator 16.15 `interval`

```
interval(u)
```

This operator returns the interval between the previous and present tick of the clock of the expression, in which this operator is called. The optional argument *u* is only used for clock inference, see section 16.7. At the first tick of the clock the following is returned:

1. If the specified clock interval is a parameter expression, this value is returned.

2. Otherwise the start value of the variable specifying the interval is returned.
3. For an event clock the additional `startInterval` argument to the event clock constructor is returned.

The return value of `interval` is a scalar `Real` number.

[Example: A discrete PI controller is parameterized with the parameters of a continuous PI controller, in order that the discrete block is robust against changes in the sample period. This is achieved by discretizing a continuous PI controller (here with an implicit Euler method):

```

block ClockedPI
  parameter Real T "Time constant of continuous PI controller";
  parameter Real k "Gain of continuous PI controller";
  input Real u;
  output Real y;
  Real x(start = 0);
  protected
  Real Ts = interval(u);
equation
  /* Continuous PI equations: der(x) = u / T; y = k * (x + u);
   * Discretization equation: der(x) = (x - previous(x)) / Ts;
   */
  when Clock() then
    x = previous(x) + Ts / T * u;
    y = k * (x + u);
  end when;
end ClockedPI;
    
```

A continuous-time model is inverted, discretized and used as feedforward controller for a PI controller (`der`, `previous`, `interval` are used in the same partition):

```

block MixedController
  parameter Real T "Time constant of continuous PI controller";
  parameter Real k "Gain of continuous PI controller";
  input Real y_ref, y_meas;
  Real y;
  output Real yc;
  Real z(start = 0);
  Real xc(start = 1, fixed = true);
  Clock c = Clock(Clock(0.1), solverMethod="ImplicitEuler");
  protected
  Real uc;
  Real Ts = interval(uc);
equation
  /* Continuous-time, inverse model */
  uc = sample(y_ref, c);
  der(xc) = uc;
  /* PI controller */
  z = if firstTick() then 0 else
  previous(z) + Ts / T * (uc - y_meas);
  y = xc + k * (xc + uc);
  yc = hold(y);
end MixedController;
    
```

]

16.11 Semantics

The execution of sub partitions requires exact time management for proper synchronization. The implication is that testing a `Real`-valued time variable to determine sampling instants is not possible. One possible method is to use counters to handle sub-sampling scheduling,

```
Clock_i_j_ticks =
```

```

if pre(Clock_i_j_ticks) < subSamplingFactor_i_j then
  1 + pre(Clock_i_j_ticks)
else
  1;

```

and to test the counter to determine when the sub-clock is ticking:

```

Clock_i_j_activated =
  BaseClock_i_activated and Clock_i_j_ticks >= subSamplingFactor_i_j;

```

The `Clock_i_j_activated` flag is used as the guard for the sub partition equations.

[Consider the following example:

```

model ClockTicks
  Integer second = sample(1, Clock(1));
  Integer seconds(start = -1) = mod(previous(seconds) + second, 60);
  Integer milliSeconds(start = -1) =
    mod(previous(milliSeconds) + superSample(second, 1000), 1000);
  Integer minutes(start = -1) =
    mod(previous(minutes) + subSample(second, 60), 60);
end ClockTicks;

```

A possible implementation model is shown below using Modelica 3.2 semantics. The base-clock is determined to 0.001 seconds and the sub-sampling factors to 1000 and 60000.

```

model ClockTicksWithModelica32
  Integer second;
  Integer seconds(start = -1);
  Integer milliSeconds(start = -1);
  Integer minutes(start = -1);

  Boolean BaseClock_1_activated;
  Integer Clock_1_1_ticks(start = 59999);
  Integer Clock_1_2_ticks(start = 0);
  Integer Clock_1_3_ticks(start = 999);
  Boolean Clock_1_1_activated;
  Boolean Clock_1_2_activated;
  Boolean Clock_1_3_activated;
equation
  // Prepare clock tick
  BaseClock_1_activated = sample(0, 0.001);
  when BaseClock_1_activated then
    Clock_1_1_ticks =
      if pre(Clock_1_1_ticks) < 60000 then 1 + pre(Clock_1_1_ticks) else 1;
    Clock_1_2_ticks =
      if pre(Clock_1_2_ticks) < 1 then 1 + pre(Clock_1_2_ticks) else 1;
    Clock_1_3_ticks =
      if pre(Clock_1_3_ticks) < 1000 then 1 + pre(Clock_1_3_ticks) else 1;
  end when;
  Clock_1_1_activated = BaseClock_1_activated and Clock_1_1_ticks >= 60000;
  Clock_1_2_activated = BaseClock_1_activated and Clock_1_2_ticks >= 1;
  Clock_1_3_activated = BaseClock_1_activated and Clock_1_3_ticks >= 1000;

  // -----
  // Sub partition execution
  when {Clock_1_3_activated} then
    second = 1;
  end when;
  when {Clock_1_1_activated} then
    minutes = mod(pre(minutes) + second, 60);
  end when;
  when {Clock_1_2_activated} then
    milliSeconds = mod(pre(milliSeconds) + second, 1000);
  end when;

```

```
when {Clock_1_3_activated} then  
  seconds = mod(pre(seconds) + second, 60);  
end when;  
end ClockTicksWithModelica32;
```

]

Chapter 17

State Machines

This chapter defines language elements to define clocked state machines. These state machines have a similar modeling power as Statecharts (Harel 1987) and have the important feature that at one clock tick, there is only one assignment to every variable (for example, it is an error if state machines are executed in parallel and they assign to the same variable at the same clock tick; such errors are detected during translation). Furthermore, it is possible to activate and deactivate clocked equations and blocks at a clock tick. An efficient implementation will only evaluate the equations and blocks that are active at the current clock tick. With other Modelica language elements, this important feature cannot be defined.

The semantics of the state machines defined in this chapter is inspired by mode automata and is basically the one from Lucid Synchrone 3.0 (Pouzet 2006). Note, safety critical control software in aircrafts is often defined with such kind of state machines. The following properties are different to Lucid Synchrone 3.0:

- Lucid Synchrone has two kinds of transitions: *strong* and *weak* transitions. Strong transitions are executed before the actions of a state are evaluated and weak transitions are executed after the actions of a state are evaluated. This can lead to surprising behavior, because the actions of a state are skipped if it is activated by a weak transition and exited by a true strong transition.

For this reason, the state machines in this chapter use *immediate* (= the same as *strong*) and *delayed* transitions. Delayed transitions are *immediate* transitions where the condition is automatically delayed with an implicit `previous`.

- Parallel state machines can be explicitly synchronized with a language element (similarly as parallel branches in Sequential Function Charts). This often occurring operation can also be defined in Statecharts or in Lucid Synchrone state machines but only indirectly with appropriate conditions on transitions.
- Modelica blocks can be used as states. They might contain clocked or clocked discretized continuous-time equations (in the latter case, the equations are integrated between the previous and the next clock tick, if the corresponding state is active).

17.1 Transitions

Any Modelica block instance without continuous-time equations or algorithms can potentially be a state of a state machine. A cluster of instances which are coupled by `transition` statements makes a state machine. All parts of a state machine must have the same clock. All transitions leaving one state must have different priorities. One and only one instance in each state machine must be marked as initial by appearing in an `initialState` statement.

The special kinds of connect-equations listed below are used to define define a state machine.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>transition(from, to, condition, ...)</code>	State machine transition between states	Operator 17.1
<code>initialState(state)</code>	State machine initial state	Operator 17.2

The `transition` and `initialState` equations can only be used in equations, and cannot be used inside if-equations with conditions that are not parameter expressions, or in when-equations.

The operators listed below are used to query the status of the state machine.

<i>Expression</i>	<i>Description</i>	<i>Details</i>
<code>activeState(state)</code>	Predicate for active state	Operator 17.3
<code>ticksInState()</code>	Ticks since activation	Operator 17.4
<code>timeInState()</code>	Time since activation	Operator 17.5

Operator 17.1 transition

```
transition(from, to, condition,
           immediate=imm, reset=reset, synchronize=synch, priority=prio)
```

Arguments *from* and *to* are block instances, and *condition* is a `Boolean` argument. The optional arguments *immediate*, *reset*, and *synchronize* are of type `Boolean`, have parameter variability and a default of `true`, `true`, `false` respectively. The optional argument *priority* is of type `Integer`, has parameter variability and a default of 1.

This operator defines a transition from instance *from* to instance *to*. The *from* and *to* instances become states of a state machine. The transition fires when *condition* = `true` if *imm* = `true` (this is called an *immediate transition*) or `previous(condition)` when *imm* = `false` (this is called a *delayed transition*). Argument *priority* defines the priority of firing when several transitions could fire. In this case the transition with the smallest value of *priority* fires. It is required that $prio \geq 1$ and that for all transitions from the same state, the priorities are different. If *reset* = `true`, the states of the target state are reinitialized, i.e. state machines are restarted in initial state and state variables are reset to their start values. If *synch* = `true`, any transition is disabled until all state machines of the from-state have reached final states, i.e. states without outgoing transitions. For the precise details about firing a transition, see section 17.3.

Operator 17.2 initialState

```
initialState(state)
```

Argument *state* is the block instance that is defined to be the initial state of a state machine. At the first clock tick of the state machine, this state becomes active.

Operator 17.3 activeState

```
activeState(state)
```

Argument *state* is a block instance. The operator returns `true` if this instance is a state of a state machine and this state is active at the actual clock tick. If it is not active, the operator returns `false`.

It is an error if the instance is not a state of a state machine.

Operator 17.4 ticksInState

```
ticksInState()
```

Returns the number of ticks of the clock of the state machine for which the currently active state has maintained its active state without interruption, i.e. without local or hierarchical transitions from this state. In the case of a self-transition to the currently active state or to an active enclosing state, the number is reset to one.

This function can only be used in state machines.

Operator 17.5 timeInState

```
timeInState()
```

Returns the time duration as `Real` in [s] for which the currently active state has maintained its active state without interruption, i.e. without local or hierarchical transitions from this state. In the case of a self-transition to the currently active state or to an active enclosing state, the time is reset to zero.

This function can only be used in state machines.

[Example: If there is a transition with `immediate=false` from state A1 to A2 and the condition is `ticksInState() >= 5`, and A1 became active at 10ms, and the clock period is 1ms, then A1 will be active at 10ms, 11ms, 12ms, 13ms, 14ms, and will be not active at 15 ms.

```

block State end State;
State A0;
State A1; // Becomes active at 10ms
State A2;
equation
  initialState(A0);
  transition(A0, A1, sample(time, Clock(1, 1000)) > 0.0095);
  transition(A1, A2, ticksInState() >= 5, immediate=false);

```

17.2 State Machine Graphics

[Figure 17.1 shows the recommended layout of a state machine.]




Figure 17.1: Recommended layout of a simple state machine. For the 5 transitions, the settings are as follows, from left to right: `immediate = true, false, true, false, true`; `reset = true, true, false, false, true`; `synchronize = false, false, false, false, true`; `priority = 1, 2, 3, 4, 5`.

The annotation for graphics of `transition` has the following structure: `annotation(Line(...), Text(...))`; and for `initialState()`: `graphical-primitives(Line(...))`; with `Line` and `Text` annotations defined in chapter 18.

[Example:

```

transition(state2, state1, x < 10, immediate=true, reset=true, synchronize=
  false, priority=1)
  annotation (
    Line(
      points = {{-40, -16}, {-36, -4}, {-32, 8}, {-40, 26}, {-40, 32}, {-46, 50}},
      color = {175, 175, 175},
      thickness = 0.25,
      smooth = Smooth.Bezier),
    Text(
      string = "%condition",
      extent = {{4, -4}, {4, -10}},
      fontSize = 10,
      textStyle = {TextStyle.Bold},
      textColor = {95, 95, 95},
      horizontalAlignment = TextAlignment.Left),
  );

```

]

The `Text` annotation representing the transition condition can use the notation `%condition` to refer to the condition expression.

The extent of the `Text` is interpreted relative to either the first point of the `Line`, in the case of `immediate=false`, or the last point (`immediate=true`).

In addition to the line defined by the points of the `Line` annotation, a perpendicular line is used to represent the transition. This line is closer to the first point if `immediate=false` otherwise closer to the last point.

If the condition text is somewhat distant from the perpendicular line, a dimmed straight line joins the transition text and the perpendicular line. (See the rightmost transition above.)

If `reset=true`, a filled arrow head is used otherwise an open arrow head. For `synchronize=true`, an inverse “fork” symbol is used in the beginning of the arrow. (See the rightmost transition above.)

The value of the `priority` attribute is prefixing the condition text followed by a colon if `priority > 1`.

The `initialState` line has a filled arrow head and a bullet at the opposite end of the initial state (as shown above).

17.3 State Machine Semantics

For the purpose of defining the semantics of state machines, assume that the data of all transitions are stored in an array of records:

```
record Transition
  Integer from;
  Integer to;
  Boolean immediate = true;
  Boolean reset = true;
  Boolean synchronize = false;
  Integer priority = 1;
end Transition;
```

The transitions are sorted with lowest priority number last in the array; and the priorities must be unique for each value of `from`. The states are enumerated from 1 and up. The transition conditions are stored in a separate array `c[:]` since they are time varying.

The semantics model is a discrete-time system with inputs `{c[:], active, reset}` with `t` being an array corresponding to the inputs to the transition operator, outputs `{activeState, activeReset, activeResetStates[:]}` and states `{nextState, nextReset, nextResetStates[:]}`. For a top level state machine, `active` is always true. For sub-state machines, `active` is true only when the parent state is active. For a top level state machine, `reset` is true at the first activation only. For sub-state machine, `reset` is propagated from the state machines higher up.

17.3.1 State Activation

The state update starts from `nextState`, i.e., what has been determined to be the next state at the previous time. `selectedState` takes into account if a reset of the state machine is to be done.

```
output Integer selectedState = if reset then 1 else previous(nextState);
```

The integer fired is calculated as the index of the transition to be fired by checking that `selectedState` is the from-state and the condition is true for an immediate transition or `previous(condition)` is true for a delayed transition. The `max` function returns the index of the transition with highest priority or 0.

```
Integer fired =
  max(
    if (if t[i].from == selectedState
      then (if t[i].immediate then c[i] else previous(c[i]))
      else false)
    then i
    else 0
```

```
for i in 1:size(t,1);
```

The start value of `c` is false. This definition would require that the previous value is recorded for all transitions conditions. Below is described an equivalent semantics which just require to record the value of one integer variable delayed.

The integer `immediate` is calculated as the index of the immediate transition to potentially be fired by checking that `selectedState` is the from-state and the condition is true. The `max` function returns the index of the transition with true condition and highest priority or 0.

```
Integer immediate =
  max(
    if (if t[i].immediate and t[i].from == selectedState
        then c[i]
        else false)
    then i
    else 0
  for i in 1:size(t,1));
```

In a similar way, the `Integer delayed` is calculated as the index for a potentially delayed transition, i.e. a transition taking place at the next clock tick. In this case the from-state needs to be equal to `nextState`:

```
Integer delayed =
  max(
    if (if not t[i].immediate and t[i].from == nextState
        then c[i]
        else false)
    then i
    else 0
  for i in 1:size(t,1));
```

The transition to be fired is determined as follows, taking into account that a delayed transition might have higher priority than an immediate:

```
Integer fired = max(previous(delayed), immediate);
```

`nextState` is set to the found transitions to-state:

```
Integer nextState =
  if active then
    (if fired > 0
     then t[fired].to
     else selectedState)
  else
    previous(nextState);
```

In order to define synchronize transitions, each state machine must determine which are the final states, i.e. states without from-transitions and to determine if the state machine is in a final state currently:

```
Boolean finalStates[nStates] = {
  max(if t[j].from == i then 1 else 0 for j in 1:size(t,1)) == 0
  for i in 1:nStates};
Boolean stateMachineInFinalState = finalStates[activeState];
```

To enable a synchronize transition, all the `stateMachineInFinalState` conditions of all state machines within the meta state must be true. An example is given below in the semantic example model.

17.3.2 Reset Handling

A state can be reset for two reasons:

- The whole state machine has been reset from its context.
In this case, all states must be reset, and the initial state becomes active.

- A reset transition has been fired.
Then, its target state is reset, but not other states.

The first reset mechanism is handled by the `activeResetStates` and `nextResetStates` vectors.

The state machine reset flag is propagated and maintained to each state individually:

```
output Boolean activeResetStates[nStates] = {if reset then true else previous(
  nextResetStates[i]) for i in 1:nStates};
```

until a state is eventually executed, then its corresponding reset condition is set to false:

```
Boolean nextResetStates[nStates] =
  if active then
    {if activeState == i then false else activeResetStates[i] for i in 1:
      nStates}
  else
    previous(nextResetStates)
```

The second reset mechanism is implemented with the `selectedReset` and `nextReset` variables. If no reset transition is fired, the `nextReset` is set to false for the next cycle.

17.3.3 Activation handling

When a state is suspended its equations should not be executed, and its variables keep their values – including state-variables in clocked discretized continuous-time equations.

The execution of a sub-state machine has to be suspended when its enclosing state is not active. This activation flag is given as a `Boolean` input `active`. When this flag is true, the sub-state machine maintains its previous state, by guarding the equations of the state variables `nextState`, `nextReset` and `nextResetStates`.

17.3.4 Semantics Summary

The entire semantics model is given below:

```
model StateMachineSemantics "Semantics of state machines"
  parameter Integer nStates;
  parameter Transition t[:] "Array of transition data sorted in priority";
  input
  Boolean c[size(t,1)] "Transition conditions sorted in priority";
  input Boolean active "true if the state machine is active";
  input Boolean reset "true when the state
machine should be reset";
  Integer selectedState = if reset then 1 else previous(nextState);
  Boolean selectedReset = if reset then true else previous(nextReset);
  // For strong (immediate) and weak (delayed) transitions
  Integer immediate = max(if (if t[i].immediate and t[i].from == selectedState
    then c[i] else false) then i else 0 for i in 1:size(t,1));
  Integer delayed = max(if (if not t[i].immediate and t[i].from == nextState
    then c[i] else false) then i else 0 for i in 1:size(t,1));
  Integer fired = max(previous(delayed), immediate);
  output Integer activeState = if reset then 1 elseif fired > 0 then t[fired].
    to else selectedState;
  output Boolean activeReset = if reset then true elseif fired > 0 then t[fired
    ].reset else selectedReset;

  // Update states
  Integer nextState = if active then activeState else previous(nextState);
  Boolean nextReset = if active then false else previous(nextReset);
  // Delayed resetting of individual states
  output Boolean activeResetStates[nStates] = {if reset then true else
    previous(nextResetStates[i]) for i in 1:nStates};
  Boolean nextResetStates[nStates] = if active then {if activeState == i then
    false
```

```

    else activeResetStates[i] for i in 1:nStates} else previous(nextResetStates
    );
    Boolean finalStates[nStates] = {max(if t[j].from == i then 1 else 0 for j in
    1:size(t,1)) == 0 for i in 1:nStates};
    Boolean stateMachineInFinalState = finalStates[activeState];
  end StateMachineSemantics;

```

17.3.5 Merging Variable Definitions

[When a state class uses an **outer output** declaration, the equations have access to the corresponding variable declared **inner**. Special rules are then needed to maintain the single assignment rule since multiple definitions of such outer variables in different mutually exclusive states needs to be merged.]

In each state, the outer output variables are solved for and for each such variable a single definition is formed:

```

v := if activeState(state1) then expre1 elseif activeState(state2) then expre2
    elseif ... else last(v)

```

`last` is special internal semantic operator returning its input. It is just used to mark for the sorting that the incidence of its argument should be ignored. A start value must be given to the variable if not assigned in the initial state.

A new assignment equation is formed which might be merged on higher levels in nested state machines.

17.3.6 Merging Connections to Multiple Outputs

[The causal connection semantics of Modelica for non-state machines are generalized to states of state machines, using the fact that only one state is active at a time.]

It is possible to connect outputs each coming from different states of state machines together – and connect this with other causal connectors. These outputs are combined seen as one source of the signal, and give the following constraint equations,

```

u1 = u2 = ... = y1 = y2 = ...

```

with y_i being outputs from different states of the state-machine and u_i being other causal variables. The semantics is defined similarly to section 17.3.5:

```

v = if activeState(state1) then
    y1
    elseif activeState(state2) then
    y2
    elseif ...
    else
    last(v);
u1 = v
u2 = v
...

```

17.3.7 Example




Figure 17.2: Example of a hierarchical state machine.

[Example: Consider the hierarchical state machine in figure 17.2. The model demonstrates the following properties:

- `state1` is a meta state with two parallel state machines in it.
- `stateA` declares `v` as `outer output`. `state1` is on an intermediate level and declares `v` as `inner outer output`, i.e. matches lower level `outer v` by being `inner` and also matches higher level `inner v` by being `outer`. The top level declares `v` as `inner` and gives the start value.
- `count` is defined with a start value in `state1`. It is reset when a reset transition (`v >= 20`) is made to `state1`.
- `stateX` declares the local variable `w` to be equal to `v` declared as `inner input`.

- `stateY` declares a local counter `j`. It is reset at start and as a consequence of the reset transition (`v >= 20`) to `state1`: When the reset transition (`v >= 20`) fires, then the variables of the active states are reset immediately (so `count` from `state1`, and `i` from `stateX`). The variables of other states are only reset at the time instants when these states become active. So `j` in `StateY` is reset to 0, when the transition `stateX.i > 20` fires (after `state1` became active again, so after the reset transition `v >= 20`).
- Synchronizing the exit from the two parallel state machines of `state1` is done by checking that `stateD` and `stateY` are active using the `activeState` function.

The Modelica code (without annotations) is:

```

block HierarchicalAndParallelStateMachine
  inner Integer v(start=0);

  State1 state1;
  State2 state2;
equation
  initialState(state1);
  transition(state1, state2, activeState(state1.stateD) and
    activeState(state1.stateY), immediate=false);
  transition(state2, state1, v >= 20, immediate=false);

public
  block State1
    inner Integer count(start=0);
    inner outer output Integer v;

    block StateA
      outer output Integer v;
    equation
      v = previous(v) + 2;
    end StateA;
    StateA stateA;

    block StateB
      outer output Integer v;
    equation
      v = previous(v) - 1;
    end StateB;
    StateB stateB;

    block StateC
      outer output Integer count;
    equation
      count = previous(count) + 1;
    end StateC;
    StateC stateC;

    block StateD
    end StateD;
    StateD stateD;

  equation
    initialState(stateA);
    transition(stateA, stateB, v >= 6, immediate=false);
    transition(stateB, stateC, v == 0, immediate=false);
    transition(stateC, stateA, true, immediate=false, priority=2);
    transition(stateC, stateD, count >= 2, immediate=false);

public
  block StateX
    outer input Integer v;
    Integer i(start=0);
  
```

```

    Integer w; // = v;
    equation
    i = previous(i) + 1;
    w = v;
    end StateX;
    StateX stateX;

    block StateY
    Integer j(start=0);
    equation
    j = previous(j) + 1;
    end StateY;
    StateY stateY;

    equation
    initialState(stateX);
    transition(stateX, stateY, stateX.i > 20, immediate=false,
    reset=false);
    end State1;

    block State2
    outer output Integer v;
    equation
    v = previous(v) + 5;
    end State2;
    end HierarchicalAndParallelStateMachine;

```

Figure 17.3 shows the behavior of the state machine.




Figure 17.3: State machine behavior, as reflected by the variable v .

The transition from `state1` to `state2` could have been done with a `synchronize` transition with condition `=true` instead. The semantically equivalent model is shown below:

```

block HierarchicalAndParallelStateMachine
  extends StateMachineSemantics(
    nStates=2,
    t={Transition(from=1, to=2, immediate=false, synchronize=true),
      Transition(from=2, to=1, immediate=false)},
    c={true, v >= 20});
  Boolean init(start=true) = sample(false);

  block State1
    Boolean active;
    Boolean reset;
    outer input Integer v_previous;
    inner output Integer v;
    inner Integer count(start=0);

```

```

inner Integer count_previous = if reset then 0 else previous(count);

block StateMachineOf_stateA
  extends StateMachineSemantics(
    nStates=4,
    t={Transition(from=1, to=2, immediate=false),
      Transition(from=2, to=3, immediate=false),
      Transition(from=3, to=1, immediate=false),
      Transition(from=3, to=4, immediate=false)},
    c={v >= 6, v==0, true, count >= 2});
  outer input Integer v_previous;
  outer output Integer v;
  outer input Integer count_previous;
  outer output Integer count;
equation
  inFinalState = true; // no macro states
  if activeState == 1 then
    // equations for stateA
    v = v_previous + 2;
    count = count_previous;
  elseif activeState == 2 then
    // equations for stateB
    v = v_previous - 1;
    count = count_previous;
  elseif activeState == 3 then
    // equations for stateC
    v = v_previous;
    count = count_previous + 1;
  else // if activeState == 4 then
    // equations for stateD
    v = v_previous;
    count = count_previous;
  end if;
end StateMachineOf_stateA;

StateMachineOf_stateA stateMachineOf_stateA(active=active,
  reset=reset);

block StateMachineOf_stateX
  extends StateMachineSemantics(
    nStates=2,
    t={Transition(from=1, to=2, immediate=false, reset=false)},
    c={i>25});
  outer input Integer v;
  Integer i(start=0);
  Integer i_previous;
  Integer j(start=0);
  Integer j_previous;
  Integer w;
equation
  inFinalState = true; // no macro states
  if activeState == 1 then
    // equations for stateX
    i_previous = if activeReset or
      activeResetStates[1] then 0 else previous(i);
    j_previous = previous(j);
    i = i_previous + 1;
    w = v;
    j = j_previous;
  else // if activeState == 2 then
    // equations for stateY
    i_previous = previous(i);
    j_previous = if activeReset or activeResetStates[2] then 0 else

```



```

previous(j);
  i = i_previous;
  w = previous(w);
  j = j_previous + 1;
end if;
end StateMachineOf_stateX;

StateMachineOf_stateX stateMachineOf_stateX(active=active,
reset=reset);
Boolean inFinalState = stateMachineOf_stateA.stateMachineInFinalState and
stateMachineOf_stateX.stateMachineInFinalState;
end State1;

State1 state1;
Integer v(start=0);
inner Integer v_previous = if reset then 0 else previous(v);
equation
active = true;
reset = previous(init);
if activeState == 1 then
  // equations for state1
inFinalState = state1.inFinalState;
state1.active = true;
state1.reset = activeReset or activeResetStates[1];
v = state1.v;
else // if activeState == 2 then
  // equations for state2
inFinalState = true; // not macro state
state1.active = false;
state1.reset = false;
v = previous(v) + 5;
end if;
end HierarchcialAndParallelStateMachine;

```

]

Chapter 18

Annotations

Annotations are intended for storing extra information about a model, such as graphics, documentation or versioning, etc. A Modelica tool is free to define and use other annotations, in addition to those defined here, according to section 18.1. Annotations must be represented with constructs according to the Modelica grammar (for replaceable class declarations with a constraining-clause also refer to section 7.3.2.1). The specification in this document defines the semantic meaning if a tool implements any of these annotations.

18.1 Vendor-Specific Annotations

A vendor may – anywhere inside an annotation – add specific, possibly undocumented, annotations which are not intended to be interpreted by other tools. The only requirement is that any tool shall save files with all vendor-specific annotations (and all annotations from this chapter) intact. Two variants of vendor-specific annotations exist; one simple and one hierarchical. Double underscore concatenated with a vendor name as initial characters of the identifier are used to identify vendor-specific annotations.

[*Example:*

```
annotation (
  Icon(coordinateSystem(extent = {{-100, -100}, {100, 100}}),
    graphics = {__NameOfVendor(Circle(center = {0, 0}, radius = 10))});
```

This introduces a new graphical primitive Circle using the hierarchical variant of vendor-specific annotations.

```
annotation (
  Icon(coordinateSystem(extent = {{-100, -100}, {100, 100}}),
    graphics = {Rectangle(extent = {{-5, -5}, {7, 7}},
      __NameOfVendor_shadow = 2)});
```

This introduces a new attribute __NameOfVendor_shadow for the Rectangle primitive using the simple variant of vendor-specific annotations.]

18.2 Annotations for Documentation

The Documentation annotation has the following contents, where the info and revisions annotations are described in section 18.2.1, and the figures annotation is described in section 18.2.2:

```
record Documentation
  String info = "" "Description of the class";
  String revisions = "" "Revision history";
  Figure[:] figures = {}; "Simulation result figures";
end Documentation;
```

How the tool interprets the information in Documentation is unspecified.

18.2.1 Class Description and Revision History

Inside the `Documentation` annotation, the `info` annotation gives a textual description of the class, and the `revisions` annotation gives a revision history.

[The `revisions` documentation may be omitted in printed documentation.]

If the string starts with the tag `<html>` or `<HTML>` the entire string is HTML encoded (and is assumed to end with `</html>` or `</HTML>` and shall be rendered as HTML even if the end-tags are missing), otherwise the entire string is rendered as is. The HTML encoded content may contain links. For external links, see section 13.5. Links to Modelica classes may be defined with the HTML link command using scheme `Modelica` (using its lower case form in the URI, see section 13.5), e.g.,

```
<a href="modelica://MultiBody.Tutorial">MultiBody.Tutorial</a>
```

Together with scheme `Modelica` the (URI) fragment specifiers `#diagram`, `#info`, `#text`, `#icon` may be used to reference different layers. User-defined fragment specifiers (anchors) may also be used, and they may be renamed when generating HTML (in particular to avoid collisions). Example:

```
<a href="modelica://MultiBody.Joints.Revolute#info">Revolute</a>
```

18.2.2 Annotations for Figures

Inside the `Documentation` annotation, each element of the `figures` annotation array has the following content:

```
record Figure
  String title = "" "Title meant for display";
  String identifier "Identifier meant for programmatic access";
  String group = "" "Name of plot group";
  Boolean preferred = false "Automatically display figure after simulation";
  Plot[:] plots "Plots";
  String caption "Figure caption";
end Figure;
```

A `Figure` is a graphical container that can contain several `plots` described by `Plot` annotations:

```
record Plot
  String title "Title meant for display";
  String identifier "Identifier meant for programmatic access";
  Curve[:] curves "Plot curves";
  Axis x "X axis properties";
  Axis y "Y axis properties";
end Plot;
```

A `Plot` can contain several `curves`, see section 18.2.2.2, that all share a common x and y axis with properties described in section 18.2.2.1.

Both `Figure` and `Plot` can have an optional title. When the `Figure` `title` is the empty string (the default), the tool must produce a non-empty title based on the figure content. On the other hand, the `Plot` `title` has a tool-dependent default, but the default may be the empty string. When the `Plot` `title` is the empty string, no title should be shown. The plot title is not to be confused with the plot *label* which is never empty, see below. Variable replacements, as described in section 18.2.2.4, can be used in the title of `Figure` and `Plot`.

The `identifier` in `Figure` and `Plot` is an optional `String` identifier, and is intended to identify the `Figure` and `Plot` for programmatic access. The `figures` annotation is inherited in the sense that each class has a collection of figures comprised by the contents of the `figures` annotation in the class itself, as well as the `figures` annotations from any base classes. A `Figure` must be uniquely identified by its `identifier` and a class having it in its collection. This means that a `Figure` `identifier` must be unique among all `Figure` annotations within the same `figures` annotation as well as among all `figures` annotations from inherited classes. A `Plot` `identifier` on the other hand is only required to be unique among the `plots` in the the same `Figure` annotation.

[For **Figure**, this makes it possible to reference the plot from a tool-specific scripting environment. For **Plot**, this makes it possible to reference the plot in the figure caption, which becomes useful when the **Figure** contains more than one **Plot**.]

Even though a **Figure** annotation can be shared through inheritance between classes in a class hierarchy, note that each simulated class provides its own data to be displayed in the figure.

Every **Plot** has an automatically generated *label* which is required to be shown as soon as at least one **Plot** in the **Figure** has an **identifier**. A tool is free to choose both labeling scheme (such as *a*, *b*, ..., or *i*, *ii*, ...), placement in the plot, and styling in the plot itself as well as in other contexts.

When a **Figure** defines a non-empty **group**, it is used to organize figures similar to how **group** is used in the **Dialog** annotation (see section 18.7). However, leaving **group** at the default of an empty string does not mean that a group will be created automatically, but that the figure resides outside of any group. The **group** is both the key used for grouping, and the name of the group for display purposes.

The **preferred** attribute of **Figure** indicates whether the figure should be given preference when automatically determining which figures to show, and a class may define any number of **preferred** figures. For example, a tool might choose to automatically show all preferred figures when the class is simulated.

The **caption** attribute of **Figure** can use the restricted form of text markup described in section 18.2.2.5 as well as the variable replacements described in section 18.2.2.4.

18.2.2.1 Axis Properties

Properties may be defined for each **Plot** axis:

```
record Axis
  Real min "Axis lower bound, in 'unit'";
  Real max "Axis upper bound, in 'unit'";
  String unit = "" "Unit of axis tick labels";
  String label "Axis label";
end Axis;
```

When an axis bound isn't provided, the tool computes one automatically.

An empty **unit** means that the axis is unitless, and each expression plotted against it may use its own unit determined by the tool. The tool is responsible for conveying the information about choice of unit for the different variables, for instance by attaching this information to curve legends.

The Modelica tool is responsible for showing that values at the axis tick marks are expressed in **unit**, so the axis **label** shall not contain this information.

[When **unit** is empty, and axis bounds are to be determined automatically, a natural choice of unit could be the variable's **displayUnit**. When axis bounds are specified by the user, on the other hand, a tool may choose a unit for the variable such that the range of the variable values (expressed in the chosen unit) fit nicely with the range of the unitless axis.]

If a tool does not recognize the **unit**, it is recommended to issue a warning and treat the **unit** as if it was empty, as well as ignore any setting for **min** and **max**.

When an axis label isn't provided, the tool produces a default label. Providing the empty string as axis label means that no label should be shown. Variable replacements, as described in section 18.2.2.4, can be used in the **label** of **Axis**. The Modelica tool is responsible for showing the unit used for values at the axis tick marks, so the axis **label** shall not contain the unit.

18.2.2.2 Plot Curves

The actual data to plot is specified in the **curves** of a **Plot**:

```
record Curve
  expression x = time "X coordinate values";
  expression y "Y coordinate values";
  String legend "Legend";
end Curve;
```

The mandatory `x` and `y` expressions are restricted to be component references referring to a scalar variable or `time`. It is an error if `x` or `y` does not designate a scalar variable. When the `unit` of an `Axis` is non-empty, it is an error if the unit of the corresponding `Curve` expression (i.e., a variable's `unit`, or second for `time`) is incompatible with the axis unit.

When `legend` isn't provided, the tool produces a default based on `x` and/or `y`. Providing the empty string as `legend` means that the curve shall be omitted from the plot legend. Variable replacements, as described in section 18.2.2.4, can be used in the `legend` of `Curve`

18.2.2.3 Escape sequences

In an attribute inside a figure where the variable replacements of section 18.2.2.4 or the text markup of section 18.2.2.5 can be used, the following use of *text markup escape sequences* applies. These escape sequences are applied after the application of other markup, and is not applied at all inside some of the other markup, see details for the respective markup.

The percent character `'%'` shall be encoded `%%`. The following are all the recognized escape sequences:

<i>Sequence</i>	<i>Encoded character</i>	<i>Comment</i>
<code>%%</code>	<code>'%'</code>	Only way to encode character.
<code>%]</code>	<code>']'</code>	Escape sequence only needed in link text.

[With the percent character being encoded as `%%`, the behavior of `%` appearing in any other way than the escape sequences above, for variable replacement (section 18.2.2.4), or for link markup (section 18.2.2.5) is undefined, and thus possible to define in the future without breaking backward compatibility.]

18.2.2.4 Variable Replacements

In the places listed in table 18.1 where text for display is defined, the final value of a result variable can be embedded by referring to the variable as `%{inertial.w}`. This is similar to the `Text` graphical primitive in section 18.6.5.5.

Table 18.1: Attributes that can use variable replacements.

<i>Attribute</i>	<i>Annotation</i>
<code>title</code>	Figure and Plot
<code>caption</code>	Figure
<code>legend</code>	Curve
<code>label</code>	Axis

In `%{variable}`, text markup escape sequences don't apply inside the *variable*, which has the form of `component-reference` in the grammar (appendix B.2.7). This means that a complete `component-reference` shall be scanned before looking for the terminating closing brace.

[Example: The variable replacement `%{'%%'}` references the variable `'%%'`, not the variable `'%'`.]

[Example: The variable replacement `%{foo . '}'bar{'}` makes a valid reference to the variable `foo . '}'bar{'}`.]

Note that expansion to the final value means that expansion is not restricted to parameters and constants, so that values to be shown in a caption can be determined during simulation.

[By design, neither `%class` nor `%name` is supported in this context, as this information is expected to already be easily accessible (when applicable) in tool-specific ways. (Titles making use of `%class` or `%name` would then only lead to ugly duplication of this information.)]

18.2.2.5 Text Markup in Captions

In addition to variable replacements, a very restricted form of text markup is used for the `caption`. Note that the text markup escape sequences described in section 18.2.2.3 generally apply inside `caption`, with one exception given below for links.

Links take the form `%[text](link)`, where the `[text]` part is optional. Text markup escape sequences don't apply inside the `link`, which can be in either of the following forms:

- A URI. Well established schemes such as `https://github.com/modelica/ModelicaSpecification` or `modelica://Modelica`, as well as lesser known schemes may be used. (A tool that has no special recognition of a scheme can try sending the URI to the operating system for interpretation.)
- A `variable:id`, where `id` is a component reference in the form of `component-reference` in the grammar, such as `inertia1.w`.
- A `plot:id`, where `id` is the identifier of a Plot in the current Figure.

When `[text]` is omitted, a Modelica tool is free to derive a default based on the `link`.

[Note that for the character ']' to appear in text, it needs to be encoded as the escape sequence %], or it would be interpreted as the terminating delimiter of the [text].

Similarly, the closing parenthesis ')' must be handled with care in link in order to not be interpreted as the terminating delimiter of the (link).

- *For a URI, a closing parenthesis must be URL encoded in order to not be interpreted as the end of the (link). For example, the URL in `%(http://example.org/(tryme))` is just `http://example.org/(tryme`, and the entire link is followed by a stray closing parenthesis. To make it work, one has to use URL encoding: `%(http://example.org/%28tryme%29)` (using URL encoding of the opening parenthesis just for symmetry, and note that the % of the percent-encoded sequences are not subject to text markup escape sequences).*
- *For a `variable:.`, no special treatment is needed, as the component reference syntax of the `<id >` allows parentheses to appear without risk of misinterpretation inside a quoted identifier. For example, `%(variable:'tryme!')` has a parenthesis in `'tryme!'` that must not be mistaken for the end of the (link).*
- *For a `plot:.`, there is currently no way to reference a plot with ')' in its identifier.*

]

The styling of the link text, as well as the link action, is left for each Modelica tool to decide.

[For example, `%(inertia1.w)` could be displayed as the text `inertia1.w` formatted with upright monospaced font, and have a pop-up menu attached with menu items for plotting the variable, setting its start value, or investigating the equation system from which it is solved. On the other hand, `%(angular velocity](inertia1.w)` could be formatted in the same style as the surrounding text, except some non-intrusive visual clue about it being linked.]

A sequence of one or more newlines (encoded either literally or using the `\n` escape sequence) means a paragraph break. (A line break within a paragraph is not supported, and any paragraph break before the first paragraph or after the last paragraph has no impact.)

18.3 Annotations for Code Generation

The annotations listed below, appearing directly inside `annotation(...)`, can influence the code generation.

<i>Annotation</i>	<i>Description</i>	<i>Details</i>
<code>Evaluate</code>	Use parameter value for symbolic processing	Annotation 18.1
<code>HideResult</code>	Don't show component's simulator result	Annotation 18.2
<code>Inline</code>	Inline function	Annotation 18.3
<code>LateInline</code>	Inline after all symbolic transformations	Annotation 18.4
<code>InlineAfterIndexReduction</code>	Inline after index reduction	Annotation 18.5
<code>GenerateEvents</code>	Generate events for zero crossings in function	Annotation 18.6
<code>smoothOrder</code>	Function smoothness guarantee	Annotation 18.7

Annotation 18.1 Evaluate

```
"Evaluate" "=" ( false | true )
```

The annotation `Evaluate` can occur in the component declaration, its type declaration, or a base-class of the type-declaration. In the case of multiple conflicting annotations it is handled similarly to modifiers (e.g., an `Evaluate` annotation on the component declaration takes precedence). In the case of hierarchical components it is applied to all components, and if both of them have annotations then `Evaluate = true` takes precedence. The annotation `Evaluate` only has effect for a component declared with the prefix `parameter`.

If `Evaluate = true`, the model developer proposes to utilize the value for the symbolic processing. In that case, it is not possible to change the parameter value after symbolic pre-processing.

If `Evaluate = false`, the model developer proposes to not utilize the value of the corresponding parameter for the symbolic processing.

[Evaluate is for example used for axis of rotation parameters in the Modelica.Mechanics.MultiBody library in order to improve the efficiency of the generated code.]

Annotation 18.2 HideResult

```
"HideResult" "=" ( false | true )
```

`HideResult = true` defines that the model developer proposes to not show the simulator results of the corresponding component.

`HideResult = false` defines that the developer proposes to show the corresponding component.

[For example, a tool is not expected to provide means to plot a variable with `HideResult = true`. If a variable is declared in a protected section, a tool might not include it in a simulation result. By setting `HideResult = false`, the modeler would like to have the variable in the simulation result, even if in the protected section.]

HideResult is for example used in the connectors of the Modelica.StateGraph library to not show variables to the modeler that are of no interest to him and would confuse him.]

Annotation 18.3 Inline

```
"Inline" "=" ( false | true )
```

Has only an effect within a function declaration.

If `Inline = true`, the model developer proposes to inline the function. This means, that the body of the function is included at all places where the function is called.

If `Inline = false`, the model developer proposes to not inline the function.

[Inline = true is for example used in Modelica.Mechanics.MultiBody.Frames and in functions of Modelica.Media to have no overhead for function calls such as resolving a vector in a different coordinate system and at the same time the function can be analytically differentiated, e.g., for index reduction needed for mechanical systems.]

Annotation 18.4 LateInline

```
"LateInline" "=" ( false | true )
```

Has only an effect within a function declaration.

If `LateInline = true`, the model developer proposes to inline the function after all symbolic transformations have been performed.

[Late inlining is especially useful for differentiation and inversion of functions; for efficiency reasons it is then useful to replace all function calls with identical input arguments by one function call, before the inlining.]

If `LateInline = false`, the model developer proposes to not inline the function after symbolic transformations have been performed.

`Inline = true, LateInline = false` is identical to `Inline = true`.

`Inline = true, LateInline = true` is identical to `LateInline = true`.

`Inline = false, LateInline = true` is identical to `LateInline = true`.

[This annotation is for example used in `Modelica.Media.Water.IF97_Uutilities.T_props_ph` to provide in combination with common subexpression elimination the automatic caching of function calls. Furthermore, it is used in order that a tool is able to propagate specific enthalpy over connectors in the `Modelica.Fluid` library.]

Annotation 18.5 `InlineAfterIndexReduction`

```
"InlineAfterIndexReduction" "=" ( false | true )
```

Has only an effect within a function declaration.

If `true`, the model developer proposes to inline the function after the function is differentiated for index reduction, and before any other symbolic transformations are performed. This annotation cannot be combined with annotations `Inline` and `LateInline`.

Annotation 18.6 `GenerateEvents`

```
"GenerateEvents" "=" ( false | true )
```

Has only an effect within a function declaration

If `GenerateEvents = true`, the model developer proposes that crossing functions in the function should generate events (one possibility of doing this is to inline the function and generate events for the inlined function).

[This annotation is for example used in `Modelica.Media.Water.IF97_Uutilities.phase_dT` to indicate that the output should generate an event when it changes.]

Annotation 18.7 `smoothOrder`

```
"smoothOrder" "=" UNSIGNED-NUMBER ")"
"smoothOrder" "(" "normallyConstant" "=" IDENT
    { "," "normallyConstant" "=" IDENT } ")" "="
    UNSIGNED-NUMBER ")"
```

This annotation has only an effect within a function declaration.

`smoothOrder` defines the number of differentiations of the function, in order that all of the differentiated outputs are continuous provided all input arguments and their derivatives up to order `smoothOrder` are continuous.

[This means that the function is at least $C^{\text{smoothOrder}}$. `smoothOrder = 1` means that the function can be differentiated at least once in order that all output arguments are still continuous, provided the input arguments are continuous. If a tool needs the derivative of a function, e.g. for index reduction or to compute an analytic Jacobian, the function can be differentiated analytically at least `smoothOrder` times.]

The optional argument `normallyConstant` of `smoothOrder` defines that the function argument `IDENT` is usually constant.

[A tool might check whether the actual argument to `IDENT` is a parameter expression at the place where the function is called. If this is the case, the derivative of the function might be constructed under the assumption that the corresponding argument is constant, to enhance efficiency. Typically, a tool would generate at most two different derivative functions of a function: One, under the assumption that all `normallyConstant` arguments are actually constant. And one, under the assumption that all input arguments are time varying. Based on the actual arguments of the function call either of the two derivative functions is used.]

This annotation is used by many functions of the `Modelica.Fluid` library, such as `Modelica.Fluid.Dissipation.PressureLoss.StraightPipe.dp_laminar_DP`, since geometric arguments to these functions are usually constant.]

18.4 Annotations for Simulations

These annotations define how models can be checked, translated, and simulated.

18.4.1 Annotations for Simulation Experiments

```

experiment-annotation :
  annotation "(" "experiment" [ "(" [experimentOption
    {"," experimentOption}] ")" ] ")"

experimentOption :
  "StartTime" "=" [ "+" | "-" ] UNSIGNED-NUMBER |
  "StopTime" "=" [ "+" | "-" ] UNSIGNED-NUMBER |
  "Interval" "=" UNSIGNED-NUMBER |
  "Tolerance" "=" UNSIGNED-NUMBER
  
```

The `experiment` annotation defines the default start time (`StartTime`) in [s], the default stop time (`StopTime`) in [s], the suitable time resolution for the result grid (`Interval`) in [s], and the default relative integration tolerance (`Tolerance`) for simulation experiments to be carried out with the model or block at hand. If `StartTime` is not specified it is assumed to be 0.0.

18.4.2 Annotation for Test Cases

```

testcase-annotation :
  annotation "(" "TestCase" "(" "shouldPass" "=" ( false | true ) ")" ")"
  
```

If `shouldPass` is `false` it indicates that the translation or the simulation of the model should fail. If a tools checks a package where classes have `shouldPass=false` they should not generate errors, and checking may even be skipped. On the other hand, models with `shouldPass=false` may be useful for creation of negative tests in tool-specific ways. Similarly as a class with `obsolete-annotation`, a class with `TestCase` annotation (regardless of the value of `shouldPass`) shall not be used in other models, unless those models also have a `TestCase` annotation.

[The intent of the test-case can be included in the documentation of the class. This annotation can both be used for models intended as test-cases for implementations, and for models explaining detectable errors.]

18.5 Annotation for single use of class

For state machines it is useful to have single instances of local classes. This can be done using:

```

annotation(singleInstance=true)
  
```

The annotation `singleInstance` in a class indicates that there should only be one component instance of the class, and it should be in the same scope as the class is defined. The intent is to remove the class when the component is removed and to prevent duplication of the component.

18.6 Annotations for Graphical Objects

A graphical representation of a class consists of two abstraction layers, icon layer and diagram layer showing graphical objects, component icons, connectors and connection lines. The icon representation typically visualizes the component by hiding hierarchical details. The hierarchical decomposition is described in the diagram layer showing icons of subcomponents and connections between these.

Graphical annotations described in this chapter ties into the Modelica grammar as follows.

```

graphical-annotations :
  annotation "(" [ layer-annotations ] ")"

layer-annotations :
  ( icon-layer | diagram-layer ) [ "," layer-annotations ]
  
```

Layer descriptions (start of syntactic description):

```

icon-layer :
  "Icon" "(" [ coordsys-specification "," ] graphics ")"
  
```

```

diagram-layer :
  "Diagram" "(" [ coordsys-specification "," ] graphics ")"

```

[Example:

```

annotation(
  Icon(coordinateSystem(extent = {{-100, -100}, {100, 100}}),
    graphics = {Rectangle(extent = {{-100, -100}, {100, 100}}),
      Text(extent = {{-100, -100}, {100, 100}},
        textString = "Icon")});

```

]

The graphics is specified as an ordered sequence of graphical primitives, which are described below. First base-class contents is drawn according to the order of the extends-clauses, and then graphical primitives are drawn according to the order such that later objects can cover earlier ones.

[Note that the ordered sequence is syntactically a valid Modelica annotation, although there is no mechanism for defining an array of heterogeneous objects in Modelica.]

These Icon, Diagram, and Documentation annotations are only allowed directly in classes (e.g. not on components or connections). The allowed annotations for a short class definition is the union of the allowed annotations in classes and on extends-clauses.

18.6.1 Common Definitions

The following common definitions are used to define graphical annotations in the later sections.

```

type DrawingUnit = Real(final unit="mm");
type Point = DrawingUnit[2] "{x, y}";
type Extent = Point[2] "Defines a rectangular area {{x1, y1}, {x2, y2}}";

```

The interpretation of `unit` is with respect to printer output in natural size (not zoomed).

All graphical entities have a visible attribute which indicates if the entity should be shown.

```

partial record GraphicItem
  Boolean visible = true;
  Point origin = {0, 0};
  Real rotation(quantity="angle", unit="deg")=0;
end GraphicItem;

```

The `origin` attribute specifies the origin of the graphical item in the coordinate system of the layer in which it is defined. The origin is used to define the geometric information of the item and for all transformations applied to the item. All geometric information is given relative the `origin` attribute, which by default is `{0, 0}`.

The `rotation` attribute specifies the rotation of the graphical item counter-clockwise around the point defined by the `origin` attribute.

18.6.1.1 Coordinate Systems

Each of the layers has its own coordinate system. A coordinate system is defined by the coordinates of two points, the left (`x1`) lower (`y1`) corner and the right (`x2`) upper (`y2`) corner, where the coordinates of the first point shall be less than the coordinates of the second point.

The attribute `preserveAspectRatio` specifies a hint for the shape of components of the class, but does not actually influence the rendering of the component. If `preserveAspectRatio` is true, changing the extent of components should preserve the current aspect ratio of the coordinate system of the class.

The attribute `initialScale` specifies the default component size as `initialScale` times the size of the coordinate system of the class. An application may use a different default value of `initialScale`.

The attribute `grid` specifies the spacing between grid points which can be used by tools for alignment of points in the coordinate system, e.g. "snap-to-grid". Its use and default value is tool-dependent.

```

record CoordinateSystem
  Extent extent;
  Boolean preserveAspectRatio=true;
  Real initialScale = 0.1;
  DrawingUnit grid[2];
end CoordinateSystem;
    
```

[Example: A coordinate system for an icon could for example be defined as:

```
CoordinateSystem(extent = {{-10, -10}, {10, 10}});
```

i.e. a coordinate system with width 20 units and height 20 units.]

The coordinate systems for the icon and diagram layers are by default defined as follows; where the array of `GraphicsItem` represents an ordered list of graphical primitives.

```

record Icon "Representation of the icon layer"
  CoordinateSystem coordinateSystem(extent = {{-100, -100}, {100, 100}});
  GraphicItem[:] graphics;
end Icon;

record Diagram "Representation of the diagram layer"
  CoordinateSystem coordinateSystem(extent = {{-100, -100}, {100, 100}});
  GraphicItem[:] graphics;
end Diagram;
    
```

The coordinate system (including `preserveAspectRatio`) of a class is defined by the following priority:

1. The coordinate system annotation given in the class (if specified).
2. The coordinate systems of the first base-class where the extent on the extends-clause specifies a null-region (if any). Note that null-region is the default for base-classes, see section 18.6.3.
3. The default coordinate system `CoordinateSystem(extent = {{-100, -100}, {100, 100}})`.

18.6.1.2 Graphical Properties

Properties of graphical objects and connection lines are described using the following attribute types.

```

type Color = Integer[3](min = 0, max = 255) "RGB representation";
constant Color Black = zeros(3);
type LinePattern = enumeration(None, Solid, Dash, Dot, DashDot, DashDotDot);
type FillPattern = enumeration(None, Solid, Horizontal, Vertical,
                               Cross, Forward, Backward, CrossDiag,
                               HorizontalCylinder, VerticalCylinder, Sphere);
type BorderPattern = enumeration(None, Raised, Sunken, Engraved);
type Smooth = enumeration(None, Bezier);
type EllipseClosure = enumeration(None, Chord, Radial);
    
```

The `LinePattern` attribute `Solid` indicates a normal line, `None` an invisible line, and the other attributes various forms of dashed/dotted lines.

The `FillPattern` attributes `Horizontal`, `Vertical`, `Cross`, `Forward`, `Backward` and `CrossDiag` specify fill patterns drawn with the line color over the fill color.

The attributes `HorizontalCylinder`, `VerticalCylinder` and `Sphere` specify gradients that represent a horizontal cylinder, a vertical cylinder and a sphere, respectively. The gradient goes from line color to fill color.

The border pattern attributes `Raised`, `Sunken` and `Engraved` represent frames which are rendered in a tool-dependent way — inside the extent of the filled shape.




Figure 18.1: Line with `smooth = Bezier`. The four line points P_1, \dots, P_4 result in two quadratic splines and two straight line segments.

The `smooth` attribute specifies that a line can be drawn as straight line segments (`None`) or using a spline (`Bezier`), where the line's points specify control points of a quadratic Bezier curve, see figure 18.1.

For lines with only two points, the `smooth` attribute has no effect.

For lines with three or more points (P_1, P_2, \dots, P_n), the middle point of each line segment ($P_{12}, P_{23}, \dots, P_{(n-1)n}$) becomes the starting point and ending points of each quadratic Bezier curve. For each quadratic Bezier curve, the common point of the two line segment becomes the control point. For instance, point P_2 becomes the control point for the Bezier curve starting at P_{12} and ending at P_{23} . A straight line is drawn between the starting point of the line and the starting point of the first quadratic Bezier curve, as well as between the ending point of the line and the ending point of the last quadratic Bezier curve.

In the illustration above, the square points (P_1, P_2, P_3 , and P_4) represent the points that define the line, and the circle points (P_{12}, P_{23} , and P_{34}) are the calculated middle points of each line segment. Points P_{12}, P_2 , and P_{23} define the first quadratic Bezier curve, and the points P_{23}, P_3 , and P_{34} define the second quadratic Bezier curve. Finally a straight line is drawn between points P_1 and P_{12} as well as between P_{34} and P_4 .

The values of the `EllipseClosure` enumeration specify if and how the endpoints of an elliptical arc are to be joined (see section 18.6.5.4).

```
type Arrow = enumeration(None, Open, Filled, Half);
type TextStyle = enumeration(Bold, Italic, UnderLine);
type TextAlignment = enumeration(Left, Center, Right);
```

Filled shapes have the following attributes for the border and interior.

```
record FilledShape "Style attributes for filled shapes"
  Color lineColor = Black "Color of border line";
  Color fillColor = Black "Interior fill color";
  LinePattern pattern = LinePattern.Solid "Border line pattern";
  FillPattern fillPattern = FillPattern.None "Interior fill pattern";
  DrawingUnit lineThickness = 0.25 "Line thickness";
end FilledShape;
```

The extent/points of the filled shape describe the theoretical zero-thickness filled shape, and the actual rendered border is then half inside and half outside the extent.

18.6.2 Component Instance

A component instance can be placed within a diagram or icon layer. It has an annotation with a `Placement` modifier to describe the placement. Placements are defined in term of coordinate systems transformations:

```
record Transformation
  Point origin = {0, 0};
  Extent extent;
  Real rotation(quantity="angle", unit="deg")=0;
end Transformation;
```

The `origin` attribute defines the position of the component in the coordinate system of the enclosing class. The `extent` defines the position, size and flipping of the component, relative to the `origin` attribute. The `extent` is defined relative to the `origin` attribute of the component instance. Given an extent $\{\{x1, y1\}, \{x2, y2\}\}$, $x2 < x1$ defines horizontal flipping and $y2 < y1$ defines vertical flipping around the center of the object.

The `rotation` attribute specifies rotation of the extent around the point defined by the `origin` attribute.

The graphical operations are applied in the order: scaling, flipping and rotation.

```

record Placement
  Boolean visible = true;
  Transformation transformation "Placement in the diagram layer";

  Boolean iconVisible "Visible in icon layer; for public connector";
  Transformation iconTransformation "Placement in the icon layer; for public
  connector";
end Placement;

```

If no `iconTransformation` is given the `transformation` is also used for placement in the icon layer. If no `iconVisible` is given for a public connector the `visible` is also used for visibility in the icon layer.

[A connector can be shown in both an icon layer and a diagram layer of a class. Since the coordinate systems typically are different, placement information needs to be given using two different coordinate systems. More flexibility than just using scaling and translation is needed since the abstraction views might need different visual placement of the connectors. The attribute `transformation` gives the placement in the diagram layer and `iconTransformation` gives the placement in the icon layer. When a connector is shown in a diagram layer, its diagram layer is shown to facilitate opening up a hierarchical connector to allow connections to its internal subconnectors.]

For connectors, the icon layer is used to represent a connector when it is shown in the icon layer of the enclosing model. The diagram layer of the connector is used to represent it when shown in the diagram layer of the enclosing model. Protected connectors are only shown in the diagram layer. Public connectors are shown in both the diagram layer and the icon layer. Non-connector components are only shown in the diagram layer.

18.6.3 Extends clause

Each extends-clause (and short-class-definition, as stated in section 18.6) may have layer specific annotations which describe the rendering of the base class' icon and diagram layers in the subclass.

```

record IconMap
  Extent extent = {{0, 0}, {0, 0}};
  Boolean primitivesVisible = true;
end IconMap;

record DiagramMap
  Extent extent = {{0, 0}, {0, 0}};
  Boolean primitivesVisible = true;
end DiagramMap;

```

All graphical objects are by default inherited from a base class. If the `primitivesVisible` attribute is false, components and connections are visible but graphical primitives are not.

- If the extent of the extends-clause defines a null region (the default), the base class contents is mapped to the same coordinates in the derived class, and the coordinate system (including `preserveAspectRatio`) can be inherited as described in section 18.6.1.1.
- If the extent of the extends-clause defines a non-null region, the base class coordinate system is mapped to the region specified by the attribute `extent`, if `preserveAspectRatio` is true for the base class the mapping shall preserve the aspect ratio. The base class coordinate system (and `preserveAspectRatio`) is not inherited.

[Example:

```

model A
  extends B annotation(
    IconMap(extent = {{-100, -100}, {100, 100}}, primitivesVisible = false),
    DiagramMap(extent = {{-50, -50}, {0, 0}}, primitivesVisible = true)
  );
end A;

model B
  extends C annotation(DiagramMap(primitivesVisible = false));
  ...
end B;

```

In this example the diagram of A contains the graphical primitives from A and B (but not from C since they were hidden in B) – the ones from B are rescaled, and the icon of A contains the graphical primitives from A (but neither from B nor from C).]

18.6.4 Connections

A connection is specified with an annotation containing a `Line` primitive and optionally a `Text` primitive, as specified below.

[Example:

```

connect(a.x, b.x)
  annotation(Line(points = {{-25, 30}, {10, 30}, {10, -20}, {40, -20}}));

```

]

The optional `Text` primitive defines a text that will be written on the connection line. It has the following definition (*it is not equal to the `Text` primitive as part of graphics – the differences are marked as bold lines*):

```

record Text
  extends GraphicItem;
  extends FilledShape;
  Extent extent;
  String string;
  Real fontSize = 0 "unit pt";
  String fontName;
  TextStyle textStyle[:];
  Color textColor = lineColor;
  TextAlignment horizontalAlignment =
    if index < 0 then TextAlignment.Right else TextAlignment.Left;
  Integer index;
end Text;

```

The `index` is one of the points of `Line` (numbered 1, 2, 3, ... where negative numbers count from the end, thus -1 indicate the last one). The `string` may use the special symbols "`%first`" and "`%second`" to indicate the connectors in the connect-equation.

The `extent` and `rotation` are relative to the `origin` (default `{0, 0}`) and the `origin` is relative to the point on the `Line`.

The `textColor` attribute defines the color of the text. The text is drawn with transparent background and no border around the text (and without outline). The contents inherited from `FilledShape` is deprecated, but kept for compatibility reasons. The default value for `horizontalAlignment` is deprecated. Having a zero size for the `extent` is deprecated and is handled as if upper part is moved up an appropriate amount.

[Example:

```

connect(controlBus.axisControlBus1, axis1.axisControlBus) annotation (
  Text(string="%first", index=-1, extent=[-6,3; -6,7]),

```

```

Line(points
  ={{-80, -10},{-80, -14.5},{-79, -14.5},{-79, -17},{-65, -17},{-65, -65},{-25, -65}})
);
    
```

Draws a connection line and adds the text axisControlBus1 ending at $\{-6, 3\} + \{-25, -65\}$ and 4 vertical units of space for the text. Using a height of zero, such as `extent = [-6, 3; -6, 3]` is deprecated, but gives similar result.]

18.6.5 Graphical primitives

This section describes the graphical primitives that can be used to define the graphical objects in an annotation.

18.6.5.1 Line

A line is specified as follows:

```

record Line
  extends GraphicItem;
  Point points[:];
  Color color = Black;
  LinePattern pattern = LinePattern.Solid;
  DrawingUnit thickness = 0.25;
  Arrow arrow[2] = {Arrow.None, Arrow.None} "{start arrow, end arrow}";
  DrawingUnit arrowSize = 3;
  Smooth smooth = Smooth.None "Spline";
end Line;
    
```

Note that the Line primitive is also used to specify the graphical representation of a connection.

For arrows:

- The arrow is drawn with an aspect ratio of 1/3 for each arrow half, i.e., if the arrow-head is 3 mm long an arrow with `Half` will extend 1 mm from the mid-line and with `Open` or `Filled` extend 1 mm to each side, in total making the base 2 mm wide.
- The `arrowSize` gives the width of the arrow (including the imagined other half for `Half`) so that `lineThickness = 10` and `arrowSize = 10` will touch at the outer parts.
- All arrow variants overlap for overlapping lines.
- The lines for the `Open` and `Half` variants are drawn with `lineThickness`.

18.6.5.2 Polygon

A polygon is specified as follows:

```

record Polygon
  extends GraphicItem;
  extends FilledShape;
  Point points[:];
  Smooth smooth = Smooth.None "Spline outline";
end Polygon;
    
```

The polygon is automatically closed, if the first and the last points are not identical.

18.6.5.3 Rectangle

A rectangle is specified as follows:

```

record Rectangle
  extends GraphicItem;
  extends FilledShape;
  BorderPattern borderPattern = BorderPattern.None;
  Extent extent;
  DrawingUnit radius = 0 "Corner radius";
    
```

```
end Rectangle;
```

The `extent` attribute specifies the bounding box of the rectangle. If the `radius` attribute is specified, the rectangle is drawn with rounded corners of the given radius.

18.6.5.4 Ellipse

An ellipse is specified as follows:

```
record Ellipse
  extends GraphicItem;
  extends FilledShape;
  Extent extent;
  Real startAngle(quantity="angle", unit="deg")=0;
  Real endAngle(quantity="angle", unit="deg")=360;
  EllipseClosure closure = if startAngle == 0 and endAngle == 360
  then EllipseClosure.Chord
  else EllipseClosure.Radial;
end Ellipse;
```

The `extent` attribute specifies the bounding box of the ellipse.

Partial ellipses can be drawn using the `startAngle` and `endAngle` attributes. These specify the endpoints of the arc prior to the stretch and rotate operations. The arc is drawn counter-clockwise from `startAngle` to `endAngle`, where `startAngle` and `endAngle` are defined counter-clockwise from 3 o'clock (the positive x-axis).

The closure attribute specifies whether the endpoints specified by `startAngle` and `endAngle` are to be joined by lines to the center of the extent (`closure = EllipseClosure.Radial`), joined by a single straight line between the end points (`closure = EllipseClosure.Chord`), or left unconnected (`closure = EllipseClosure.None`). In the latter case, the ellipse is treated as an open curve instead of a closed shape, and the `fillPattern` and `fillColor` are not applied (if present, they are ignored).

The default closure is `EllipseClosure.Chord` when `startAngle` is 0 and `endAngle` is 360, or `EllipseClosure.Radial` otherwise.

[The default for a closed ellipse is not `EllipseClosure.None`, since that would result in `fillColor` and `fillPattern` being ignored, making it impossible to draw a filled ellipse. `EllipseClosure.Chord` is equivalent in this case, since the chord will be of zero length.]

18.6.5.5 Text

A text string is specified as follows:

```
record Text
  extends GraphicItem;
  extends FilledShape;
  Extent extent;
  String textString;
  Real fontSize = 0 "unit pt";
  String fontName;
  TextStyle textStyle[:];
  Color textColor=lineColor;
  TextAlignment horizontalAlignment = TextAlignment.Center;
end Text;
```

The `textColor` attribute defines the color of the text. The text is drawn with transparent background and no border around the text (and without outline). The contents inherited from `FilledShape` is deprecated, but kept for compatibility reasons.

There are a number of common macros that can be used in the text, and they should be replaced when displaying the text as follows (in order such that the earliest ones have precedence, and using the longest sequence of identifier characters – alphanumeric and underscore):

- %% replaced by %

- `%name` replaced by the name of the component (i.e., the identifier for it in the enclosing class).
- `%class` replaced by the name of the class (only the last part of the hierarchical name).
- `%par` and `{%par}` replaced by the value of the parameter `par`. If the value is numeric, tools shall display the value with `displayUnit`, formatted according to bipm-specification. E.g., for

```
parameter Real t(unit="s", displayUnit="ms") = 0.1
```

tools shall display *100 ms*. The intent is that the text is easily readable, thus if `par` is of an enumeration type, replace `%par` by the item name, not by the full name.

[*Example: If `par` = "Modelica.Blocks.Types.Enumeration.Periodic", then `%par` should be displayed as Periodic.*]

The form `{%par}` allows component-references and is required for quoted identifiers, and can be directly followed by a letter. Thus `{w}x{h}` gives the value of `w` directly followed by `x` and the value of `h`, while `wxh` gives the value of the parameter `wxh`. If the parameter does not exist it is an error.

The style attribute `fontSize` specifies the font size. If the `fontSize` attribute is 0 the text is scaled to fit its extent. Otherwise, the size specifies the absolute size. The text is vertically centered in the extent.

If the `extent` specifies a box with zero width and positive height the height is used as height for the text (unless `fontSize` attribute is non-zero – which specifies the absolute size), and the text is not truncated (the `horizontalAlignment` is still used in this case).

[*A zero-width extent is convenient for handling texts where the width is unknown.*]

If the string `fontName` is empty, the tool may choose a font. The font names "serif", "sans-serif", and "monospace" shall be recognized. If possible the correct font should be used – otherwise a reasonable match, or treat as if `fontName` was empty.

The style attribute `textStyle` specifies variations of the font.

18.6.5.6 Bitmap

A bitmap image is specified as follows:

```
record Bitmap
  extends GraphicItem;
  Extent extent;
  String fileName "Name of bitmap file";
  String imageSource "Base64 representation of bitmap";
end Bitmap;
```

The `Bitmap` primitive renders a graphical bitmap image. The data of the image can either be stored on an external file or in the annotation itself. The image is scaled to fit the extent. Given an extent `{x1, y1}, {x2, y2}`, `x2 < x1` defines horizontal flipping and `y2 < y1` defines vertical flipping around the center of the object.

The graphical operations are applied in the order: scaling, flipping and rotation.

When the attribute `fileName` is specified, the string refers to an external file containing image data. The mapping from the string to the file is specified for some URIs in section 13.5. The supported file formats include PNG, BMP, JPEG, and SVG.

When the attribute `imageSource` is specified, the string contains the image data, and the image format is determined based on the contents. The image is represented as a Base64 encoding of the image file format (see RFC 4648, <http://tools.ietf.org/html/rfc4648>).

The image is uniformly scaled (preserving the aspect ratio) so it exactly fits within the extent (touching the extent along one axis). The center of the image is positioned at the center of the extent.

18.6.6 Variable Graphics and Schematic Animation

Any value (coordinates, color, text, etc.) in graphical annotations can be dependent on class variables using the `DynamicSelect` expression. `DynamicSelect` has the syntax of a function call with two arguments, where the first argument specifies the value of the editing state and the second argument the value of the non-editing state. The first argument must be a literal expression. The second argument may contain references to variables to enable a dynamic behavior.

[Example: The level of a tank could be animated by a rectangle expanding in vertical direction and its color depending on a variable overflow:

```

annotation (Icon(graphics =
  {Rectangle(
    extent =
      DynamicSelect({{0, 0}, {20, 20}},
                   {{0, 0}, {20, level}}),
    fillColor =
      DynamicSelect({0, 0, 255},
                   if overflow then {255, 0, 0} else {0, 0, 255})
  )}
));

```

18.6.7 User input

It is possible to interactively modify variables during a simulation. The variables may either be parameters, discrete-time variables or states. New numeric values can be given, a mouse click can change a `Boolean` variable or a mouse movement can change a `Real` variable. Input fields may be associated with a `GraphicItem` or a component as an array named `interaction`. The `interaction` array may occur as an attribute of a graphic primitive, an attribute of a component annotation or as an attribute of the layer annotation of a class.

18.6.7.1 Mouse input

A `Boolean` variable can be changed when the cursor is held over a graphical item or component and the selection button is pressed if the interaction annotation contains `OnMouseDownSetBoolean`:

```

record OnMouseDownSetBoolean
  Boolean variable "Name of variable to change when mouse button pressed";
  Boolean value "Assigned value";
end OnMouseDownSetBoolean;

```

[Example: A button can be represented by a rectangle changing color depending on a `Boolean` variable on and toggles the variable when the rectangle is clicked on:

```

annotation (Icon(
  graphics = {Rectangle(
    extent = [0, 0; 20, 20],
    fillColor = if on then {255, 0, 0} else {0, 0, 255}}),
  interaction = {OnMouseDownSetBoolean(on, not on)});

```

In a similar way, a variable can be changed when the mouse button is *released*:

```

record OnMouseUpSetBoolean
  Boolean variable "Name of variable to change when mouse button released";
  Boolean value "Assigned value";
end OnMouseUpSetBoolean;

```

Note that several interaction objects can be associated with the same graphical item or component.

[Example:

```
interaction = {OnMouseDownSetBoolean(on, true),
              OnMouseUpSetBoolean(on, false)}
```

]

The `OnMouseMoveXSetReal` interaction object sets the variable to the position of the cursor in X direction in the local coordinate system mapped to the interval defined by the `minValue` and `maxValue` attributes.

```
record OnMouseMoveXSetReal
  Real xVariable "Name of variable to change when cursor moved in x direction";
  Real minValue;
  Real maxValue;
end OnMouseMoveXSetReal;
```

The `OnMouseMoveYSetReal` interaction object works in a corresponding way as the `OnMouseMoveXSetReal` object but in the Y direction.

```
record OnMouseMoveYSetReal
  Real yVariable "Name of variable to change when cursor moved in y direction";
  Real minValue;
  Real maxValue;
end OnMouseMoveYSetReal;
```

18.6.7.2 Edit input

The `OnMouseDownEditInteger` interaction object presents an input field when the graphical item or component is clicked on. The field shows the actual value of the variable and allows changing the value. If a too small or too large value according to the `min` and `max` parameter values of the variable is given, the input is rejected.

```
record OnMouseDownEditInteger
  Integer variable "Name of variable to change";
end OnMouseDownEditInteger;
```

The `OnMouseDownEditReal` interaction object presents an input field when the graphical item or component is clicked on. The field shows the actual value of the variable and allows changing the value. If a too small or too large value according to the `min` and `max` parameter values of the variable is given, the input is rejected.

```
record OnMouseDownEditReal
  Real variable "Name of variable to change";
end OnMouseDownEditReal;
```

The `OnMouseDownEditString` interaction object presents an input field when the graphical item or component is clicked on. The field shows the actual value of the variable and allows changing the value.

```
record OnMouseDownEditString
  String variable "Name of variable to change";
end OnMouseDownEditString;
```

18.7 Annotations for the Graphical User Interface

This section describes the annotations that are used to define properties of the graphical user interface.

```
preferred-view-annotation:
  annotation (" preferredView "=" ("info" | "diagram" | "text") ")")
```

The `preferredView` annotation defines the default view when selecting the class. `info` means info layer, i.e., the documentation of the class, `diagram` means diagram layer and `text` means the Modelica text layer.

```
documentation-class-annotation:
  annotation (" DocumentationClass "=" true ")")
```

Only allowed as class annotation on any kind of class and implies that this class and all classes within it are treated as having the annotation `preferredView = "info"`. If the annotation `preferredView` is explicitly set for a class, it has precedence over a `DocumentationClass` annotation.

[A tool may display such classes in special ways. For example, the description texts of the classes might be displayed instead of the class names, and if no icon is defined, a special information default icon may be displayed in the package browser.]

```
annotation(defaultComponentName = "name")
```

When creating a component of the given class, the recommended component name is *name*.

```
annotation(defaultComponentPrefixes = "prefixes")
```

When creating a component, it is recommended to generate a declaration of the form

```
type-prefix type-specifier component-declaration
```

The following prefixes may be included in the string `prefixes`: `inner`, `outer`, `replaceable`, `constant`, `parameter`, `discrete`.

[In combination with `defaultComponentName` it can be used to make it easy for users to create `inner` components matching the `outer` declarations; see also example below. If the prefixes contain `inner` or `outer` and the default name cannot be used (e.g., since it is already in use) it is recommended to give a diagnostic.]

```
annotation(missingInnerMessage = "message")
```

When an `outer` component of the class does not have a corresponding `inner` component, the literal string message may be used as part of a diagnostic message (together with appropriate context), see section 5.4.

[Example:

```
model World
  ...
  annotation(defaultComponentName = "world",
    defaultComponentPrefixes = "inner replaceable",
    missingInnerMessage = "The World object is missing");
end World;
```

When an instance of `model World` is dragged in to the diagram layer, the following declaration is generated:

```
inner replaceable World world;
```

]

A simple type or component of a simple type may have:

```
annotation(absoluteValue=false);
```

If `false`, then the variable defines a relative quantity, and if `true` an absolute quantity.

[When converting between units (in the user-interface for plotting and entering parameters), the offset must be ignored for a variable defined with annotation `absoluteValue = false`. This annotation is used in the Modelica Standard Library, for example in `Modelica.Units.SI` for the type definition `TemperatureDifference`.]

A model or block definition may contain:

```
annotation(defaultConnectionStructurallyInconsistent=true)
```

If `true`, it is stated that a default connection will result in a structurally inconsistent model or block¹. A "default connection" is constructed by instantiating the respective `model` or `block` and for every input

¹For the precise definition of *structurally inconsistent*, see Pantelides (1988).

u providing an equation $0 = f(u)$, and for every (potential, flow) pair of the form (v, i) , providing an equation of the form $0 = f(v, i)$.

[It is useful to check all models/blocks of a Modelica package in a simple way. One check is to default connect every model/block and to check whether the resulting class is structurally consistent (which is a stronger requirement than being balanced). It is rarely needed; but is for example used in Modelica.Blocks.Math.InverseBlockConstraints, in order to prevent a wrong error message. Additionally, when a user defined model is structurally inconsistent, a tool should try to pinpoint in which class the error is present. This annotation avoids then to show a wrong error message.]

A class may have the following annotation:

```
annotation(obsolete = "message");
```

It indicates that the class ideally should not be used anymore and gives a message indicating the recommended action. This annotation is not inherited, the assumption is that if a class uses an obsolete class (as a base-class or as the class of one of the components) that shall be updated – ideally without impacting users of the class. If that is not possible the current class can have also have an obsolete annotation.

A component declaration may have the following annotation:

```
annotation(unassignedMessage = "message");
```

When the variable to which this annotation is attached in the declaration cannot be computed due to the structure of the equations, the string "message" can be used as a diagnostic message.

[When using BLT partitioning, this means if a variable a or one of its aliases $b = a$ or $b = -a$ cannot be assigned, the message is displayed. This annotation is used to provide library specific error messages.]

[Example:

```
connector Frame "Frame of a mechanical system"
...
flow Modelica.Units.SI.Force f[3]
annotation(unassignedMessage =
  "All Forces cannot be uniquely calculated. The reason could be that the
  mechanism contains a planar loop or that joints constrain the same motion.
  For planar loops, use in one revolute joint per loop the option
  PlanarCutJoint=true in the Advanced menu.
");
end Frame;
```

]

A component declaration or a short replaceable class definition may have the following annotation:

```
annotation(Dialog(enable = true,
  tab = "General",
  group = "",
  showStartAttribute = false,
  colorSelector = false,
  groupImage="modelica://MyPackage/Resources/Images/load.png",
  connectorSizing = false));
```

For a short replaceable class definition only the fields tab, group, enable and groupImage are allowed.

The annotations tab and group define the placement of the component or of variables in a dialog with optional tab and group specification, where the empty string (default) means tool-specific group. The idea is that a tool may as default place parameters in the group "Parameters" in the tab "General", but add e.g., variables with showStartAttribute=true to another group. If enable = false, the input field may be disabled and no input can be given. If showStartAttribute = true the dialog should allow the user to set the start-value and the fixed attribute for the variable instead of the value-attribute.

[The showStartAttribute = true is primarily intended for non-parameter values and avoids introducing a separate parameter for the start-value of the variable.]

The order of parameters within each group and the order of the groups and tabs are according to the declaration order, where inherited elements are added at the place of the extends.

If `colorSelector = true`, it indicates that an rgb-value selector can be presented for a vector of three elements and generate values 0..255 (the annotation should be useable both for vectors of `Integer` and `Real`).

The annotation `groupImage` references an image using an URI (see section 13.5), and the image is intended to be shown together with the parameter-group (only one image per group is supported). Disabling the input field will not disable the image.

The background of the `groupImage` and any image used in HTML-documentation is recommended to be transparent (intended to be a light color) or white.

The value of the `connectorSizing` annotation must be a literal `false` or `true` value. If `connectorSizing = false`, this annotation has no effect. If `connectorSizing = true`, the corresponding variable must be declared with the `parameter` prefix, must be a subtype of a scalar `Integer` and must have a literal default value of zero.

[The reason why `connectorSizing` must be given a literal value is that if the value is an expression, the `connectorSizing` functionality is conditional and this will then lead easily to wrong models.]

The default value of the variable must be zero since this annotation is designed for a parameter that is used as vector dimension, and the dimension of the vector should be zero when the component is dragged or redeclared. Furthermore, when a tool does not support the `connectorSizing` annotation, dragging will still result in a correct model.]

If `connectorSizing = true`, a tool may set the parameter value in a modifier automatically, if used as dimension size of a vector of connectors.

[The `connectorSizing` annotation is used in cases where connections to a vector of connectors shall be made and a new connection requires to resize the vector and to connect to the new index (unary connections). The annotation allows a tool to perform these two actions in many cases automatically. This is, e.g., very useful for state machines and for certain components of fluid libraries.]

Annotation Dialog is defined as:

```

record Dialog
  parameter String tab = "General";
  parameter String group = "Parameters";
  parameter Boolean enable = true;
  parameter Boolean showStartAttribute = false;
  parameter Boolean colorSelector = false;
  parameter Selector loadSelector;
  parameter Selector saveSelector;
  parameter String groupImage = "";
  parameter Boolean connectorSizing = false;
end Dialog;

record Selector
  parameter String filter="";
  parameter String caption="";
end Selector;

```

A parameter dialog is a sequence of tabs with a sequence of groups inside them.

A `Selector` displays a file dialog to select a file. Setting `filter` only shows files that fulfill the given pattern defined by `text1 (*.ext1);;text2 (*.ext2)` to show only files with file extension `ext1` or `ext2` and displaying a description text `text1` and `text2`, respectively. Parameter `caption` is the text displayed in the dialog menu. Parameter `loadSelector` is used to select an existing file for reading, whereas parameter `saveSelector` is used to define a file for writing.

[Example:

```



model DialogDemo
  parameter Boolean b = true "Boolean parameter";

```

```

parameter Modelica.Units.SI.Length length "Real parameter with unit";
parameter Integer nInports=0
  annotation(Dialog(connectorSizing=true));
parameter Real r1 "Real parameter in Group 1"
  annotation(Dialog(group="Group 1"));
parameter Real r2 "Disabled Real parameter in Group 1"
  annotation(Dialog(group="Group 1",enable = not b));
parameter Real r3 "Real parameter in Tab 1"
  annotation(Dialog(tab="Tab 1"));
parameter Real r4 "Real parameter in Tab 1 and Group 2"
  annotation(Dialog(tab="Tab 1", group="Group 2"));
StepIn stepIn[nInports];
...
end DialogDemo;
    
```

When clicking on an instance of model `DialogDemo`, a menu pops up that may have the following layout (other layouts are also possible, this is vendor specific). Note, parameter `nInports` is not present in the menu since it has the `connectorSizing` annotation and therefore it should not be modified by the user (an alternative is to show parameter `nInports` in the menu but with disabled input field):



]

[The following part is non-normative text and describes a useful way to handle the `connectorSizing` annotation in a tool (still a tool may use another strategy and/or may handle other cases than described below). The recommended rules are clarified at hand of the following example which represents a connector and a model from the `Modelica.StateGraph` library:

```

connector StepIn // Only 1:1 connections are possible since input used
  output Boolean occupied;
  input Boolean set;
end StepIn;

model Step
  // nIn cannot be set in the parameter dialog (but maybe shown)
  parameter Integer nIn=0 annotation(Dialog(connectorSizing=true));
  StepIn inPorts[nIn];
  ...
end Step;
    
```

If the parameter is used as dimension size of a vector of connectors, it is automatically updated according to the following rules:

1. If a new connection line is drawn between one outside and one inside vector of connectors both dimensioned with (`connectorSizing`) parameters, a connection between the two vectors is performed and the (`connectorSizing`) parameter is propagated from connector to component. Other types of outside connections do not lead to an automatic update of a (`connectorSizing`) parameter. Example: Assume there is a connector `inPorts` and a component `step1`:

```

parameter Integer nIn=0 annotation(Dialog(connectorSizing=true));
StepIn inPorts[nIn];
Step step1(nIn=0);
    
```

Drawing a connection line between connectors `inPorts` and `step1.inPorts` results in:

```

parameter Integer nIn=0 annotation(Dialog(connectorSizing=true));
StepIn inPorts[nIn];
    
```

```

Step step1(nIn=nIn); // nIn=0 changed to nIn=nIn
equation
  connect(inPorts, step1.inPorts); // new connect equation
    
```

2. If a connection line is deleted between one outside and one inside vector of connectors both dimensioned with (`connectorSizing`) parameters, the connect equation is removed and the (`connectorSizing`) parameter of the component is set to zero or the modifier is removed. Example: Assume the connection line in (3) is removed. This results in:

```

parameter Integer nIn=0 annotation(Dialog(connectorSizing=true));
StepIn inPorts[nIn];
Step step1; // modifier nIn=nIn is removed
    
```

3. If a new connection line is drawn to an inside connector with `connectorSizing` and case 1 does not apply then, the parameter is incremented by one and the connection is performed for the new highest index. Example: Assume that 3 connections are present and a new connection is performed. The result is:

```

Step step1(nIn=4); // index changed from nIn=3 to nIn=4
equation
  connect(.., step1.inPorts[4]); // new connect equation
    
```

In some applications, like state machines, the vector index is used as a priority, e.g., to define which transition is firing if several transitions become active at the same time instant. It is then not sufficient to only provide a mechanism to always connect to the last index. Instead, some mechanism to select an index conveniently should be provided.

4. If a connection line is deleted to an inside connector with `connectorSizing` and case 2 does not apply then, then the (`connectorSizing`) parameter is decremented by one and all connections with index above the deleted connection index are also decremented by one. Example: Assume there are 4 connections:

```

Step step1(nIn=4);
equation
  connect(a1, step1.inPorts[1]);
  connect(a2, step1.inPorts[2]);
  connect(a3, step1.inPorts[3]);
  connect(a4, step1.inPorts[4]);
    
```

and the connection from `a2` to `step1.inPorts[2]` is deleted. This results in

```

Step step1(nIn=3);
equation
  connect(a1, step1.inPorts[1]);
  connect(a3, step1.inPorts[2]);
  connect(a4, step1.inPorts[3]);
    
```

These rules also apply if the connectors and/or components are defined in superclass. Example: Assume that `step1` is defined in superclass `CompositeStep` with 3 connections, and a new connection is performed in a subclass. The result is:

```

extends CompositeStep(step1(nIn=4)); // new modifier nIn=4
equation
  connect(.., step1.inPorts[4]); // new connect equation
    
```

]

18.8 Annotations for Version Handling

A top-level package or model can specify the version of top-level classes it uses, its own version number, and if possible how to convert from previous versions. This can be used by a tool to guarantee that consistent versions are used, and if possible to upgrade usage from an earlier version to a current one.

18.8.1 Version Numbering

Version numbers are of the forms:

- Main release versions: `"" UNSIGNED-INTEGER { "." UNSIGNED-INTEGER } ""`
Example: "2.1"
- Pre-release versions: `"" UNSIGNED-INTEGER { "." UNSIGNED-INTEGER } " " {S-CHAR} ""`
Example: "2.1 Beta 1"
- Un-ordered versions: `"" NON-DIGIT {S-CHAR} ""`
Example: "Test 1"

The main release versions are ordered using the hierarchical numerical names, and follow the corresponding pre-release versions. The pre-release versions of the same main release version are internally ordered alphabetically.

18.8.2 Version Handling

In a top-level class, the version number and the dependency to earlier versions of this class are defined using one or more of the following annotations:

- `version = CURRENT-VERSION-NUMBER`
Defines the version number of the model or package. All classes within this top-level class have this version number.
- `conversion(noneFromVersion = VERSION-NUMBER)`
Defines that user models using the `VERSION-NUMBER` can be upgraded to the `CURRENT-VERSION-NUMBER` of the current class without any changes.
- `conversion(from(version = Versions, [to=VERSION-NUMBER,] Convert))`
where `Versions` is `VERSION-NUMBER | {VERSION-NUMBER, VERSION-NUMBER, ...}` and `Convert` is `script="..." | change={conversionRule(), ..., conversionRule()}`
Defines that user models using the `VERSION-NUMBER` or any of the given `VERSION-NUMBER` can be upgraded to the given `VERSION-NUMBER` (if the `to`-tag is missing this is the `CURRENT-VERSION-NUMBER`) of the current class by applying the given conversion rules. The script consists of an unordered sequence of `conversionRule()`; (and optionally Modelica comments). The `conversionRule` functions are defined in section 18.8.2.1.
[The to-tag is added for clarity and optionally allows a tool to convert in multiple steps.]
- `uses(IDENT (version = VERSION-NUMBER [, versionBuild=INTEGER] [, dateModified=STRING]))`
Defines that classes within this top-level class uses version `VERSION-NUMBER` of classes within the top-level class `IDENT`.

The annotations `uses` and `conversion` may contain several different sub-entries.

[Example:

```

package Modelica
...
  annotation(version="3.1",
    conversion(noneFromVersion="3.1 Beta 1",
      noneFromVersion="3.1 Beta 2",
      from(version={"2.1", "2.2", "2.2.1"},
        script="convertTo3.mos"),
      from(version="1.5",
        script="convertFromModelica1_5.mos")));
end Modelica;

model A
...
  annotation(version="1.0",
    uses(Modelica(version="1.5")));
end A;

```

```

model B
  ...
  annotation(uses(Modelica(version="3.1 Beta 1")));
end B;

```

In this example the model A uses an older version of the Modelica library and can be upgraded using the given script, and model B uses an older version of the Modelica library but no changes are required when upgrading.]

18.8.2.1 Conversion rules

There are a number of functions: `convertClass`, `convertClassIf`, `convertElement`, `convertModifiers`, `convertMessage` defined as follows. The calls of these functions do not directly convert, instead they define conversion rules as below. The order between the function calls does not matter, instead the longer paths (in terms number of hierarchical names) are used first as indicated below, and it is an error if there are any ambiguities.

The conversion should generate correct Modelica models using the new version of the library corresponding to the old version.

[Whenever possible tools should preserve the original style of the model, e.g. use of imports.]

These functions can be called with literal strings or array of strings and vectorize according to section 12.4.6.

All of these convert-functions only use inheritance among user models, and not in the library that is used for the conversion – thus conversions of base-classes will require multiple conversion-calls; this ensures that the conversion is independent of the new library structure. The name of the class used as argument to `convertElement` and `convertModifiers` is similarly the old name of the class, i.e. the name before it is possibly converted by `convertClass`.

[Specifying conversions using the old name of a class allows the conversion to be done without access to the old version of the library (by suitable modifications of the lookup). Another alternative is to use the old version of the library during the conversion.]

convertClass("OldClass", "NewClass") Convert class `OldClass` to `NewClass`.

Match longer path first, so if converting both `A` to `C` and `A.B` to `D` then `A.F` is converted to `C.F` and `A.B.E` to `D.E`. This is considered before `convertMessage` for the same `OldClass`.

[Example: Consider the following as part of a conversion script:

```

convertClass("Modelica.SIunits", "Modelica.Units.SI")
convertClass("Modelica.SIunits.Icons", "Modelica.Units.Icons")

```

This ensures that for example `Modelica.SIunits.Length` is converted to `Modelica.Units.SI.Length` and `Modelica.SIunits.Icons` is converted to `Modelica.SIunits.Icons`.]

convertClassIf("OldClass", "oldElement", "whenValue", "NewClass") Convert class `OldClass` to `NewClass` if the literal modifier for `oldElement` has the value `whenValue`, and also remove the modifier for `oldElement`.

These are considered before `convertClass` and `convertMessage` for the same `OldClass`.

The old element should be of a `Boolean`, `Integer`, `String`, or enumeration type and the match is based on the literal value of the modifier. For string elements the value argument to `convertClassIf` shall be up-quoted, e.g. `"\My String\"`, and for enumeration literals only the enumeration literal part of the old value matters, e.g., `red` for `"Colors.red"`.

convertElement("OldClass", "OldName", "NewName") In `OldClass`, convert element `OldName` to `NewName`. Both `OldName` and `NewName` normally refer to components, but they may also refer to class-parameters, or hierarchical names. For hierarchical names, the longest match is used first.

For replaceable classes in packages (and replaceable classes in other classes) `convertElement` shall be used if the class is renamed within the package (or class), whereas `convertClass` shall only be used if the class is placed outside of the package (or class).

[The latter case indicates a problem with overuse of replaceable classes in the previous design of the library.]

[Example: Consider the following as part of a conversion script:

```
convertElement({"Modelica.Mechanics.MultiBody.World",
              "Modelica.Mechanics.MultiBody.World.gravityAcceleration"},
             "mue", "mu")
```

This implies that

```
Modelica.Mechanics.MultiBody.World world(mue=2);
function f=Modelica.Mechanics.MultiBody.World.gravityAcceleration(mue=4);
```

is converted to:

```
Modelica.Mechanics.MultiBody.World world(mu=2);
function f=Modelica.Mechanics.MultiBody.World.gravityAcceleration(mu=4);
```

]

convertModifiers

```
convertModifiers("OldClass",
               {"OldModifier1=default1", "OldModifier2=default2", ...},
               {"NewModifier1=...%OldModifier2%...", "NewModifier2=...", ...}
               [, simplify=true])
```

Normal case; if any modifier among `OldModifier` exist then replace all of them with the list of `NewModifiers`. The `...%OldModifier2%...` indicate an expression that may involve the values of the old modifiers (tools are responsible for adding parenthesis if needed). The lists of old and new modifiers can have different lengths. The defaults (if present) are used if there are multiple `OldModifier` and not all are set in the component instance. The defaults are optional if there is at most one `OldModifier` element, and should otherwise be provided.

If `simplify` is specified and true then perform obvious simplifications to clean up the new modifier; otherwise leave as is.

[Note: `simplify` is primarily intended for converting enumerations and emulated enumerations that naturally lead to large nested if-expressions. The simplifications may also simplify parts of the original expression.]

If the modifiers contain literal string values they must be quoted.

Behaviour in unusual cases:

- if `NewModifier` list is empty then the modifier is just removed
- If `OldModifier` list is empty it is added for all uses of the class
- If `OldModifieri` is `cardinality(a) = 0` the conversion will only be applied for a component `comp` if there are no inside connections to `comp.a`. This can be combined with other modifiers that are handled in the usual way.
- If `OldModifieri` is `cardinality(a) = 1` the conversion will only be applied for a component `comp` if there are any inside connections to `comp.a`.

The converted modifiers and existing modifiers are merged such that the existing modifiers take precedence over the result of `convertModifiers`. A diagnostics is recommended if this merging removes some modifiers unless those modifiers are identical or it is the special case of an empty `OldModifier` list.

[This can be used to handle the case where the default value was changed.]

Converting modifiers with cardinality is used to remove the deprecated operator `cardinality` from model libraries, and replace tests on cardinality in models by parameters explicitly enabling the different cases. The case where the old class is used as a base-class, and there exist outside connections to `a`, and there is `convertModifiers` involving the cardinality of `a` is not handled.

[Having a parameter for explicitly enabling the different cases means that instead of model A internally testing if its connector B is connected, there will be a parameter for enabling connector B, and the conversion ensures that each component of model A will have this parameter set accordingly.]

In case a parameter is simply renamed it is preferable to use `convertElement`, since that also handles e.g. binding equations using the parameter.]

[Example: The conversion

```

convertClass("Modelica.Thermal.FluidHeatFlow.Components.IsolatedPipe",
            "Modelica.Thermal.FluidHeatFlow.Components.Pipe")
convertModifiers({"Modelica.Thermal.FluidHeatFlow.Components.IsolatedPipe"},
               fill("",0), {"useHeatPort=false"})

convertClass("Modelica.StateGraph.Temporary.NumericValue",
            "Modelica.Blocks.Interaction.Show.RealValue")
convertModifiers("Modelica.StateGraph.Temporary.NumericValue",
               {"Value"}, {"number=%Value%"})
convertModifiers("Modelica.StateGraph.Temporary.NumericValue",
               {"hideConnector"}, {"use_numberPort=not %hideConnector%})

convertModifiers("Modelica.Blocks.Math.LinearDependency",
               {"y0=0", "k1=0", "k2=0"}, {"y0=%y0%", "k1=%y0%*%k1%", "k2=%y0%*%k2%"},
               true)
convertClass(
    "Modelica.Electrical.Machines.BasicMachines.QuasiStationaryDCMachines",
    "Modelica.Electrical.Machines.BasicMachines.QuasiStaticDCMachines")
convertElement("Modelica.Electrical.Machines.Interfaces.PartialBasicDCMachine",
               "quasiStationary", "quasiStatic")
convertElement("Modelica.Electrical.Machines.BasicMachines.
    QuasiStationaryDCMachines.DC_ElectricalExcited",
               "quasiStationary", "quasiStatic")
    
```

converts

```

Modelica.Thermal.FluidHeatFlow.Components.IsolatedPipe pipe1;
Modelica.StateGraph.Temporary.NumericValue tempValue(Value=10,
    hideConnector=true);
Modelica.Blocks.Math.LinearDependency linearDep(y0=2, k2=1);
model A
    import Modelica.Electrical.Machines.BasicMachines;
    extends BasicMachines.QuasiStationaryDCMachines.DC_ElectricalExcited;
end A;
model B
    extends A;
    Boolean b=quasiStationary;
end B;
    
```

to

```

Modelica.Thermal.FluidHeatFlow.Components.Pipe pipe1(useHeatPort=false);
Modelica.Blocks.Interaction.Show.RealValue(number=10, use_numberPort=not true);
Modelica.Blocks.Math.LinearDependency linearDep(y0=2, k1=0, k2=2);
model A
    import Modelica.Electrical.Machines.BasicMachines;
    extends BasicMachines.QuasiStaticDCMachines.DC_ElectricalExcited;
end A;
model B
    extends A;
    Boolean b=a.quasiStatic;
    
```

```
end B;
```

The `convertElement` call for `DC_ElectricalExcited` is needed to avoid relying on base-classes in the original library where `DC_ElectricalExcited` inherits from `PartialBasicDCMachine`. However, the inheritance among the models to convert (in this case `B` inherits from `A`) should be handled.]

`convertMessage("OldClass", "Failed Message");` For any use of `OldClass` (or element of `OldClass`) report that conversion could not be applied with the given message.

[This may be useful if there is no possibility to convert a specific class. An alternative is to construct `ObsoleteLibraryA` for problematic cases, which may be more work but allows users to directly run the models after the conversion and later convert them.]

`convertMessage("OldClass", "Failed Message", "oldElement");` For any use of `oldElement` in `OldClass` report that conversion could not be applied with the given message.

[This is useful if there is no possibility to convert a specific parameter (or other element), especially if it rarely modified. If the parameter had no impact on the model it can be removed using `convertModifiers`, see section 18.8.2.1.]

18.8.3 Mapping of Versions to File System

A top-level class, `IDENT`, with version `VERSION-NUMBER` can be stored in one of the following ways in a directory given in the `MODELICAPATH`:

- The file `IDENT ".mo"`
Example: `Modelica.mo`
- The file `IDENT " " VERSION-NUMBER ".mo"`
Example: `Modelica 2.1.mo`
- The directory `IDENT` with the file `package.mo` directly inside it
Example: `Modelica/package.mo`
- The directory `IDENT " " VERSION-NUMBER` with the file `package.mo` directly inside it
Example: `Modelica 2.1/package.mo`

This allows a tool to access multiple versions of the same package.

18.8.4 Version Date and Build Information

Besides version information, a top level class can have additionally the following top-level annotations to specify associated information to the version number:

```
String versionDate    "UTC date of first version build (in format: YYYY-MM-DD)";
Integer versionBuild  "Larger number is a more recent maintenance update";
String dateModified   "UTC date and time of the latest change to the package
                        in the following format (with one space between date
                        and time): YYYY-MM-DD hh:mm:ssZ";
String revisionId     "Revision identifier of the version management system used
                        to manage this library. It marks the latest submitted
                        change to any file belonging to the package";
```

[Example:

```
package Modelica
...
annotation(version = "3.0.1",
  versionDate = "2008-04-10",
  versionBuild = 4,
  dateModified = "2009-02-15 16:33:14Z",
  revisionId = "$Id:: package.mo 2566 2009-05-26 13:25:54Z #");
end Modelica;
```

```

model M1
  annotation(uses(Modelica(version = "3.0.1"))); // Common case
end M1

model M2
  annotation(uses(Modelica(version = "3.0.1", versionBuild = 4)));
end M2

```

The meanings of these annotations are:

- `version` is the version number of the released library, see section 18.8.2.
- `versionDate` is the date in UTC format (according to ISO 8601) when the library was released. This string is updated by the library author to correspond with the version number.
- `versionBuild` is the optional build number of the library. When a new version is released `versionBuild` should be omitted or `versionBuild = 1`. There might be bug fixes to the library that do not justify a new library version. Such maintenance changes are called a *build* release of the library. For every new maintenance change, the `versionBuild` number is increased. A `versionBuild` number *A* that is higher than `versionBuild` number *B*, is a newer release of the library. There are no conversions between the same versions with different build numbers.

Two releases of a library with the same `version` but different `versionBuild` are in general assumed to be compatible. In special cases, the `uses` clause of a model may specify `versionBuild` and/or `dateModified`. In such a case the tool is expected to give a warning if there is a mismatch between library and model.

- `dateModified` is the UTC date and time (according to ISO 8601) of the last modification of the package.
[The intention is that a Modelica tool updates this annotation whenever the package or part of it was modified and is saved on persistent storage (like file or database system).]
- `revisionId` is a tool specific revision identifier possibly generated by a source code management system (e.g. Subversion or CVS). This information exactly identifies the library source code in the source code management system.

The `versionBuild` and `dateModified` annotations can also be specified in the `uses` annotation (together with the version number).

[It is recommended that tools do not automatically store `versionBuild` and `dateModified` in the `uses` annotation.]

18.9 Annotations for Access Control to Protect Intellectual Property

This section presents annotations to define the protection and the licensing of packages. The goal is to unify basic mechanisms to control the access to a package in order to protect the intellectual property contained in it. This information is used to encrypt a package and bind it optionally to a particular target machine, and/or restrict the usage for a particular period of time.

[Protecting the intellectual property of a Modelica package is considerably more difficult than protecting code from a programming language. The reason is that a Modelica tool needs the model equations in order that it can process the equations symbolically, as needed for acausal modeling. Furthermore, if a Modelica tool generates C-code of the processed equations, this code is then potentially available for inspection by the user. Finally, the Modelica tool vendors have to be trusted, that they do not have a backdoor in their tools to store the (internally) decrypted classes in human readable format. The only way to protect against such misuse is legally binding warranties of the tool vendors.]

The intent of this section is to enable a library vendor to maintain one source version of their Modelica library that can be encrypted and used with several different Modelica tools, using different encryption formats.]

The following definitions relate to access control.

Definition 18.1. *Protection.* Define what parts of a class are visible. □

Definition 18.2. *Obfuscation.* Changing a Modelica class or generated code so that it is difficult to inspect by a user (e.g. by automatically renaming variables to non-meaningful names). □

Definition 18.3. *Encryption.* Encoding of a model or a package in a form so that the modeler cannot inspect any content of a class without an appropriate key. An encrypted package that has the `Protection` annotation is read-only; the way to modify it is to generate a new encrypted version. □

Definition 18.4. *Licensing.* Restrict the use of an encrypted package for particular users for a specified period of time. □

In this section annotations are defined for protection and licensing. Obfuscation and encryption are not standardized.

Protection and licensing are both defined inside the `Protection` annotation:

```
annotation(Protection(...));
```

18.9.1 Protection of Classes

A class may have the following annotations to define what parts of a class are visible, and only the parts explicitly listed as visible below can be accessed (if a class is encrypted and no `Protection` annotation is defined, the access annotation has the default value `Access.documentation`):

```
type Access =
  enumeration(hide, icon, documentation, diagram,
              nonPackageText, nonPackageDuplicate,
              packageText, packageDuplicate);
annotation(Protection(access = Access.documentation));
```

The items of the `Access` enumeration have the following meanings:

1. `Access.hide`
Do not show the class anywhere (it is not possible to inspect any part of the class).
2. `Access.icon`
The class can be instantiated and public parameter, constant, input, output variables as well as public connectors can be accessed, as well as the `Icon` annotation, as defined in section 18.6 (the declared information of these elements can be shown). Additionally, the class name and its description text can be accessed.
3. `Access.documentation`
Same as `Access.icon` and additionally the `Documentation` annotation (as defined in section 18.2) can be accessed. HTML-generation in the `Documentation` annotation is normally performed before encryption, but the generated HTML is intended to be used with the encrypted package. Thus the HTML-generation should use the same access as the encrypted version – even before encryption.
4. `Access.diagram`
Same as `Access.documentation` and additionally, the `Diagram` annotation, and all components and connect-equations that have a graphical annotation can be accessed.
5. `Access.nonPackageText`
Same as `Access.diagram` and additionally if it is not a package: the whole class definition can be accessed (but that text cannot be copied, i.e., you can see but not copy the source code).
6. `Access.nonPackageDuplicate`
Same as `Access.nonPackageText` and additionally if it is not a package: the class, or part of the class, can be copied.
7. `Access.packageText`
Same as `Access.diagram` (note: *not* including all rights of `Access.nonPackageDuplicate`) and additionally the whole class definition can be accessed (but that text cannot be copied, i.e., you can see but not copy the source code).

8. Access.packageDuplicate

Same as Access.packageText and additionally the class, or part of the class, can be copied.

The access annotation holds for the respective class and all classes that are hierarchically on a lower level, unless overridden by a Protection annotation with access. Overriding access=Access.hide or access=Access.packageDuplicate has no effect.

[Example: If the annotation is given on the top level of a package and at no other class in this package, then the access annotation holds for all classes in this package.]

[It is currently not standardized which result variables are accessible for plotting. It seems natural to not introduce new flags for this, but reuse the Access.XXX definition, e.g., for Access.icon only the variables can be stored in a result file that can also be inspected in the class, and for Access.nonPackageText all public and protected variables can be stored in a result file, because all variables can be inspected in the class.]

```

package CommercialFluid // Access icon , documentation , diagram
  package Examples // Access icon , documentation , diagram
    model PipeExample // Access everything , can be copied
    end PipeExample;

    package Circuits // Access icon , documentation , diagram
      model ClosedCircuit // Access everything , can be copied
      end ClosedCircuit;
    end Circuits;

    model SecretExample // No access
      annotation(Protection(access=Access.hide));
    end SecretExample;
    annotation(Protection(access=Access.nonPackageDuplicate));
  end Examples;

  package Pipe // Access icon
    model StraightPipe // Access icon
    end StraightPipe;
    annotation(Protection(access=Access.icon));
  end Pipe;

  package Vessels // Access icon , documentation , diagram
    model Tank // Access icon , documentation , diagram , text
    end Tank;
  end Vessels;
  annotation(Protection(access=Access.nonPackageText));
end CommercialFluid;
    
```

]

18.9.2 Licensing

In this section annotations within the Protection annotation are defined to restrict the usage of the encrypted package:

```

record Protection
  ...
  String features[:]=fill("", 0) "Required license features";
  record License
    String libraryKey;
    String licenseFile="" "Optional, default mapping if empty";
  end License;
end Protection;
    
```

The License annotation has only an effect on the top of an encrypted class and is then valid for the whole class hierarchy. (Usually the licensed class is a package.) The libraryKey is a secret string from the

library vendor and is the protection mechanism so that a user cannot generate his/her own authorization file since the `libraryKey` is unknown to him/her.

The `features` annotation defines the required license options. If the features vector has more than one element, then at least a license feature according to one of the elements must be present. As with the other annotations, the `features` annotation holds for the respective class and for all classes that are hierarchically on a lower level, unless further restricted by a corresponding annotation. If no license according to the `features` annotation is provided in the authorization file, the corresponding classes are not visible and cannot be used, not even internally in the package.

[Example:

```
// Requires license feature "LicenseOption"
annotation(Protection(features={"LicenseOption"}));

// Requires license features "LicenseOption1" or "LicenseOption2"
annotation(Protection(features={"LicenseOption1", "LicenseOption2"}));

// Requires license features ("LicenseOption1" and "LicenseOption2") or "
  LicenseOption3"
annotation(Protection(features={"LicenseOption1 LicenseOption2", "
  LicenseOption3"}));
```

]

In order that the protected class can be used either a tool specific license manager, or a license file (called `licenseFile`) must be present. The license file is standardized. It is a Modelica package without classes that has a `Protection` annotation of the following form which specifies a sequence of target records, which makes it natural to define start/end dates for different sets of targets individually:

```
record Authorization
  String licensor="" "Optional string to show information about the licensor";
  String libraryKey "Matching the key in the class. Must be encrypted and not
    visible";
  License license[:] "Definition of the license options and of the access
    rights";
end Authorization;

record License
  String licensee="" "Optional string to show information about the licensee";
  String id[:] "Unique machine identifications, e.g.\ MAC addresses";
  String features[:] =fill("", 0) "Activated library license features";
  String startDate="" "Optional start date in UTCformat YYYY-MM-DD";
  String expirationDate="" "Optional expiration date in UTCformat YYYY-MM-DD";
  String operations[:]=fill("",0) "Library usage conditions";
end License;
```

The format of the strings used for `libraryKey` and `id` are not specified (they are vendor specific). The `libraryKey` is a secret of the library developer. The `operations` define the usage conditions and the following are default names:

- "ExportBinary" Binary code generated from the Modelica code of the library can be included in binaries produced by a simulation tool.
- "ExportSource" Source code generated from the Modelica code of the library can be included in sources produced by a simulation tool.

Additional tool-specific names can also be used. To protect the `libraryKey` and the target definitions, the authorization file must be encrypted and must never show the `libraryKey`.

[All other information, especially `licensor` and `license` should be visible, in order that the user can get information about the license. It is useful to include the name of the tool in the authorization file name with which it was encrypted. Note, it is not useful to store this information in the annotation, because only the tool that encrypted the `Authorization` package can also decrypt it.]

[Example: (Before encryption:)]

```

// File MyLibrary\package.mo
package MyLibrary
  annotation(Protection(License(libraryKey="15783-A39323-498222-444ckk411",
    licenseFile="MyLibraryAuthorization_Tool.mo_lic), ...));
end MyLibrary;

// File MyLibrary\MyLibraryAuthorization_Tool.mo\
// (authorization file before encryption)
package MyLibraryAuthorization_Tool
  annotation(Authorization(
    libraryKey="15783-A39323-498222-444ckk411",
    licensor ="Organization A\nRoad, Country",
    license={
      License(licensee="Organization B, Mr. X",
        id={"lic:1269"}), // tool license number
      License(licensee="Organization C, Mr. Y",
        id={"lic:511"}, expirationDate="2010-06-30",
        operations={"ExportBinary"}),
      License(licensee="Organization D, Mr. Z",
        id={"mac:0019d2c9bfe7"}) // MAC address
    }));
end MyLibraryAuthorization_Tool;

```

]

18.10 Annotations for Functions

18.10.1 Function Derivative Annotations

See section 12.7.1

18.10.2 Inverse Function Annotation

See section 12.8.

18.10.3 External Function Annotations

See section 12.9.4.

18.11 Annotation Choices for Modifications and Redeclarations

See section 7.3.4.

18.12 Annotation for External Libraries and Include Files

See section 12.9.4.

Chapter 19

Unit Expressions

Unless otherwise stated, the syntax and semantics of unit expressions in Modelica conform with the international standards *International System of Units (SI)* by BIPM superseding parts of ISO 31/0-1992 *General principles concerning quantities, units and symbols* and ISO 1000-1992 *SI units and recommendations for the use of their multiples and of certain other units*. Unfortunately, these standards do not define a formal syntax for unit expressions. There are recommendations and Modelica exploits them.

Note that this document uses the American spelling meter, whereas the SI specification from BIPM uses the British spelling metre.

Examples for the syntax of unit expressions used in Modelica: "N.m", "kg.m/s²", "kg.m.s⁻²" "1/rad", "mm/s".

19.1 The Syntax of Unit Expressions

```
unit_expression:
    unit_numerator [ "/" unit_denominator ]

unit_numerator:
    "1" | unit_factors | "(" unit_expression ")"

unit_denominator:
    unit_factor | "(" unit_expression ")"
```

The unit of measure of a dimension free quantity is denoted by "1". The SI standard does not define any precedence between multiplications and divisions. The SI standard does not allow multiple units to the right of the division-symbol (/) since the result is ambiguous; either the divisor shall be enclosed in parentheses, or negative exponents used instead of division, for example, "J/(kg.K)" may be written as "J.kg⁻¹.K⁻¹".

```
unit_factors:
    unit_factor [ unit_mulop unit_factors ]

unit_mulop:
    "."
```

The SI standard specifies that a multiplication operator symbol is written as space or as a dot. The SI standard requires that this *dot* is a bit above the base line: ‘·’, which is not part of ASCII. The ISO standard also prefers ‘·’, but Modelica supports the ISO alternative ‘.’, which is an ordinary *dot* on the base line.

For example, Modelica does not support "Nm" for newton-meter, but requires it to be written as "N.m".

```
unit_factor:
    unit_operand [ unit_exponent ]

unit_exponent:
```

```
[ "+" | "-" ] integer
```

The SI standard uses super-script for the exponentiation, and does thus not define any operator symbol for exponentiation. A `unit_factor` consists of a `unit_operand` possibly suffixed by a possibly signed integer number, which is interpreted as an exponent. There must be no spacing between the `unit_operand` and a possible `unit_exponent`.

```
unit_operand:
  unit_symbol | unit_prefix unit_symbol

unit_prefix:
  Y | Z | E | P | T | G | M | k | h | da | d | c | m | u | n | p | f | a | z
  | y
```

A `unit_symbol` is a string of letters. A basic support of units in Modelica should know the basic and derived units of the SI system. It is possible to support user defined unit symbols. In the base version Greek letters is not supported, but full names must then be written, for example "Ohm".

A `unit_operand` should first be interpreted as a `unit_symbol` and only if not successful the second alternative assuming a prefixed operand should be exploited. There must be no spacing between the `unit_symbol` and a possible `unit_prefix`. The values of the prefixes are according to the ISO standard. The letter `u` is used as a symbol for the prefix micro.

[A tool may present Ohm as Ω and the prefix `u` as μ , and similarly `m2` as m^2 .]

19.2 Examples

The unit expression "m" means meter and not milli (10^{-3}), since prefixes cannot be used in isolation. For millimeter use "mm" and for square meter, m^2 , write "m2".

The expression "mm2" means $(10^{-3}m)^2 = 10^{-6}m^2$. Note that exponentiation includes the prefix.

The unit expression "T" means tesla, but note that the letter T is also the symbol for the prefix tera which has a multiplier value of 10^{12} .

Chapter 20

The Modelica Standard Library

In order that a modeler can quickly build up system models, it is important that libraries of the most commonly used components are available, ready to use, and sharable between applications. For this reason, the Modelica Association develops and maintains a growing *Modelica Standard Library* called `package Modelica`. For an overview of the current version see <https://github.com/modelica/ModelicaStandardLibrary>. This is a free library that can be used without essential restrictions, e.g., in commercial Modelica simulation environments. The Modelica Standard Library is tool-neutral, and relies on a small library, `ModelicaServices`, that each conformant tool must implement to handle tool-specific couplings e.g. for animation. Furthermore, other people and organizations are developing free and commercial Modelica libraries. For information about these libraries and for downloading the free libraries see <https://www.modelica.org/libraries/>.

Appendix A

Glossary

algorithm section: part of a class definition consisting of the keyword `algorithm` followed by a sequence of statements. Like an equation, an algorithm section relates variables, i.e. constrains the values that these variables can take simultaneously. In contrast to an equation section, an algorithm section distinguishes inputs from outputs: An algorithm section specifies how to compute output variables as a function of given input variables. A Modelica processor may actually invert an algorithm section, i.e. compute inputs from given outputs, e.g. by search (generate and test), or by deriving an inverse algorithm symbolically. (See chapter 11.)

array or array variable: a component whose components are array elements. For an array, the ordering of its components matters: The k th element in the sequence of components of an array x is the array element with index k , denoted $x[k]$. All elements of an array have the same type. An array element may again be an array, i.e. arrays can be nested. An array element is hence referenced using n indices in general, where n is the number of dimensions of the array. Special cases are matrix ($n = 2$) and vector ($n = 1$). `Integer` indices start with 1, not zero. (See chapter 10.)

array constructor: an array can be built using the `array` function – with the shorthand `{a, b, ...}` – and can also include an iterator to build an array of expressions. (See section 10.4.)

array element: a component contained in an array. An array element has no identifier. Instead they are referenced by array access expressions called indices that use enumeration values or positive integer index values. (See chapter 10.)

assignment: a statement of the form $x := \text{expr}$. The expression `expr` must not have higher variability than x . (See section 11.2.1.)

attribute: a component contained in a scalar component, such as `min`, `max`, and `unit`. All attributes are predefined and attribute values can only be defined using a modification, such as in `Real x(unit="kg")`. Attributes cannot be accessed using dot notation, and are not constrained by equations and algorithm sections. E.g. in `Real x(unit="kg") = y`; only the values of x and y are declared to be equal, but not their unit attributes, nor any other attribute of x and y . (See section 4.8.)

base class: class A is called a base class of B, if class B extends class A. This relation is specified by an extends clause in B or in one of B's base classes. A class inherits all elements from its base classes, and may modify all non-final elements inherited from base classes. (See section 7.1.)

binding equation: Either a declaration equation or an element modification for the value of the variable. In a non-function a component with a binding equation has its value bound to some expression. (See section 8.1.) In a function an input component uses the value of the binding equation as default value, a non-input component starts with that value. (See section 12.4.4.)

class: a description that generates an object called instance. The description consists of a class definition, a modification environment that modifies the class definition, an optional list of dimension expressions if the class is an array class, and a lexically enclosing class for all classes. (See section 4.5.)

class type or **inheritance interface:** property of a class, consisting of a number of attributes and a set of public or protected elements consisting of element name, element type, and element attributes.

(See section 6.3.2.)

component or **variable**: an instance (object) generated by a component declaration. Special kinds of components are scalar, array, and attribute. (See section 4.4.)

component declaration: an element of a class definition that generates a component. A component declaration specifies (1) a component name, i.e., an identifier, (2) the class to be flattened in order to generate the component, and (3) an optional `Boolean` parameter expression. Generation of the component is suppressed if this parameter expression evaluates to false. A component declaration may be overridden by an element-redeclaration. (See section 4.4.)

component reference: An expression containing a sequence of identifiers and indices. A component reference is equivalent to the referenced object, which must be a component. A component reference is resolved (evaluated) in the scope of a class (or expression for the case of a local iterator variable). A scope defines a set of visible components and classes. Example reference: `Ele.Resistor.u[21].r` (See section 4.4 and section 10.6.9.)

declaration assignment: Deprecated name for binding equation in function.

declaration equation: Equation of the form `x = expression` defined by a component declaration. The expression must not have higher variability than the declared component `x`. Unlike other equations, a declaration equation can be overridden (replaced or removed) by an element modification. (See section 4.4.2.1.)

derived class or **subclass** or **extended class**: class B is called derived from A, if B extends A. (See chapter 7.)

element: part of a class definition, one of: class definition, component declaration, or extends clause. Component declarations and class definitions are called named elements. An element is either inherited from a base class or local.

element modification: part of a modification, overrides the declaration equation in the class used by the instance generated by the modified element. Example: `vcc(unit="V")=1000`. (See section 7.2.)

element-redeclaration: part of a modification, replaces one of the named elements possibly used to build the instance generated by the element that contains the redeclaration. Example: `redeclare type Voltage = Real(unit="V")` replaces `type Voltage`. (See section 7.3.)

encapsulated: a class that does not depend on where it is placed in the package-hierarchy, since its lookup is stopped at the encapsulated boundary. (See section 5.3.1.)

equation: part of a class definition. A scalar equation relates scalar variables, i.e. constrains the values that these variables can take simultaneously. When $n-1$ variables of an equation containing n variables are known, the value of the n th variable can be inferred (solved for). In contrast to a statement in an algorithm section, an equation does not define for which of its variable it is to be solved. Special cases are: initial equations, instantaneous equations, declaration equations. (See chapter 8.)

event: something that occurs instantaneously at a specific time or when a specific condition occurs. Events are for example defined by the condition occurring in a when clause, if clause, or if expression. (See section 8.5.)

expression: a term built from operators, function references, components, or component references (referring to components) and literals. Each expression has a type and a variability. (See chapter 3.)

extends clause: an unnamed element of a class definition that uses a name and an optional modification to specify a base class of the class defined using the class definition. (See chapter 7.)

flattening: the computation that creates a flattened class of a given class, where all inheritance, modification, etc. has been performed and all names resolved, consisting of a flat set of equations, algorithm sections, component declarations, and functions. (See section 5.6.)

function: a class of the specialized class function. (See chapter 12.)

function subtype or **function compatible interface**: A is a function subtype of B iff A is a subtype of B and the additional arguments of function A that are not in function B are defined in such a way (e.g. additional arguments need to have default values), that A can be called at places where B is called. (See section 6.6.)

- identifier** or **ident**: an atomic (not composed) name. Example: `Resistor` (See section 2.3.)
- index** or **subscript**: An expression, typically of `Integer` type or the colon symbol (`:`), used to reference a component (or a range of components) of an array. (See section 10.5.)
- inheritance interface** or **class type**: property of a class, consisting of a number of attributes and a set of public or protected elements consisting of element name, element type, and element attributes. (See section 6.3.2.)
- instance**: the object generated by a class. An instance contains zero or more components (i.e. instances), equations, algorithms, and local classes. An instance has a type. Basically, two instances have same type, if their important attributes are the same and their public components and classes have pair wise equal identifiers and types. More specific type equivalence definitions are given e.g. for functions.
- instantaneous**: An equation or statement is instantaneous if it holds only at events, i.e., at single points in time. The equations and statements of a when-clause are instantaneous. (See section 8.3.5 and section 11.2.7.)
- interface**: see type. (See section 6.3.)
- literal**: a real, integer, boolean, enumeration, or string literal. Used to build expressions. (See section 2.4.)
- matrix**: an array where the number of dimensions is 2. (See chapter 10.)
- modification**: part of an element. Modifies the instance generated by that element. A modification contains element modifications and element redeclarations. (See section 7.2.)
- modification environment**: the modification environment of a class defines how to modify the corresponding class definition when flattening the class. (See section 7.2.2.)
- name**: Sequence of one or more identifiers. Used to reference a class. A class name is resolved in the scope of a class, which defines a set of visible classes. Example name: `Ele.Resistor`. (See section 2.3.2.)
- operator record**: A record with user-defined operations; defining e.g. multiplication and addition see chapter 14.
- partial**: a class that is incomplete and cannot be instantiated in a simulation model; useful e.g. as a base-class. (See section 4.4.2.)
- partial flattening**: first find the names of declared local classes and components. Modifiers, if present, are merged to the local elements and redeclarations are performed. Then base-classes are looked up, flattened and inserted into the class. See also flattening, which additionally flattens local elements and performs modifications. (See section 5.6.)
- plug-compatibility**: see restricted subtyping and section 6.5.
- predefined type**: one of the types `Real`, `Boolean`, `Integer`, `String` and types defined as `enumeration` types. The component declarations of the predefined types define attributes such as `min`, `max`, and `unit`. (See section 4.8.)
- prefix**: property of an element of a class definition which can be present or not be present, e.g. `final`, `public`, `flow`. (See section 4.4.2.2.)
- primitive type**: one of the built-in types `RealType`, `BooleanType`, `IntegerType`, `StringType`, `EnumType`. The primitive types are used to define attributes and value of predefined types and enumeration types. (See section 4.8.)
- redeclare**: the modifier that changes a replaceable element. (See section 7.3)
- replaceable**: an element that can be replaced by a different element having a compatible interface. (See section 7.3)
- restricted subtyping** or **plug-compatibility**: a type A is a restricted subtype of type B iff A is a subtype of B, and all public components present in A but not in B must be default-connectable. This is used to avoid introducing, via a redeclaration, an un-connected connector in the object/class of type A at a level where a connection is not possible. (See section 6.5.)
- scalar** or scalar variable: a variable that is not an array.

simple type: `Real`, `Boolean`, `Integer`, `String` and enumeration types

specialized class: one of: `model`, `connector`, `package`, `record`, `block`, `function`, `type`. The class restriction of a class represents an assertion regarding the content of the class and restricts its use in other classes. For example, a class having the `package` class restriction must only contain classes and constants. (See section 4.6.)

subtype or **interface compatible:** relation between types. `A` is a subtype of (interface compatible with) `B` iff a number of properties of `A` and `B` are the same and all important elements of `B` have corresponding elements in `A` with the same names and their types being subtypes of the corresponding element types in `B`. See also restricted subtyping and function restricted subtyping. (See section 6.4.)

supertype: relation between types. The inverse of subtype. `A` is a subtype of `B` means that `B` is a supertype of `A`. (See section 6.4.)

transitively nonreplaceable: a class reference is considered transitively non-replaceable if there are no replaceable elements in the referenced class, or any of its base classes or constraining types transitively at any level. (See section 6.3.1.)

type or **interface:** property of an instance, expression, consisting of a number of attributes and a set of public elements consisting of element name, element type, and element attributes. Note: The concept of class type is a property of a class definition. (See section 6.3.)

variability: property of an expression: one of

- **continuous:** an expression that may change its value at any point in time.
- **discrete:** may change its value only at events during simulation.
- **parameter:** constant during the entire simulation, recommended to change for a component.
- **constant:** constant during the entire simulation (can be used in a package) .

Assignments `x := expr` and binding equations `x = expr` must satisfy a variability constraint: The expression must not have a higher variability than component `x`. (See section 3.8.)

variable: synonym for component. (See section 4.4.)

vector: an array where the number of dimensions is 1. (See chapter 10.)

Appendix B

Modelica Concrete Syntax

B.1 Lexical conventions

The following syntactic meta symbols are used (extended BNF):

```
[ ] optional
{ } repeat zero or more times
| or
"text" The text is treated as a single token (no white-space between any
      characters)
```

The following lexical units are defined (the ones in boldface are the ones used in the grammar, the rest are just internal to the definition of other lexical units):

```
IDENT = NONDIGIT { DIGIT | NONDIGIT } | Q-IDENT
Q-IDENT = "'" { Q-CHAR | S-ESCAPE } "'"
NONDIGIT = "_" | letters "a" to "z" | letters "A" to "Z"
STRING = "" { S-CHAR | S-ESCAPE } ""
S-CHAR = see below
Q-CHAR = NONDIGIT | DIGIT | "!" | "#" | "$" | "%" | "&" | "(" | ")"
        | "*" | "+" | "," | "-" | "." | "/" | ":" | ";" | "<" | ">" | "="
        | "?" | "@" | "[" | "]" | "^" | "{" | "}" | "|" | "~" | " " | ""
S-ESCAPE = "\'" | "\" | "\?" | "\\\"
        | "\a" | "\b" | "\f" | "\n" | "\r" | "\t" | "\v"
DIGIT = "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9"
UNSIGNED-INTEGERS = DIGIT { DIGIT }
UNSIGNED-REAL =
  UNSIGNED-INTEGERS "." [ UNSIGNED-INTEGERS ]
  | UNSIGNED-INTEGERS [ "." [ UNSIGNED-INTEGERS ] ]
  ( "e" | "E" ) [ "+" | "-" ] UNSIGNED-INTEGERS
  | "." UNSIGNED-INTEGERS [ ( "e" | "E" ) [ "+" | "-" ] UNSIGNED-INTEGERS ]
```

S-CHAR is any member of the Unicode character set (<http://www.unicode.org>; see section 13.4 for storing as UTF-8 on files) except double-quote '"', and backslash '\'.

For identifiers the redundant escapes ('\?' and '\') are the same as the corresponding non-escaped variants ('?' and '). The single quotes are part of an identifier. E.g. 'x' and x are different IDENTs.

Note:

- White-space and comments can be used between separate lexical units and/or symbols, and also separates them. Each lexical unit will consume the maximum number of characters from the input stream. White-space and comments cannot be used inside other lexical units, except for STRING and Q-IDENT where they are treated as part of the STRING or Q-IDENT lexical unit.
- String constant concatenation "a" "b" becoming "ab" (as in C) is replaced by the "+" operator in Modelica.

- Modelica uses the same comment syntax as C++ and Java (i.e., `//` signals the start of a line comment and `/* ... */` is a multi-line comment); comments may contain any Unicode character. Modelica also has structured comments in the form of annotations and string comments.
- Each description-string or string in annotations (= STRING with production annotation-clause in the grammar) may contain any member of the Unicode character set. All other strings have to contain only the sub-set of Unicode characters identical with the 7-bit US-ASCII character set.

[As a consequence, operators like ‘>’ or ‘<’, and external functions only operate on ASCII strings and not on Unicode-strings. Within a description-string the tags <HTML> and </HTML> or <html> and </html> define optionally begin and end of content that is HTML encoded.]

- Boldface denotes keywords of the Modelica language. Keywords are reserved words and shall not be used as identifiers.
- Productions use hyphen as separator both in the grammar and in the text. Previously the grammar used underscore.

B.2 Grammar

B.2.1 Stored Definition – Within

```
stored-definition :
  [ within [ name ] ";" ]
  { [ final ] class-definition ";" }
```

B.2.2 Class Definition

```
class-definition :
  [ encapsulated ] class-prefixes class-specifier

class-prefixes :
  [ partial ]
  ( class
    | model
    | [ operator ] record
    | block
    | [ expandable ] connector
    | type
    | package
    | [ pure | impure ] [ operator ] function
    | operator
  )

class-specifier :
  long-class-specifier | short-class-specifier | der-class-specifier

long-class-specifier :
  IDENT description-string composition end IDENT
  | extends IDENT [ class-modification ] description-string composition
  end IDENT

short-class-specifier :
  IDENT "=" base-prefix type-specifier [ array-subscripts ]
  [ class-modification ] description
  | IDENT "=" enumeration "(" ( [ enum-list ] | ":" ) ")" description

der-class-specifier :
  IDENT "=" der "(" type-specifier "," IDENT { "," IDENT } ")" description

base-prefix :
  [ input | output ]
```

```

enum-list :
  enumeration-literal { "," enumeration-literal }

enumeration-literal :
  IDENT description

composition :
  element-list
  { public element-list
    | protected element-list
    | equation-section
    | algorithm-section
  }
  [ external [ language-specification ]
  [ external-function-call ] [ annotation-clause ] ";" ]
  [ annotation-clause ";" ]

language-specification :
  STRING

external-function-call :
  [ component-reference "=" ]
  IDENT "(" [ expression-list ] ")"

element-list :
  { element ";" }

element :
  import-clause
  | extends-clause
  | [ redeclare ]
  [ final ]
  [ inner ] [ outer ]
  ( class-definition
  | component-clause
  | replaceable ( class-definition | component-clause )
  [ constraining-clause description ]
  )

import-clause :
  import
  ( IDENT "=" name
  | name [ "." | "*" | "{" import-list "}" ] )
  )
  description

import-list :
  IDENT { "," IDENT }
  
```

B.2.3 Extends

```

extends-clause :
  extends type-specifier [ class-modification ] [ annotation-clause ]

constraining-clause :
  constrainedby type-specifier [ class-modification ]
  
```

B.2.4 Component Clause

```

component-clause :
  type-prefix type-specifier [ array-subscripts ] component-list

type-prefix :
  
```

```

[ flow | stream ]
[ discrete | parameter | constant ]
[ input | output ]

component-list :
  component-declaration { "," component-declaration }

component-declaration :
  declaration [ condition-attribute ] description

condition-attribute :
  if expression

declaration :
  IDENT [ array-subscripts ] [ modification ]

```

B.2.5 Modification

```

modification :
  class-modification [ "=" expression ]
  | "=" expression
  | ":=" expression

class-modification :
  "(" [ argument-list ] ")"

argument-list :
  argument { "," argument }

argument :
  element-modification-or-replaceable
  | element-redeclaration

element-modification-or-replaceable :
  [ each ] [ final ] ( element-modification | element-replaceable )

element-modification :
  name [ modification ] description-string

element-redeclaration :
  redeclare [ each ] [ final ]
  ( short-class-definition | component-clause1 | element-replaceable )

element-replaceable :
  replaceable ( short-class-definition | component-clause1 )
  [ constraining-clause ]

component-clause1 :
  type-prefix type-specifier component-declaration1

component-declaration1 :
  declaration description

short-class-definition :
  class-prefixes short-class-specifier

```

B.2.6 Equations

```

equation-section :
  [ initial ] equation { equation ";" }

algorithm-section :
  [ initial ] algorithm { statement ";" }

```

```

equation :
  ( simple-expression "=" expression
    | if-equation
    | for-equation
    | connect-clause
    | when-equation
    | component-reference function-call-args
  )
  description

statement :
  ( component-reference ( ":" expression | function-call-args )
    | "(" output-expression-list ")" ":"
      component-reference function-call-args
    | break
    | return
    | if-statement
    | for-statement
    | while-statement
    | when-statement
  )
  description

if-equation :
  if expression then
    { equation ";" }
  { elseif expression then
    { equation ";" }
  }
  [ else
    { equation ";" }
  ]
  end if

if-statement :
  if expression then
    { statement ";" }
  { elseif expression then
    { statement ";" }
  }
  [ else
    { statement ";" }
  ]
  end if

for-equation :
  for for-indices loop
    { equation ";" }
  end for

for-statement :
  for for-indices loop
    { statement ";" }
  end for

for-indices :
  for-index { "," for-index }

for-index :
  IDENT [ in expression ]

while-statement :

```

```

while expression loop
{ statement ";" }
end while

when-equation :
  when expression then
    { equation ";" }
  { elseif expression then
    { equation ";" }
  }
end when

when-statement :
  when expression then
    { statement ";" }
  { elseif expression then
    { statement ";" }
  }
end when

connect-clause :
  connect "(" component-reference "," component-reference ")"

```

B.2.7 Expressions

```

expression :
  simple-expression
  | if expression then expression
  { elseif expression then expression }
  else expression

simple-expression :
  logical-expression [ ":" logical-expression [ ":" logical-expression ] ]

logical-expression :
  logical-term { or logical-term }

logical-term :
  logical-factor { and logical-factor }

logical-factor :
  [ not ] relation

relation :
  arithmetic-expression [ relational-operator arithmetic-expression ]

relational-operator :
  "<" | "<=" | ">" | ">=" | "==" | "<>"

arithmetic-expression :
  [ add-operator ] term { add-operator term }

add-operator :
  "+" | "-" | ".+" | "-."

term :
  factor { mul-operator factor }

mul-operator :
  "*" | "/" | ".*" | "./"

factor :
  primary [ ("^" | ".^") primary ]

```

```

primary :
  UNSIGNED-NUMBER
  | STRING
  | false
  | true
  | ( component-reference | der | initial | pure ) function-call-args
  | component-reference
  | "(" output-expression-list ")"
  | "[" expression-list { ";" expression-list } "]"
  | "{" array-arguments }"
  | end

UNSIGNED-NUMBER :
  UNSIGNED-INTEGGER | UNSIGNED-REAL

type-specifier :
  ["."] name

name :
  IDENT { "." IDENT }

component-reference :
  ["."] IDENT [ array-subscripts ] { "." IDENT [ array-subscripts ] }

function-call-args :
  "(" [ function-arguments ] ")"

function-arguments :
  expression [ "," function-arguments-non-first | for for-indices ]
  | function-partial-application [ "," function-arguments-non-first ]
  | named-arguments

function-arguments-non-first :
  function-argument [ "," function-arguments-non-first ]
  | named-arguments

array-arguments :
  expression [ "," array-arguments-non-first | for for-indices ]

array-arguments-non-first :
  expression [ "," array-arguments-non-first ]

named-arguments: named-argument [ "," named-arguments ]

named-argument: IDENT "=" function-argument

function-argument :
  function-partial-application | expression

function-partial-application :
  function type-specifier "(" [ named-arguments ] ")"

output-expression-list :
  [ expression ] { "," [ expression ] }

expression-list :
  expression { "," expression }

array-subscripts :
  "[" subscript { "," subscript } "]"

subscript :

```



```
"|" | expression  
  
description :  
  description-string [ annotation-clause ]  
  
description-string :  
  [ STRING { "+" STRING } ]  
  
annotation-clause :  
  annotation class-modification
```

Appendix C

Modelica DAE Representation

In this appendix, the mapping of a Modelica model into an appropriate mathematical description form is discussed.

In a first step, a Modelica translator transforms a hierarchical Modelica simulation model into a “flat” set of Modelica “statements”, consisting of the equation and algorithm sections of all used components by:

- Expanding all class definitions (flattening the inheritance tree) and adding the equations and assignment statements of the expanded classes for every instance of the model.
- Replacing all connect-equations by the corresponding equations of the connection set (see section 9.2).
- Mapping all algorithm sections to equation sets.
- Mapping all when-clauses to equation sets (see section 8.3.5).

As a result of this transformation process, a set of equations is obtained consisting of differential, algebraic and discrete equations of the following form where ($v := [p; t; \dot{x}; x; y; z; m; \text{pre}(z); \text{pre}(m)]$):

$$0 = f_x(v, c) \tag{C.1a}$$

$$0 = \begin{cases} f_z(v, c) & \text{at events} \\ z - \text{pre}(z) & \text{otherwise} \end{cases} \tag{C.1b}$$

$$m := f_m(v, c) \tag{C.1c}$$

$$c := f_c(\text{relation}(v)) \tag{C.1d}$$

and where

- p : Modelica variables declared as `parameter` or `constant`, i.e., variables without any time-dependency.
- t : Modelica variable `time`, the independent (real) variable.
- $x(t)$: Modelica variables of type `Real`, appearing differentiated.
- $y(t)$: Continuous-time modelica variables of type `Real` that do not appear differentiated (= algebraic variables).
- $z(t_e)$: Discrete-time modelica variables of type `Real`. These variables change their value only at event instants t_e . $\text{pre}(z)$ are the values of z immediately before the current event occurred.
- $m(t_e)$: Modelica variables of discrete-valued types (`Boolean`, `Integer`, etc) which are unknown. These variables change their value only at event instants t_e . $\text{pre}(m)$ are the values of m immediately before the current event occurred.
- $c(t_e)$: The conditions of all if-expressions generated including when-clauses after conversion, see section 8.3.5).
- $\text{relation}(v)$: A relation containing variables v_i , e.g. $v_1 > v_2, v_3 \geq 0$.

For simplicity, the special cases of `noEvent` and `reinit` are not contained in the equations above and are not discussed below.

The key difference between the two groups of discrete-time variables z and m here is how they are determined. The interpretation of the solved form of (C.1c) is that given values for everything else, there is a closed-form solution for m in the form of a sequence of assignments to each of the variables of m in turn – there must be no cyclic dependencies between the equations used to solve for m . Further, each of the original model equations behind (C.1c) must be given in solved form, at most requiring flipping sides of the equation to obtain the used assignment form. The interpretation of the non-solved form of (C.1b) at events, on the other hand, is that at events, the discrete-time `Real` variables z are solved together with the continuous-time variables using (C.1a) and (C.1b).

[Example: The following model is illegal since there is no equation in solved form that can be used in (C.1c) to solve for the discrete-valued variable y :

```

model M
  Boolean x;
  Boolean y;
equation
  x = time >= 1.0;
  not y = x; /* Equation in solved form, but not with respect to y. */
end M;

```

]

The generated set of equations is used for simulation and other analysis activities. Simulation means that an initial value problem is solved, i.e., initial values have to be provided for the states x , section 8.6. The equations define a DAE (Differential Algebraic Equations) which may have discontinuities, a variable structure and/or which are controlled by a discrete-event system. Such types of systems are called *hybrid DAEs*. Simulation is performed in the following way:

1. The DAE (C.1a) is solved by a numerical integration method. In this phase the conditions c of the if- and when-clauses, as well as the discrete-time variables z and m are kept constant. Therefore, (C.1a) is a continuous function of continuous variables and the most basic requirement of numerical integrators is fulfilled.
2. During integration, all relations from (C.1d) are monitored. If one of the relations changes its value an event is triggered, i.e., the exact time instant of the change is determined and the integration is halted. As discussed in section 8.5, relations which depend only on time are usually treated in a special way, because this allows determining the time instant of the next event in advance.
3. At an event instant, (C.1) is a mixed set of algebraic equations which is solved for the `Real`, `Boolean` and `Integer` unknowns.
4. After an event is processed, the integration is restarted with 1.

Note, that both the values of the conditions c as well as the values of z and m (all discrete-time `Real`, `Boolean` and `Integer` variables) are only changed at an event instant and that these variables remain constant during continuous integration. At every event instant, new values of the discrete-time variables z and m , as well as of new initial values for the states x , are determined. The change of discrete-time variables may characterize a new structure of a DAE where elements of the state vector x are *disabled*. In other words, the number of state variables, algebraic variables and residue equations of a DAE may change at event instants by disabling the appropriate part of the DAE. For clarity of the equations, this is not explicitly shown by an additional index in (C.1).

At an event instant, including the initial event, the model equations are reinitialized according to the following iteration procedure:

```

known variables: x, t, p
unknown variables: dx/dt, y, z, m, pre(z), pre(m), c

// pre(z) = value of z before event occurred
// pre(m) = value of m before event occurred
loop
  solve (1) for the unknowns, with pre(z) and pre(m) fixed

```

```

if z == pre(z) and m == pre(m) then break
pre(z) := z
pre(m) := m
end loop

```

Solving (C.1) for the unknowns is non-trivial, because this set of equations contains not only `Real`, but also discrete-valued unknowns. Usually, in a first step these equations are sorted and in many cases the discrete-valued unknowns m can be just computed by a forward evaluation sequence. In some cases, there remain systems of equations involving m due to cyclic dependencies with y and z (e.g. for ideal diodes, Coulomb friction elements), and specialized algorithms have to be used to solve them.

Due to the construction of the equations by *flattening* a Modelica model, the hybrid DAE (C.1) contains a huge number of sparse equations. Therefore, direct simulation of (C.1) requires sparse matrix methods. However, solving this initial set of equations directly with a numerical method is both unreliable and inefficient. One reason is that many Modelica models, like the mechanical ones, have a DAE index of 2 or 3, i.e., the overall number of states of the model is less than the sum of the states of the sub-components. In such a case, every direct numerical method has the difficulty that the numerical condition becomes worse, if the integrator step size is reduced and that a step size of zero leads to a singularity. Another problem is the handling of idealized elements, such as ideal diodes or Coulomb friction. These elements lead to mixed systems of equations having both `Real` and `Boolean` unknowns. Specialized algorithms are needed to solve such systems.

To summarize, symbolic transformation techniques are needed to transform (C.1) into a set of equations which can be numerically solved reliably. Most important, the algorithm of Pantelides should to be applied to differentiate certain parts of the equations in order to reduce the index. Note, that also explicit integration methods, such as Runge-Kutta algorithms, can be used to solve (C.1a), after the index of (C.1a) has been reduced by the Pantelides algorithm: During continuous integration, the integrator provides x and t . Then, (C.1a) is a linear or nonlinear system of equations to compute the algebraic variables y and the state derivatives dx/dt and the model returns dx/dt to the integrator by solving these systems of equations. Often, (C.1a) is just a linear system of equations in these unknowns, so that the solution is straightforward. This procedure is especially useful for real-time simulation where usually explicit one-step methods are used.

Appendix D

Derivation of Stream Equations

This appendix contains a derivation of the equation for stream connectors from chapter 15.

D.1 Reasons for avoiding the actual mixing enthalpy in connector definitions

Consider a connection set with n connectors, and denote the mass flow rates $\mathbf{m_flow}$ by \tilde{m} . The mixing enthalpy is defined by the mass balance (the general mass-balance for a component has $\dot{m} = \sum \tilde{m}$ which simplifies for the mixing enthalpy where $m = 0$ and thus $\dot{m} = 0$)

$$0 = \sum_{j=1}^n \tilde{m}_j$$

and similarly the energy balance

$$0 = \sum_{j=1}^n \tilde{H}_j$$

with

$$\tilde{H}_j = \tilde{m}_j \begin{cases} h_{\text{mix}} & \text{if } \tilde{m}_j > 0 \\ h_{\text{outflow},j} & \text{if } \tilde{m}_j \leq 0 \end{cases}$$

Herein, mass flow rates are positive when entering models (exiting the connection set). The specific enthalpy represents the specific enthalpy inside the component, close to the connector, for the case of outflow. Expressed with variables used in the balance equations we arrive at:

$$h_{\text{outflow},j} = \begin{cases} \frac{\tilde{H}_j}{\tilde{m}_j} & \text{if } \tilde{m}_j < 0 \\ \text{arbitrary} & \text{if } \tilde{m}_j \geq 0 \end{cases}$$

While these equations are suitable for device-oriented modeling, the straightforward usage of this definition leads to models with discontinuous residual equations, which violates the prerequisites of several solvers for nonlinear equation systems. This is the reason why the actual mixing enthalpy is not modelled directly in the model equations. The stream connectors provide a suitable alternative.




Figure D.1: Exemplary connection set with three connected components and a common mixing enthalpy.

D.2 Rationale for the formulation of inStream

For simplicity, the derivation of `inStream` is shown at hand of 3 model components that are connected together. The case for N connections follows correspondingly.

The energy and mass balance equations for the connection set for 3 components are (see above):

$$\begin{aligned}
 0 = \tilde{m}_1 \cdot \begin{cases} h_{\text{mix}} & \text{if } \tilde{m}_1 > 0 \\ h_{\text{outflow},1} & \text{if } \tilde{m}_1 \leq 0 \end{cases} \\
 + \tilde{m}_2 \cdot \begin{cases} h_{\text{mix}} & \text{if } \tilde{m}_2 > 0 \\ h_{\text{outflow},2} & \text{if } \tilde{m}_2 \leq 0 \end{cases} \quad (\text{D.1a})
 \end{aligned}$$

$$\begin{aligned}
 + \tilde{m}_3 \cdot \begin{cases} h_{\text{mix}} & \text{if } \tilde{m}_3 > 0 \\ h_{\text{outflow},3} & \text{if } \tilde{m}_3 \leq 0 \end{cases} \\
 0 = \tilde{m}_1 + \tilde{m}_2 + \tilde{m}_3 \quad (\text{D.1b})
 \end{aligned}$$

The balance equations are implemented using a max operator in place of the piecewise expressions, taking care of the different flow directions:

$$\begin{aligned}
 0 = \max(\tilde{m}_1, 0)h_{\text{mix}} - \max(-\tilde{m}_1, 0)h_{\text{outflow},1} \\
 + \max(\tilde{m}_2, 0)h_{\text{mix}} - \max(-\tilde{m}_2, 0)h_{\text{outflow},2} \\
 + \max(\tilde{m}_3, 0)h_{\text{mix}} - \max(-\tilde{m}_3, 0)h_{\text{outflow},3} \quad (\text{D.2a})
 \end{aligned}$$

$$\begin{aligned}
 0 = \max(\tilde{m}_1, 0) - \max(-\tilde{m}_1, 0) \\
 + \max(\tilde{m}_2, 0) - \max(-\tilde{m}_2, 0) \\
 + \max(\tilde{m}_3, 0) - \max(-\tilde{m}_3, 0) \quad (\text{D.2b})
 \end{aligned}$$

Equation (D.2a) is solved for h_{mix}

$$h_{\text{mix}} = \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1} + \max(-\tilde{m}_2, 0)h_{\text{outflow},2} + \max(-\tilde{m}_3, 0)h_{\text{outflow},3}}{\max(\tilde{m}_1, 0) + \max(\tilde{m}_2, 0) + \max(\tilde{m}_3, 0)}$$

Using (D.2b), the denominator can be changed to:

$$h_{\text{mix}} = \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1} + \max(-\tilde{m}_2, 0)h_{\text{outflow},2} + \max(-\tilde{m}_3, 0)h_{\text{outflow},3}}{\max(-\tilde{m}_1, 0) + \max(-\tilde{m}_2, 0) + \max(-\tilde{m}_3, 0)}$$

Above it was shown that an equation of this type does not yield properly formulated model equations. In the streams concept we therefore decide to split the energy balance, which consists of different branches depending on the mass flow direction. Consequently, separate energy balances are the result; each valid for specific flow directions.

In a model governing equations have to establish the specific enthalpy of fluid leaving the model based on the specific enthalpy of fluid flowing into it. Whenever the mixing enthalpy is *used* in a model it is therefore the mixing enthalpy under the assumption of fluid flowing into said model.

We establish this quantity using a dedicated operator $\text{inStream}(h_{\text{outflow},i}) = h_{\text{mix}}$ assuming that $\tilde{m}_i \geq 0$. This leads to three different incarnations of (n in the general case). This is illustrated in the figure below. For the present example of three components in a connection set, this means the following.

$$\begin{aligned} \text{inStream}(h_{\text{outflow},1}) &= \frac{\max(-\tilde{m}_2, 0)h_{\text{outflow},2} + \max(-\tilde{m}_3, 0)h_{\text{outflow},3}}{\max(-\tilde{m}_2, 0) + \max(-\tilde{m}_3, 0)} \\ \text{inStream}(h_{\text{outflow},2}) &= \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1} + \max(-\tilde{m}_3, 0)h_{\text{outflow},3}}{\max(-\tilde{m}_1, 0) + \max(-\tilde{m}_3, 0)} \\ \text{inStream}(h_{\text{outflow},3}) &= \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1} + \max(-\tilde{m}_2, 0)h_{\text{outflow},2}}{\max(-\tilde{m}_1, 0) + \max(-\tilde{m}_2, 0)} \end{aligned}$$




Figure D.2: Exemplary connection set with three connected components.

In the general case of a connection set with n components, similar considerations lead to the following.

$$\text{inStream}(h_{\text{outflow},i}) = \frac{\sum_{j=1,\dots,n;j \neq i} \max(-\tilde{m}_j, 0)h_{\text{outflow},j}}{\sum_{j=1,\dots,n;j \neq i} \max(-\tilde{m}_j, 0)}$$

D.3 Special cases covered by inStream definition

D.3.1 Stream connector is not connected ($N = 1$)

For this case, the return value of `inStream` is arbitrary. Therefore, it is set to the outflow value.

D.3.2 Connection of 2 stream connectors, one to one connections ($N = 2$)

$$\text{inStream}(h_{\text{outflow},1}) = \frac{\max(-\tilde{m}_2, 0)h_{\text{outflow},2}}{\max(-\tilde{m}_2, 0)} = h_{\text{outflow},2}$$

$$\text{inStream}(h_{\text{outflow},2}) = \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1}}{\max(-\tilde{m}_1, 0)} = h_{\text{outflow},1}$$

In this case, `inStream` is continuous (contrary to h_{mix}) and does not depend on flow rates. The latter result means that this transformation may remove nonlinear systems of equations, which requires that either simplifications of the form $a * b/a = b$ must be provided, or that this case is treated directly.

D.3.3 Connection of 3 stream connectors where one mass flow rate is identical to zero

The case where $N = 3$ and $\tilde{m}_3 = 0$ occurs when a one-port sensor (like a temperature sensor) is connected to two connected components. For the sensor, the `min` attribute of the mass flow rate should be set to zero (no fluid exiting the component via this connector). This simplification (and similar ones) can also be used if a tool determines that a mass flow rate is zero or non-negative. It is also possible to generalize this to the case where more than one sensor is connected. The suggested implementation results in the following equations, and as indicated the last formula can be simplified further by using $\tilde{m}_3 = 0$:

$$\begin{aligned} \text{inStream}(h_{\text{outflow},1}) &= h_{\text{outflow},2} \\ \text{inStream}(h_{\text{outflow},2}) &= h_{\text{outflow},1} \\ \text{inStream}(h_{\text{outflow},3}) &= \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1} + \max(-\tilde{m}_2, 0)h_{\text{outflow},2}}{\max(-\tilde{m}_1, 0) + \max(-\tilde{m}_2, 0)} \\ &= \begin{cases} h_{\text{outflow},2} & \text{if } \tilde{m}_1 \geq 0 \\ h_{\text{outflow},1} & \text{if } \tilde{m}_1 < 0 \text{ and } \tilde{m}_3 = 0 \end{cases} \end{aligned}$$




Figure D.3: Example series connection of multiple models with stream connectors.

For the two components with finite mass flow rates (not the sensor), the properties discussed for two connected components still hold. The connection set equations reflect that the sensor does not any influence by discarding the flow rate of the latter. In several cases a non-linear equation system is removed by this transformation. However, `inStream` results in a discontinuous equation for the sensor, which is consistent with modeling the convective phenomena only. The discontinuous equation is uncritical, if the sensor variable is not used in a feedback loop with direct feedthrough, since the discontinuous equation is then not part of an algebraic loop. Otherwise, it is advisable to regularize or filter the sensor signal.

D.3.4 Connection of 3 stream connectors where two mass flow rates are positive (ideal splitting junction for uni-directional flow)

If uni-directional flow is present and an ideal splitter is modelled, the required flow direction should be defined in the connector instance with the `min` attribute (the `max` attribute could be also defined, however it does not lead to simplifications):

```

model m2
  Fluidport c(m_flow(min=0));
  ...
end m2;
    
```

Consider the case of and all other mass flow rates positive (with the `min` attribute set accordingly). Connecting `m1.c` with `m2.c` and `m3.c`, such that

```

m2.c.m_flow.min = 0; // max(-m2.c.m_flow,0) = 0
m3.c.m_flow.min = 0; // max(-m3.c.m_flow,0) = 0
    
```

results in the following equation:

$$\text{inStream}(h_{\text{outflow},1}) = \frac{\max(-\tilde{m}_2, 0)h_{\text{outflow},2} + \max(-\tilde{m}_3, 0)h_{\text{outflow},3}}{\max(-\tilde{m}_2, 0) + \max(-\tilde{m}_3, 0)} = \frac{0}{0}$$

`inStream` cannot be evaluated for a connector, on which the mass flow rate has to be negative by definition. The reason is that the value is arbitrary, which is why it is defined as follows.

$$\text{inStream}(h_{\text{outflow},1}) := h_{\text{outflow},1}$$

For the remaining connectors, `inStream` reduces to a simple result.

$$\begin{aligned} \text{inStream}(h_{\text{outflow},2}) &= \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1} + \max(-\tilde{m}_3, 0)h_{\text{outflow},3}}{\max(-\tilde{m}_1, 0) + \max(-\tilde{m}_3, 0)} = h_{\text{outflow},1} \\ \text{inStream}(h_{\text{outflow},3}) &= \frac{\max(-\tilde{m}_1, 0)h_{\text{outflow},1} + \max(-\tilde{m}_2, 0)h_{\text{outflow},2}}{\max(-\tilde{m}_1, 0) + \max(-\tilde{m}_2, 0)} = h_{\text{outflow},1} \end{aligned}$$

Again, the previous non-linear algebraic system of equations is removed. This means that utilizing the information about uni-directional flow is very important.

To summarize, if all mass flow rates are zero, the balance equations for stream variables (D.1) and for flows (D.2) are identically fulfilled. In such a case, any value of h_{mix} fulfills (D.1), i.e., a unique mathematical solution does not exist. This specification only requires that a solution fulfills the balance equations. Additionally, a recommendation is given to compute all unknowns in a unique way, by providing an explicit formula for `inStream`. Due to the definition, that only flows where the corresponding `min` attribute is neither zero nor positive enter this formula, a meaningful physical result is always obtained, even in case of zero mass flow rate. As a side effect, non-linear equation systems are automatically removed in special cases, like sensors or uni-directional flow, without any symbolic transformations (no equation must be analyzed; only the `min` attributes of the corresponding flow variables).

Appendix E

Modelica Revision History

This appendix describes the history of the Modelica Language Design, and its contributors. This appendix is just present for historic reasons and is not normative. The current version of this document is available from <https://www.modelica.org/documents>.

E.1 Modelica 3.4

Modelica 3.4 was released April 10, 2017. The Modelica 3.4 specification was edited by Hans Olsson.

E.1.1 Main changes in Modelica 3.4

The following Modelica Change Proposals are backward compatible extensions added in 3.4:

- Automatic conversions between different versions (MCP-0014), section 18.8.2. Ticket #1622.
- Flattening is clearly specified (MCP-0019), section 5.3.1 and section 5.6. Ticket #1829.
- Convert from `Integer` to Enumeration (MCP-0022), primarily section 4.8.5.3. Ticket #1842.
- Explicitly casting a Model to Record (MCP-0023), section 12.6.1. Ticket #1953.
- Initialization of Clocked Continuous States (MCP-0024), section 16.8.1, section 16.9 and section 16.10. Ticket #2007.
- An added option to Ellipse Annotation to draw only an arc (MCP-0026), section 18.6.5.4. Ticket #2045.
- Allowing mixed `Real` and non-`Real` Record Derivatives (MCP-0028), section 12.7.1. Ticket #2134.

The definition of pure functions was refined, in particular to restore backwards compatibility with Modelica 3.2, section 12.3. Ticket #1937.

The following minor improvements were made (starting from 3.3 Revision 1):

- Clarified simulation model, section 1.2. Ticket #730.
- Clarified structural analysis, section 1.2. Ticket #588.
- Clarified meta-symbols, section 1.4 and 2.3.1. Ticket #1616.
- Typo, section 2.3.1. Ticket #1702.
- Clarified newline, section 2.4.4. Ticket #1479.
- Allow `''` and define it to be equal to `\''` (and similarly for `\?`), section 2.3.1. Ticket #1176.
- Clarified that built-in functions in the specification, section 3.7 and 12.5. Ticket #1608.
- Clarified named arguments for builtin operators, in particular `spatialDistribution`, section 3.7.4 and section 12.4.1. Ticket #2002.

- Clarified that `semiLinear` is continuous, section 3.7.4. Ticket #112.
- Corrected spelling, section 3.7.4. Ticket #1828.
- Corrected typo in code and reformulated description, section 3.7.4.2. Ticket #1588, #1729, and #2166.
- Corrected typo for events, operator 3.22. Ticket #1657.
- Clarified `sample` operator, section 3.7.5. Ticket #677.
- Additional functions give parameter expression, section 3.8.2. Ticket #1082.
- Clarified `ExternalObject`, section 4.4.2.2. Ticket #1546.
- Simplified rules type prefixes for structured components, section 4.4.2.2. Ticket #1686.
- Clarified that unexpanded bindings shall be unexpanded, section 4.4.3. Ticket #2153.
- Clarified conditional components, section 4.4.5. Ticket #2057.
- Improved example to avoid using class, section 4.5.3. Ticket #553.
- Clarify inheritance from predefines types, section 4.6. Ticket #1250.
- Allow `connector` inheriting from `operator record`, section 4.6. Ticket #1714.
- Clarify restrictions on record components, section 4.6. Ticket #1615.
- Clarify that only functions may have external-clause, section 4.6 and section 12.2. Ticket #2014.
- Clarified equation count for operator record, section 4.7. Ticket #866.
- Clarified restriction on attributes, section 4.8. Ticket #1426.
- Clarified how reserved the different built-in types are, section 4.8. Ticket #1538.
- Removed restriction on nominal and explained purpose, section 4.8.1. Ticket #1445.
- Added `unbounded` to `Real`, section 4.8.1 and section 4.8.6. Ticket #926.
- Added fixed-attribute for `String`, section 4.8.4. Ticket #1797.
- Clarified and corrected example in section 4.8.5. Ticket #1849 and #2150.
- Clarified nominal attribute, section 4.8.6 and section 15.2. Ticket #1877.
- Added that `Connections` is builtin package, section 4.8.7.4. Ticket #1883.
- Clarified lookup-order regarding import, section 5.3.1. Ticket #1573.
- Extend calling functions through component to array case, section 5.3.2 and section 3.1. Ticket #1613.
- Clarified existing use of automatic inner declarations, section 5.4 and section 18.7. Ticket #1551 and #1749.
- Removed restriction on array size for modifiers, section 6.4. Ticket #1432.
- Clarified that $2/3$ and 2^{-3} are `Real`, section 6.7. Ticket #1647.
- Clarified that the external-clause is inherited, section 7.1. Ticket #789.
- Further clarified order for multiple inheritance, section 5.6.1 and section 7.1. Ticket #2015.
- Clarified inheritance restrictions, section 7.1.3. Ticket #1451.
- Restricted merging of modifiers, section 7.2.3. Ticket #791.
- Clarified each especially for nested arrays, section 7.2.4 and section 7.2.5. Ticket #1596.
- Clarified replaceable with array sizes on types, section 7.3. Ticket #1251.
- Corrected and moved example, section 7.3 and section 7.3.2. Ticket #1034.
- Clarified `redeclare class extends B`, section 7.3.1. Ticket #462 and #709.

- Corrected example, section 7.3.2. Ticket #1725.
- Clarified description and annotation on constraining-clause, section 7.3.2.1. Ticket #512.
- Corrected typo, section 7.3.4. Ticket #1770.
- Clarified for-equation with types, section 8.3.2. Ticket #915.
- Clarified event generation, section 8.5. Ticket #2114.
- Further clarified `initial()` for when-clauses, section 8.6, and indicated that this appendix is not normative. Ticket #1852.
- Clarified using `start` attribute for parameters, section 8.6. Ticket #2136.
- Clarified that states for first order ODE, section 8.6.1. Ticket #937.
- Clarified adding `input/output` prefix for expandable connector variables, section 9.1.3. Ticket #829.
- Clarified creating elements in expandable connectors, section 9.1.3. Ticket #428.
- Corrected expandable connector example, section 9.1.3. Ticket #1763.
- Clarified that stream variables do not generate equations, section 9.2. Ticket #1584.
- Restrict that stream only connects to stream, section 9.3. Ticket #796.
- Clarified section heading, section 9.3.1. Ticket #727.
- Clarified vector arguments for operators in section 9.4.1. Ticket #1590.
- Corrected example, section 9.4.3.1. Ticket #2143.
- Clarified return type for reduction expressions, section 10.3.4. Ticket #981.
- Extended reduction expression `sum` to operator records, section 10.3.4. Ticket #1897.
- Clarified table in section 10.3.4.1. Ticket #1722.
- Match parenthesis, section 10.4.1.2. Ticket #1558.
- Recommend better alternative for generating vector, section 10.4.3. Ticket #1837.
- Defined unary operators, section 10.6.2. Ticket #2027.
- Allow missing trailing indices, sections section 10.5 and section 10.6.9. Ticket #1603.
- Clarified that element-wise division gives real result, section 10.6.6. Ticket #1119.
- Removed misleading comment, section 11.1.2. Ticket #938.
- Clarified how assignment works for multi-returning functions, section 11.2.1.1. Ticket #1921.
- Clarified no equations and initial algorithms in functions, section 12.2. Ticket #2160.
- Clarified assigning to record variables with bindings in functions, section 12.2. Ticket #2016.
- Clarified initialization of variables in functions, section 12.4.4. Ticket #1708.
- Standardize current practice of using `=` instead of `:=` for bindings in functions, section 12.2. Ticket #1595.
- Clarified function partial evaluation, section 12.4.2.1. Ticket #647.
- Clarified initialization of record components in functions, section 12.4.4. Ticket #1230.
- Clarified flexible array sizes, section 12.4.5. Ticket #2158.
- Clarified name of output for record constructor in section 12.6. Ticket #366.
- Clarified derivatives for functions in several ways, section 12.7.1. Ticket #985, #1543, #1544, #1545, #1547, #1548, #1788, #1972, #1987.
- Clarified that using `C89` and added possibility for `C89`, `C99`, and `C11`, section 12.9. Ticket #1088.

- Clarified input/output to external functions, section 12.9. Ticket #775.
- Clarified handling of `Boolean` variables for external C, section 12.9.1.1. Ticket #1846.
- Added that strings can be sent to FORTRAN 77, section 12.9.1.1. Ticket #1971.
- Allow multiple include directories, section 12.9.4. Ticket #2103.
- Allow specific libraries for different compiler versions, section 12.9.4. Ticket #1316.
- Added external warning functions, section 12.9.6. Ticker #1967.
- Clarified that pointers are only valid during each call, section 12.9.7. Ticket #1611.
- Clarified constructor and destructor, section 12.9.7. Ticket #1907.
- Clarified 'structured entity' to be 'directory', section 13.4. Ticket #922.
- Clarified handling of incorrect package.order, section 13.4.1. Ticket #1858.
- Restricted use of files with multiple classes, section 13.4.2. Ticket #1854.
- Clarified storing resources in a file-system, section 13.5. Ticket #685, #1623.
- Used correct font, section 13.5. Ticket #2061.
- Clarified that `inStream` optimizations are allowed, section 15.2. Ticket #1766.
- Corrected `actualStream` example in section 15.3. Ticket #1652.
- Corrected non-periodic rational clocks, section 16.3. Ticket #2022.
- Clarified initialization of clocked discretized continuous-time partitions, section 16.8.2. Ticket #1528.
- Defined rendering order, section 18.6. Ticket #1750.
- Clarified rotation direction, section 18.6.1. Ticket #1830.
- Clarified coordinate system definition, section 18.6.1.1. Ticket #1831.
- Defined coordinate system inheritance to be less suprising, section 18.6.1.1 and section 18.6.3. Ticket #1978.
- Clarified `lineThickness` and `borderPattern`, section 18.6.5 section 18.6.1.2. Ticket #1896.
- Corrected formatting, section 18.6.1.2. Ticket #1825.
- Made the different Text-annotations more similar, section 4.8.7.4. Ticket #1621.
- Clarified `LinePattern` in section 18.6.5. Ticket #1483.
- Corrected flipping of components and bitmaps, and clarified various aspects of bitmaps, section 18.6.2 and section 18.6.5.6. Ticket #1923.
- Clarified arrows, section 18.6.5.1. Ticket #1894.
- Clarified existing use of zero-width texts, section 18.6.5.5. Ticket #1636.
- Added specific fontnames, section 18.6.5.5. Ticket #1986.
- Added alternative for macro-expansion, section 18.6.5.5. Ticket #2148.
- Corrected example in section 18.6.7.1. Ticket #2111.
- Placed annotation last in classes, in particular section 18.7. Ticket #1009.
- Cleaned up code and formatting in section 18.7. Ticket #2041, #2042, and #2125.
- Added missing nano-prefix, section 19.1. Ticket #1261.
- Replaced the outdated contents of chapter 20 by a hyperlink. Ticket #2130.
- Added `ModelicaServices` to chapter 20. Ticket #2132.
- Restrict grammar to avoid modifiers with leading dot, appendix B.2. Ticket #1027.

- Restrict grammar for base-prefix, appendix B.2.2. Ticket #917.
- Restrict grammar for arrays, section 3.1. Ticket #809.
- Restrict grammar for function arguments (replacing semantic restriction), section 3.1. Ticket #1634.
- Corrected typo in appendix D.2. Ticket #1317.

E.1.2 Contributors to the Modelica Language 3.4

The members of the Modelica Association contributed to the Modelica 3.4 specification.

E.2 Modelica 3.3 Revision 1

Modelica 3.3 Revision 1 was released July 11, 2014. The Modelica 3.3 Revision 1 specification was edited by Hans Olsson.

E.2.1 Main changes in Modelica 3.3 Revision 1

The changes made in Modelica 3.2 Revision 2 are included, and in addition the following improvements were made:

- Clarified that `String` operator cannot use positional arguments, section 3.7.1. Ticket #1468.
- Corrected size of enumeration, section 3.7.1 and section 4.8.5.2. Ticket #1369.
- Clarified `spatialDistribution`, section 3.7.3. Ticket #1510.
- Restricted `cardinality` to give a clear definition, section 3.7.4.3. Ticket #1409.
- Clarified which constants need a value, section 3.8.1. Ticket #1220.
- Clarified type prefixes rules, section 4.4.2.2. Tickets #1196, #1221, #1301.
- Added exception for cyclic parameter bindings (already used in MSL), section 4.4.3. Ticket #1320.
- Added example for use of conditional components, section 4.4.5. Ticket #1227.
- Corrected annotation-grammar, section 4.5. Ticket #1378.
- Corrected duplicated class-definition grammar, section 4.5. Ticket #1388.
- Clarified short class definition, section 4.5.1. Ticket #527.
- Removed unusable variant for operator and operator function, section 4.6. Ticket #1459, #1497.
- Added definition of `AssertionLevel`, section 4.8.7.3. Ticket #962.
- Corrected typos at end of loops in examples, section 4.8.5 and section 8.3.2.2. Ticket #902
- Clarified temporary flattening, section 5.3.2. Ticket #1327.
- Added definition of modification equations, section 7.2. Ticket #959.
- Clarified modifiers for array elements, section 7.2.5. Ticket #1324.
- Corrected example for final element modification, section 7.2.6 Ticket #1326.
- Corrected duplicated class-definition grammar, section 7.3.1. Ticket #1388.
- Clarified modifiers on constraining type, section 7.3.2. Ticket #1033.
- Clarified when `redeclare` can be used with the same type and rules for redeclaring array types, section 7.3.3. Tickets #1252, #1281.
- Clarified default for annotation `choicesAllMatching`, section 7.3.4. Ticket #1391.
- Forbid when-statements in initial equation/algorithm (they would in most cases not be active; leading to confusion), section 8.3.5.3 and section 11.2.7.1. Ticket #1288.

- Clarified `reinit` during initialization, section 8.3.6 section 8.6. Ticket #1372.
- Clarified using start-values as guess-values; section 8.6. Tickets #1133, #1246.
- Clarified allowed use of variables in expandable connectors; section 9.1.3. Ticket #1279.
- Clarified causality for expandable connectors; section 9.1.3. Ticket #1305.
- Clarified expandable connectors in general; section 9.1.3. Ticket #1330.
- Clarified connection matching, section 9.3. Ticket #884.
- Added quantity checks for connectors (MSL already relies on this check); section 9.3. Ticket #1284.
- Clarified arrays with non-`Integer` dimensions, section 10.1. Ticket #1501.
- Clarified that `ndims` is allow for a scalar, section 10.3.1. Ticket #1303.
- Clarified number of arguments for `zeros`, `ones`, `fill`, section 10.3.3. Ticket #1351.
- Clarified `min/max`, section 10.3.4. Ticket #1036.
- Clarified array expressions using iterations allow non-simple types, section 10.4.1. Ticket #1521.
- Clarified arrays with non-`Integer` dimensions, section 10.5.1. Ticket #1463.
- Clarified calling function as specialized class, section 12.2. Ticket #1362.
- Clarified default values in functions depending on other inputs, section 12.4.1. Ticket #1346.
- Corrected syntax error in example, section 12.4.2.1. Ticket #1338.
- Clarified annotations on external functions, section 12.9. Ticket #660.
- Add possibility for sending arrays in records to external functions, section 12.9. Ticket #351.
- Corrected spelling to FORTRAN 77, section 12.9.3. Ticket #1278.
- Clarified default directories, section 12.9.4. Ticket #1456.
- Clarified constructing/destructuring external objects, section 12.9.7. Ticket #1518.
- Clarified encapsulation requirement for operator record, chapter 14. Ticket #1254.
- Clarified operator record: arrays, priority, and zero result, section 14.5. Tickets #1469, #1476, #1481.
- Added element wise operations for operator record, section 14.5. Ticket #1455.
- Improved formulation, section 16.3. Ticket #1362.
- Clarified why `noClock` exists; section 16.5.2. Ticket #1094.
- Added initial conditions to solver methods for clocked discretized continuous-time partitions; section 16.8.2. Ticket #1379.
- Added requirement that priorities must be unique for statechart transitions (the restriction was present in original description and is necessary to ensure deterministic behavior), section 17.3. Ticket #853.
- Corrected syntax in Line definition, section 18.6.5.1. Ticket #1464.
- Corrected `connectorSizing` description, section 18.7. Ticket #1441.
- Corrected license example, section 18.9.2. Ticket #1127.
- Clarified names of productions in grammar and changed to use hyphen, appendix B.1. Tickets #713 and #1033.
- Modified grammar use consistent style for import-list, appendix B.2.2. Ticket #1374.

E.2.2 Contributors to the Modelica Language, Version 3.3 Revision 1

The following members of the Modelica Association contributed to the Modelica 3.3 Revision 1 and/or Modelica 3.2 Revision 2 specification (alphabetical list):

Peter Aronsson, MathCore AB, Linköping, Sweden
Peter Fritzson, PELAB, Linköping University, Linköping, Sweden
Hilding Elmqvist, Dassault Systèmes, Lund, Sweden
Christoph Höger, Technical University of Berlin, Berlin, Germany
Gerd Kurzbach, ITI GmbH, Dresden, Germany
Jesper Mattsson, Modelon AB, Lund, Sweden
Hans Olsson, Dassault Systèmes, Lund, Sweden
Martin Otter, DLR, Oberpfaffenhofen, Germany
Adrian Pop, Linköping University, Linköping, Sweden
Elena Shmoylova, Maplesoft, Waterloo, Canada
Martin Sjölund, PELAB, Linköping University, Linköping, Sweden
Stefan Vorkoetter, Maplesoft, Waterloo, Canada

E.3 Modelica 3.3

Modelica 3.3 was released on May 9, 2012. The Modelica 3.3 specification was edited by Hans Olsson, Hilding Elmqvist and Martin Otter.

E.3.1 Main changes in Modelica 3.3

The following *backward compatible extensions* have been introduced with Modelica 3.3:

- Language elements for describing synchronous behavior suited for implementation of control systems, chapter 16.
- Language elements to define synchronous state machines, chapter 17.
- The `spatialDistribution` function for special one-dimensional flow problems, section 3.7.4.2.
- The `getInstanceName` function for diagnostic messages, section 3.7.4.6.
- Possible to call a function through an instance name, section 5.3.2.
- Can use the start-value for a parameter to give a non-zero default that should be changed, section 8.6.
- A recommended procedure for priority between initial values have been added, section 8.6.2.
- Functions can be defined without algorithm-section, section 12.2.
- Functions can be marked as `pure` or `impure` with specified semantics, section 12.3.
- The rules for `ExternalObject` have been clarified, section 12.9.7 and chapter 6.
- Multiple definition import, section 13.2.
- Additional annotations allowing:
 - Functions to generate events, section 18.3.
 - Experiments to specify a time-resolution of simulation result, section 18.4.1.
 - Single instance of class, section 18.5.
 - Text in the diagram layer can use a macro syntax, section 18.6.5.5.
 - Color selection dialog for parameters, section 18.7.
 - Conversion to specify a set of versions to convert with one script, section 18.8.2.

- Licensed libraries to define the set of allowed operations (including binary/source export), section 18.9.2.

The following changes in Modelica 3.3 are *not* backwards compatible:

- Chapter 16 *Synchronous Language Elements* from Modelica 3.2 has been removed (a more powerful functionality is instead provided with the new chapter 16 and chapter 17). Since, no released tools has yet supported the previous chapter 16, this not backwards compatible change is uncritical.
- The new `spatialDistribution` and `getInstanceName` functions could cause problems if another function with that name was already used.
- Conditional physical connectors must be connected if enabled, section 4.4.5 and section 9.3. In almost all cases they have to be connected to generate correct result, and it is not possible to check that they are connected in the models

E.3.2 Contributors to Modelica 3.3

The language elements for describing synchronous behavior, chapter 16, was mainly developed by Hilding Elmqvist, Martin Otter, and Sven Erik Mattsson. Hilding Elmqvist wrote a detailed tutorial. Sven Erik Mattsson developed a test implementation of the language elements and the needed new algorithms. Based on the prototype, tests and feedback have been provided by Martin Otter and Bernhard Thiele.

The language elements to define synchronous state machines, chapter 17, was mainly developed by Hilding Elmqvist with contributions from Francois Dupont, Sven Erik Mattsson and Fabien Gaucher. Hilding Elmqvist wrote a tutorial. Sven Erik Mattsson and Carl-Fredrik Abelson developed a test implementation. Based on the prototype, tests and feedback have been provided by Alain Thura, Emmanuel Ledinot, Claire Campan, and Martin Malmheden.

The `spatialDistribution` operator was initiated by Hubertus Tummescheit, based on an operator of Dymola, and with contributions from Hans Olsson.

The improved definition of pure and impure functions and the enhanced import-clause was proposed by Peter Fritzson.

The various smaller language improvements based on submitted tickets have been developed by a group headed by Hans Olsson and with group members Peter Fritzson, Christoph Höger, Gerd Kurzbach, Jesper Mattsson, Martin Sjölund, and Stefan Vorkoetter.

The following members of the Modelica Association participated at design meetings and contributed to the Modelica 3.3 specification:

Johan Åkesson, Lund University and Modelon AB, Lund, Sweden
 Peter Aronsson, MathCore AB, Linköping, Sweden
 Ingrid Bausch-Gall, BAUSCH-GALL GmbH, Munich, Germany
 Volker Beuter, Kämmerer AG, Germany
 Torsten Blochwitz, ITI GmbH, Dresden, Germany
 David Broman, PELAB, Linköping University, Sweden
 Dag Brück, Dassault Systèmes, Lund, Sweden
 Francesco Casella, Politecnico di Milano, Milano, Italy
 Christoph Clauss, Fraunhofer, Dresden, Germany
 Mike Dempsey, Claytex Services Limited, Leamington Spa, U.K.
 Karin Dietl, TU Hamburg-Harburg, Germany
 Francois Dupont, Dassault Systèmes, Brest, France
 Jonas Eborn, Modelon, Lund, Sweden
 Hilding Elmqvist, Dassault Systèmes, Lund, Sweden
 Guilioano Fontanella, AIT, Vienna, Austria
 Rüdiger Franke, ABB Power Generation, Mannheim, Germany
 Peter Fritzson, PELAB, Linköping University, Sweden
 Sébastien Furic, LMS International, Roanne. France

Leo Gall, BAUSCH-GALL Gmbh, Munich, Germany
Peter Harman, deltatheta uk limited, U.K.
Anton Haumer, AIT, Vienna, Austria
Dan Henriksson, Dassault Systèmes, Lund, Sweden
Christoph Höger, TU Berlin, Berlin, Germany
Christian Kral, AIT, Vienna, Austria
Gerd Kurzbach, ITI GmbH, Dresden, Germany
Kilian Link, Siemens AB, Erlangen, Germany
Krisitin Majetta, Fraunhofer, Dresden, Germany
Martin Malmheden, Dassault Systèmes, Velicy, France
Jesper Mattsson, Modelon, Lund, Sweden
Sven Erik Mattsson, Dassault Systèmes, Lund, Sweden
Eric Neuber, ITI GmbH, Dresden, Germany
Ramine Nikoukhah, Altair, France
Hans Olsson, Dassault Systèmes, Lund, Sweden
Martin Otter, DLR-RM (German Aerospace Center), Oberpfaffenhofen, Germany
Peter Pepper, Fraunhofer FIRST, Berlin, Germany
Adrian Pop, Linköping University, Linköping, Sweden
Olena Rogovchenko, PELAB, Linköping, Sweden
Stefan-Alexander Schneider, BMW, Munich, Germany
Michael Sielemann, German Aerospace Center, Oberpfaffenhofen, Germany
Martin Sjölund, PELAB; Linköping, Sweden
Kristian Stavåker, PELAB, Linköping, Sweden
Bernhard Thiele, DLR-RM (German Aerospace Center), Oberpfaffenhofen, Germany
Eric Thomas, Dassault Aviation, Paris, France
Michael Tiller, Dassault Systèmes, Velicy, France
Hubertus Tummescheit, Modelon AB, Lund, Sweden
Andreas Uhlig, ITI, Dresden Germany
Vladimir Vukovic, AIT, Vienna, Germany
Stefan Vorkoetter, MapleSoft, Waterloo, Canada
Daniel Weil, Dassault Systèmes, Grenoble, France
Hans-Jürg Wiesmann, ABB Switzerland, Corporate Research, Baden, Switzerland
Dietmar Winkler, Telemark University College, Porsgrunn, Norway
Stefan Wischhusen, XRG Simulation, Hamburg, Germany
Dirk Zimmer, DLR-RM (German Aerospace Center), Oberpfaffenhofen, Germany

E.3.3 Acknowledgments

For the design of the synchronous language elements (chapter 16) and synchronous state machines (chapter 17), and for the understanding of fine details of synchronous languages, especially from Lucid Synchrone, very helpful discussions with

Albert Benveniste, IRISA/INRIA, Rennes, France
Marc Pouzet, Laboratoire d'Informatique de l'ENS, Paris
Benoit Caillaud, IRISA, Rennes, France
Timothy Bourke, INRIA, Rennes, France

are appreciated.

E.4 Modelica 3.2 Revision 2

Modelica 3.2 Revision 2 was released 2013. The Modelica 3.2 Revision 2 specification was edited by Hans Olsson and Martin Otter.

E.4.1 Main changes in Modelica 3.2 Revision 2

The Modelica language was slightly adapted (in a backwards compatible way, except as listed below) in order that the Modelica Standard Library (MSL) version 3.2.1 is fully compliant to Modelica Language 3.2 Revision 2. This required the following improvements compared to 3.2 Revision 1:

- Possible to call a function through an instance name, section 5.3.2 (used in MSL 3.2 to compute the gravity acceleration in `Modelica.Mechanics.MultiBody.World`; this feature was also introduced in Modelica Language version 3.3 in May 2012).
- New built-in operator `Connections.rooted(A.R)` to inquire whether an overdetermined type or record instance `A.R` in a call to `Connections.branch(A.R, B.R)` is closer to the root of the spanning tree than `B.R`, section 9.4.1 (used in MSL 3.2 to avoid algebraic loops in several components such as in `Modelica.Mechanics.MultiBody.Joints.Revolute`).
- Several new annotations where vendor-specific variants were used in MSL 3.2; section 7.3.4, section 18.2, section 18.3, section 18.4.1, section 4.8.7.4, and section 18.7.
- Specified that Evaluate can also occur in types, since this is used in MSL and important for performance; section 18.3. Ticket #925.
- Macros in graphical text items, section 18.6.5.5. Ticket #659. (This feature was also introduced in Modelica Language version 3.3.)
- Initial equations are discrete – used in MSL for initialization of pre-variables, section 3.8.3. Ticket #853.
- Updated `noDerivative` to be consistent with MSL, section 12.7.1. This is an incompatibility, but the other variant was not used. Ticket #1035.
- Clarified handling of component with `missingInnerMessage`; section 18.7. Ticket #891.
- Clarified definition of protected; section 4.1. Ticket #975, #1123.

In addition several issues with the specification text were corrected:

- Clarified an unclear sequence regarding functions as input arguments in section 4.4.2.2. Ticket #1182.
- Clarified allowed binding equations for redeclarations, section 4.7. Ticket #1111.
- Unspecified enumerations now have defined semantics, section 4.8.5.4. Ticket #834.
- Transitively non-Replaceable, section 6.3.1. Ticket #854.
- Modification text improved to not refer to inherited class, section 7.2. Ticket #1042.
- Precedence for modifiers on constraining-clause clarified, section 7.3.2. Ticket #1128.
- Clarified arrays for constraining type, section 7.3.2. Ticket #1148.
- Avoid all forms of connections depending on connections, section 8.3.3. Ticket #828.
- Clarified equation count for if-equations, section 8.3.4. Ticket #888.
- Complete definition of `reinit`, section 8.3.6. Ticket #578. This forbids `reinit` in algorithms, but it was previously not well-defined.
- Clarified initializaton of `pre(vc)` for a non-discrete (that is continuous-time) Real variable `vc`, section 8.6. Ticket #1195.
- Only one way of handling arrays of connectors is now defined, chapter 9. Ticket #757.
- Example now use correct sine-source, section 9.1. Ticket #750.
- Restricted parameters in connectors, section 9.3. Ticket #768.
- Clarified type restrictions for some operators, section 10.3.4, section 10.3.5. Ticket #622.
- Clarified that if at least one array element is used on the left hand side of the assignment operator in an algorithm section, then the complete array is initialized in this section, section 11.1.2. Ticket #1190.

- Record constructor corrected to not refer to keywords that should not occur, section 12.6. Ticket #907.
- External storage of classes, section 13.4.2, section 13.4.3 and section 18.8.3. Tickets #1019, #892, #887.
- Added example and explanation for inheritance restriction on operator record; section 14.7. Ticket #1065.
- Chapter *Mapping of Models to Execution Environments* was removed; this change was already decided for Modelica 3.3 and no tool had released an implementation of this feature. Ticket #1015.
- Corrected license-example in section 18.9.2. Ticket #1127.
- Grammar was internally restructured for short-class-definition, section 4.5, appendix B.2.2. Ticket #1140.

E.4.2 Contributors to the Modelica Language, Version 3.2 Revision 2

The following members of the Modelica Association contributed to the Modelica 3.2 Revision 2 specification (alphabetical list):

Peter Aronsson, MathCore AB, Linköping, Sweden
 Peter Fritzson, PELAB, Linköping University, Linköping, Sweden
 Hilding Elmqvist, Dassault Systèmes, Lund, Sweden
 Christoph Höger, Technical University of Berlin, Berlin, Germany
 Gerd Kurzbach, ITI GmbH, Dresden, Germany
 Jesper Mattsson, Modelon AB, Lund, Sweden
 Hans Olsson, Dassault Systèmes, Lund, Sweden
 Martin Otter, DLR, Oberpfaffenhofen, Germany
 Adrian Pop, Linköping University, Linköping, Sweden
 Elena Shmoylova, Maplesoft, Waterloo, Canada
 Martin Sjölund, PELAB, Linköping University, Linköping, Sweden
 Stefan Vorkoetter, Maplesoft, Waterloo, Canada

E.5 Modelica 3.2 Revision 1

Modelica 3.2 Revision 1 was released on Feb. 29, 2012. The Modelica 3.2 Revision 1 specification was edited by Hans Olsson and Peter Fritzson.

E.5.1 Main changes in Modelica 3.2 Revision 1

The Modelica language was not changed with respect to the previous version 3.2. Only issues with the specification text have been fixed. In particular:

- Corrected typos and improved formatting.
- section 2.2 Comments:
There are 2 and not 3 kinds of comments and comments are treated as white space character.
Added definition of white space character.
- section 2.3.1 Identifiers:
The single quotes are part of the identifier, e.g., 'x'.
- section 3.6.7 Built-in Variable time:
Variable `time` is only available in models and blocks and not in the other classes.

- section 3.7.3 Built-in Mathematical Functions
Definition of `atan2` corrected.
- section 3.7.4 Special Purpose Operators
Included definition of `inStream` and `actualStream` operators from chapter 15.
- section 3.7.5 Event-Related Operators
Clarified, that the first argument of `smooth` is a scalar.
Improved the definition of `reinit`.
- section 3.8.3 Discrete-Time Expressions
Improved definition of ordered relations (`>`, `<`, `>=`, `<=`).
- section 4.4.5 Conditional Component Declaration
Clarified redeclaration of a component.
- section 4.6 Specialized Classes
Clarified that `stream` cannot be used in a record.
Clarified restrictions on elements in a `connector`.
Errors in example of operator record `Complex` corrected.
- section 4.8.5 Enumeration Types
Error in example corrected.
- section 5.5 Simultaneous Inner/Outer Declarations
Clarified inner/outer declarations.
- section 7.1 Inheritance
Clarified that the elements of a flattened base class are added at the place of the extends clause.
Equations of the flattened base class that are syntactically equivalent to equations in the flattened enclosing class are deprecated.
- section 7.2 Modifications
Element modifiers are no longer part of language, reference grammar instead of duplicating it.
- section 7.3 Redeclaration
Improved redeclarations definition and moved an example from section 7.3.1 at the right place.
- section 8.3.6 `reinit`
Improved `reinit` definition.
- section 8.6 Initialization
Clarified that only when-clauses with restricted form of `initial()` as condition will be active during initialization.
- section 10.3.4.1 Reduction Expressions
Improved definition
- section 11.2.2.2 Types as iteration ranges
Newly introduced section to improve the definition of iteration ranges
- section 12.2 Function
Added missing restrictions that `model`, `block`, `inner`, `outer` cannot be used in a function.
- section 12.4.1 Positional or Named Input Arguments of Functions
Corrected formal syntax of a function call
- section 12.4.4 Initialization and Declaration Assignments of Components in Functions
Added the restriction of acyclic bindings.
- section 12.9.1.3 Records
Mapping of arrays in records to C-structs is removed.
- section 18.6.5.6 Bitmap
Defined flipping more precisely.
- appendix B.1 Lexical conventions
More precisely defined white-space and comments.

- appendix B.2 Grammar
Improved/corrected grammar definition

E.5.2 Contributors to the Modelica Language, Version 3.2 Revision 1

The following members of the Modelica Association contributed to the Modelica 3.2 Revision 1 specification (alphabetical list):

Peter Aronsson, MathCore AB, Linköping, Sweden

Peter Fritzon, PELAB, Linköping University, Linköping, Sweden

Christoph Höger, Technical University of Berlin, Berlin, Germany

Gerd Kurzbach, ITI GmbH, Dresden, Germany

Jesper Mattsson, Modelon AB, Lund, Sweden

Hans Olsson, Dassault Systèmes, Lund, Sweden

Adrian Pop, Linköping University, Linköping, Sweden

Martin Sjölund, PELAB, Linköping University, Linköping, Sweden

Stefan Vorkoetter, Maplesoft, Waterloo, Canada

E.6 Modelica 3.2

Modelica 3.2 was released on March 24, 2010. The Modelica 3.2 specification was edited by Hans Olsson, Martin Otter and others.

E.6.1 Main changes in Modelica 3.2

The following *backward compatible extensions* have been introduced with Modelica 3.2:

- Homotopy function for making it easier to solve initialization problems (see section 3.7.4).
- Functions as formal inputs to functions (see new section 12.4.2).
- Overloaded operators have been refined (see chapter 14):
 - A new specialized class `operator record` is introduced – with specialized typing rules (the type is identified by the class name; all other Modelica classes have a structural type system where the type is only defined by the public elements). Overloaded operators can only be defined inside an `operator record`. This change fixes a flaw in Modelica 3.1, since the look-up of overloaded operators is performed by the record class name.
 - Inheritance of an `operator record` is allowed if defined via a short class definition. This removes a restriction of operator overloading in Modelica 3.1, e.g., to define derived classes with units for the record elements, like deriving `ComplexVoltage` from `Complex`.
 - New overloaded element '0' in order that operator record classes can be used as flow variables in connectors.
- Unicode support in description strings, strings in annotations and in comments in order to improve Modelica, e.g., for Arabian, Asian or Indian users (see grammar changes in appendix B.1). Modelica files are UTF-8 encoded, and can start with the UTF-8 encoded byte order mark (0xef 0xbb 0xbf) to indicate that it may contain UTF-8 characters; this is treated as white-space in the grammar (see section 13.4.2).
- Constants can once again be modified unless declared final – as this is already used in packages. (See section 3.8.1).
- Global name lookup has been introduced (e.g. `.Modelica.Constants.pi`), see section 5.3.3.
- New C functions `ModelicaVFormatMessage` and `ModelicaVFormatError`, to simplify message formatting in external functions (see section 12.9.6)

- Additional annotations allowing:
 - Inclusion of C-header and object library files in packages and referencing them with URIs. (new annotations `IncludeDirectory`, `LibraryDirectory`, and standardized platform names like `win32`; see section 12.9.4; resolves ticket #297).
 - Images in parameter dialogs (new annotation `groupImage`; see section 18.7).
 - Start and fixed attributes for variables in parameter dialogs (new annotation `showStartAttribute`; see section 18.7).
 - Access control for packages to protect intellectual property. (new annotations `Protection`, and `License`; see new section 18.9).

The following changes in Modelica 3.2 are *not* backwards compatible:

- The new built-in operator name `homotopy`. In rare cases this might give name clashes in existing models.
- Records with overloaded operations must be declared as `operator record` instead of as `record`. This is uncritical because tools that already support operator overloading will support the Modelica 3.1 form still for some time.

E.6.2 Contributors to the Modelica Language, Version 3.2

The initial version of *functions as formal inputs to functions* was proposed by Peter Fritzson.

The definition of header-files and object-libraries with the `IncludeDirectory` and `LibraryDirectory` annotations was mainly developed by Hans Olsson.

The `Protection` annotation used for access control is an improved version of existing annotations from Dymola. The `License` annotation was mainly developed by Dag Brück with improvements from Magnus Gäfvert.

The flaw in the operator overloading concept was detected by Sébastien Furic. He also proposed the basic fix, by using a nominal type system for records with overloaded operations.

The global name lookup was proposed by Stefan Vorkoetter.

The support for Unicode was initiated by Rui Gao and Hoyoun Kim.

Martin Otter, Michael Sielemann and Francesco Casella proposed `homotopy`. Michael Sielemann demonstrated with benchmark problems that non-linear solvers are not able to solve reliably initialization problems and that the `homotopy` operator is therefore needed. He provided a prototype implementation of `homotopy` and demonstrated its use on a vehicle dynamics example. Utilizing the prototype implementation, Francesco Casella demonstrated with a model of a thermal power plant with 390 iteration variables of the initialization problem, that an appropriate usage of `homotopy` allows the system to reliably initialized without providing guess values for the iteration variables. This was a strong indication that `homotopy` will indeed improve initialization in Modelica significantly.

The following members of the Modelica Association participated at design meetings and contributed to the Modelica 3.2 specification:

Johan Åkesson, Lund University and Modelon AB, Lund, Sweden

Peter Aronsson, MathCore AB, Linköping, Sweden

Bernhard Bachmann, University of Applied Sciences, Bielefeld, Germany

Jonathan Beck, Dassault Systèmes, Paris, France

Torsten Blochwitz, ITI GmbH, Dresden, Germany

David Broman, PELAB, Linköping University, Sweden

Dag Brück, Dassault Systèmes, Lund, Sweden

Francesco Casella, Politecnico di Milano, Milano, Italy

Mike Dempsey, Claytex Services Limited, Leamington Spa, U.K.

Karin Dietl, TU Hamburg-Harburg, Germany
Filippo Donida, Politecnico di Milano, Milano, Italy
Hilding Elmqvist, Dassault Systèmes, Lund, Sweden
Peter Fritzson, PELAB, Linköping University, Sweden
Sébastien Furic, LMS International, Roanne. France
Manuel Gräber, TU Braunschweig, Germany
Peter Harman, deltatheta uk limited, U.K.
Anton Haumer, AIT, Vienna, Austria
Carsten Heinrich, Institut für Luft- und Kältetechnik, Dresden, Germany
Dan Henriksson, Dassault Systèmes, Lund, Sweden
Fredrik Karlsson, PELAB, Linköping University, Sweden
Christian Kral, AIT, Vienna, Austria
Imke Krüger, TU Hamburg-Harburg, Hamburg, Germany
Gerd Kurzbach, ITI GmbH, Dresden, Germany
Kilian Link, Siemens AB, Erlangen, Germany
Sven Erik Mattsson, Dassault Systèmes, Lund, Sweden
Eric Neuber, ITI GmbH, Dresden, Germany
Hans Olsson, Dassault Systèmes, Lund, Sweden
Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany
Adrian Pop, Linköping University, Linköping, Sweden
Katrín Prölb, Modelon AB, Lund, Sweden
Michael Sielemann, German Aerospace Center, Oberpfaffenhofen, Germany
Bernhard Thiele, German Aerospace Center, Oberpfaffenhofen, Germany
Thorben Vahlenkamp, XRG Simulation, Hamburg, Germany
Eric Thomas, Dassault Aviation, Paris, France
Michael Tiller, Emmeskay, Plymouth, MI, U.S.A
Hubertus Tummescheit, Modelon AB, Lund, Sweden
Stefan Vorkoetter, Maplesoft, Waterloo, Canada
Hans-Jürg Wiesmann, ABB Switzerland, Corporate Research, Baden, Switzerland
Dietmar Winkler, Telemark University College, Porsgrunn, Norway

E.6.3 Acknowledgments

Partial financial support for the development of Modelica 3.2 by the following funding agencies has been received:

- The German Ministry BMBF has partially funded DLR, Fraunhofer and Siemens (BMBF Förderkennzeichen: 01IS07022F) within the ITEA2 project EUROSYSLIB (<https://itea3.org/project/eurosyslib.html>).
- The German Ministry BMBF has partially funded ITI GmbH (BMBF Förderkennzeichen: 01IS08002K), and the Swedish funding agency VINNOVA has partially funded Dynasim (2008-02291), within the ITEA2 project MODELISAR (<https://itea3.org/project/modelisar.html>).

- The Swedish funding agency VINNOVA has partially funded Linköping University (PELAB) within the ITEA2 project OPENPROD (<http://www.openprod.org>).
- The Swedish Research Council has partially funded Linköping University (PELAB) within the project *High-Level Debugging of Equation-Based System Modeling & Simulation Languages*.
- The German Ministry BMBF has partially funded FH Bielefeld (BMBF Förderkennzeichen: 01IS09029C) within the ITEA2 project OPENPROD (<http://www.openprod.org>).

E.7 Modelica 3.1

Modelica 3.1 was released on May 27, 2009. The Modelica 3.1 specification was edited by Francesco Casella, Rüdiger Franke, Hans Olsson, Martin Otter, and Michael Sielemann.

E.7.1 Main changes in Modelica 3.1

The following *backward compatible extensions* have been introduced with Modelica 3.1:

- Overloading of operators like '+' or '*' to allow convenient usage of user-defined data structures like complex numbers, polynomials, transfer functions. Usually, only scalar operations for one data type need to be overloaded. All other needed operations, like operations between different data types or on arrays of the new data type, can be automatically constructed by the tool.
- Stream connector concept to ensure efficient and reliable simulation of fluid systems.
- Partitioning models in parts and mapping these parts to execution environments. This allows convenient definition of, e.g., Model-in-the-Loop, Software-in-the-Loop, Hardware-in-the-Loop Simulation, from the same “logical” system, by inheriting from the logical system and setting configuration options.
- Arrays in buses (expandable connector) are much better supported. Furthermore, variables declared in an expandable connector need not to be referenced in the model and are then not available in the simulation model.
- The order of classes stored in separate files can be given.
- A restriction of balanced models was removed, so that modifiers to connector and record instances can be used and are considered for the equation count. This allows, e.g., a much easier implementation of the support connector of the Modelica.Mechanics.Rotational/Translational libraries.
- A tool that uses `missingInnerMessage` to give information may also automatically use the corresponding `inner`-component.
- URIs can be used for links in html-documentation and for the Bitmap annotation (such as: `modelica://Modelica.Mechanics/C.jpg` for image `C.jpg` that is stored in the directory of package `Modelica.Mechanics`). This allows storing resources in a package on persistent storage and to reference resources via package and resource names.
- *Annotation* `connectorSizing` to automatically enlarge a vector of connectors and connect to a free element of this vector when a connection line is drawn. This allows improving the user convenience, especially for state machine and fluid models.
- *Annotation* `inverse` to define inverses of functions. This allows a tool to solve non-linear algebraic loops by using the user-provided inverse function.
- *Annotations* `versionDate`, `versionBuild`, `dateModified`, `revisionId` to improve version handling. For example, this allows handling of maintenance (bug-fix) releases and gives a library developer the possibility to state that a particular build is needed for a used library.

The following defects have been fixed in the Modelica specification:

- Modifier with subscripts were previously allowed, but not clearly defined, not implemented in many tools, and not used in libraries. They were thus removed.
- Modelica keywords (section 2.3.3) updated.

- Clarification: Exponentiation and array range operator are non-associative (x^y^z or $a:b:c:d:e:f$ are not allowed; parentheses are required)
- Clarification: Restrictions on combining base classes (section 4.5.2).
- Clarification: Execution of an algorithm (new section 11.1.2).
- The default type for arrays sent to external functions has been clarified.
- The `iconTransformation` defaults to the (diagram) `transformation`, as was originally the intention.
- The Connection set section was rewritten:
 - Connection set is clearly defined including examples.
 - The handling of connections involving `outer` components was rewritten to ensure that models that seem to have the same connection structure generate the same equations from the connect-equations. Previously a connection between an `outer` component and an outside connector would move the connection upwards. The implication of this movement was unclear: if it turned the connector into an inside connector (as was likely the intention) it would prohibit default-connection of this normal connector leading to counter-intuitive results; and otherwise it would have no impact.
- Example of using fields was corrected.
- Example with `MatrixGain` was corrected.
- Ambiguous annotations after external-clause were corrected (as already used).
- The `reinit` operator can be used multiple times in one algorithm, and the semantics clarified.
- Made clearer that acyclic parameters also hold for one parameter equation.
- Changed in the text *attributes* to *prefix*, if a prefix is meant.

The following changes in Modelica 3.1 are *not* backwards compatible:

- A class-level annotation can only be placed before the end-statement. This is uncritical because tools can easily fix incorrect models (ignore this new rule when reading a model and use this rule when storing the model).
- New Modelica keywords `stream` and `operator`, a new built-in package `Subtask` and new-built-in operators `inStream` and `actualStream` have been introduced. In rare cases this might give name clashes in existing models.
- Modifiers on declared variables in expandable connectors are no longer allowed. This should be uncritical, because expandable connectors are usually used for signal buses where the signal is communicated to the bus from a block with a connect equation.

E.7.2 Contributors to the Modelica Language, Version 3.1

The concept of operator overloading was developed by Hans Olsson, based on work of Dag Brück, Peter Fritzson, and Martin Otter.

The streams concept was developed by Rüdiger Franke based on work from Francesco Casella and with contributions especially from Hilding Elmqvist, Sven Erik Mattson, Hans Olsson, Martin Otter and Michael Sielemann.

The concept to map models to execution environments was developed by Hilding Elmqvist, Dan Henriksson, Martin Otter, Bernhard Thiele and Sven Erik Mattson.

The following members of the Modelica Association participated at design meetings and contributed to the Modelica 3.1 specification:

Johan Akesson, Lund University and Modelon AB, Lund, Sweden

Johan Andreasson, Modelon AB, Lund, Sweden

Peter Aronsson, MathCore AB, Linköping, Sweden

Bernhard Bachmann , University of Applied Sciences, Bielefeld, Germany
Torsten Blochwitz, ITI GmbH, Dresden, Germany
David Broman, Linköping University, Linköping, Sweden
Dag Brück, Dynasim, Lund, Sweden
Francesco Casella, Politecnico di Milano, Milano, Italy
Christoph Clauß, Fraunhofer Institute for Integrated Circuits, Dresden, Germany
Karin Dietl, TU Hamburg-Harburg, Germany
Filippo Donida, Politecnico di Milano, Milano, Italy
Thomas Doumenc, Dassault Systèmes, Paris, France
Jonas Eborn, Modelon AB, Lund, Sweden
Hilding Elmqvist, Dynasim, Lund, Sweden
Rüdiger Franke, ABB Power Generation, Mannheim, Germany
Peter Fritzson, Linköping University, Sweden
Magnus Gäfvert, Modelon AB, Lund, Sweden
Manuel Gräber, TU Braunschweig, Germany
Anton Haumer, Technical Consulting & Electrical Engineering, St.Andrae-Woerden, Austria
Carsten Heinrich, Institut für Luft- und Kältetechnik, Dresden, Germany
Dan Henriksson, Dynasim, Lund, Sweden
Roland Kossel, TLK Thermo GmbH, Braunschweig, Germany
Christian Kral, arsenal research, Vienna, Austria
Gerd Kurzbach, ITI GmbH, Dresden, Germany
Kilian Link, Siemens AB, Erlangen, Germany
Sven Erik Mattsson, Dynasim, Lund, Sweden
Ramine Nikoukhah, INRIA, Paris, France
Hans Olsson, Dynasim, Lund, Sweden
Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany
Adrian Pop, Linköping University, Linköping, Sweden
Katrín Pröhl, Technical University Hamburg-Harburg, Germany
Christoph Richter, TU Braunschweig, Germany
Michael Sielemann, German Aerospace Center, Oberpfaffenhofen, Germany
Bernhard Thiele, German Aerospace Center, Oberpfaffenhofen, Germany
Michael Tiller, Ford Motor Company/Emmeskay, Dearborn/Plymouth, MI, U.S.A
Hubertus Tummescheit, Modelon AB, Lund, Sweden
Stefan Vorkoetter, Maplesoft, Waterloo, Canada
Hans-Jürg Wiesmann, ABB Switzerland, Corporate Research, Baden, Switzerland
Dietmar Winkler, TU Berlin, Germany

E.7.3 Acknowledgments

Partial financial support for the development of Modelica 3.1 by the following funding agencies has been received:

- The German Ministry BMBF has partially funded ABB, DLR, Fraunhofer and Siemens (BMBF Förderkennzeichen: 01IS07022F) within the ITEA2 project EUROSYSLIB (<https://itea3.org/project/eurosyslib.html>).
- The German Ministry BMBF has partially funded ITI GmbH (BMBF Förderkennzeichen: 01IS08002K), and the Swedish funding agency VINNOVA has partially funded Dynasim (2008-02291), within the ITEA2 project MODELISAR (<https://itea3.org/project/modelisar.html>).
- The Swedish funding agency VINNOVA has partially funded Linköping University (PELAB) within the project *Safe and Secure Modeling and Simulation*.
- The Swedish funding agency VR has partially funded Linköping University (PELAB) within the project *High-Level Debugging of Equation-Based System Modeling & Simulation Languages*.

E.8 Modelica 3.0

Modelica 3.0 was released Sept. 5, 2007. The Modelica 3.0 specification was edited by Peter Fritzson, Hans Olsson, and Martin Otter.

E.8.1 Contributors to the Modelica Language, Version 3.0

The Modelica 3.0 specification was newly structured and written by Peter Fritzson using text from the previous specification and also adding new explanatory text. This draft specification was afterwards improved by members of the Modelica Association.

The concept of *balanced models* (which is the major change of the language) was developed by Hans Olsson with contributions from Martin Otter, Hilding Elmqvist, and Sven Erik Mattsson. The original inspiration was from Mike Tiller.

This goes together with making the type interface and sub-typing cleaner and stricter (the new chapter 6). This concept fixes flaws of the language that have been pointed out by Sébastien Furic.

The graphical annotations have been redesigned and improved by Daniel Hedberg, Erik Areskog, Dag Brück, and Hilding Elmqvist with contributions from Peter Aronsson and Gerd Kurzbach.

The following members of the Modelica Association participated at design meetings and contributed to the Modelica 3.0 specification:

Peter Aronsson, MathCore AB, Linköping, Sweden

Bernhard Bachmann, University of Applied Sciences, Bielefeld, Germany

John Battah, Ford Motor Company, Dearborn, MI, U.S.A.

David Broman, Linköping University, Linköping, Sweden

Peter Bunus, Linköping University, Linköping, Sweden

Dag Brück, Dynasim, Lund, Sweden

Francesco Casella, Politecnico di Milano, Milano, Italy

Christoph Clauß, Fraunhofer Institute for Integrated Circuits, Dresden, Germany

Thomas Doumenc, Dassault Systèmes, Paris, France

Jonas Eborn, Modelon AB, Lund, Sweden

Hilding Elmqvist, Dynasim, Lund, Sweden

Rüdiger Franke, ABB Corporate Research, Ladenburg, Germany

Peter Fritzson, Linköping University, Sweden

Sebastien Furic, Imagine, Roanne, France
Anton Haumer, Technical Consulting & Electrical Engineering, St.Andrae-Woerden, Austria
Daniel Hedberg, MathCore AB, Linköping, Sweden
Carsten Heinrich, Institut für Luft- und Kältetechnik gGmbH, Dresden, Germany
Olof Johansson, Linköping University, Linköping, Sweden
Christian Kral, arsenal research, Vienna, Austria
Roland Kossel, TLK Thermo GmbH, Braunschweig, Germany
Gerd Kurzbach, ITI GmbH, Dresden, Germany
Christian Kral, arsenal research, Vienna, Austria
Kilian Link, Siemens AB, Erlangen, Germany
José Diaz Lopez, Dynasim AB, Lund, Sweden
Karin Lund, Fachhochschule Ulm, Germany
Håkan Lundvall, Linköping University, Linköping, Sweden
Ludwig Marvan, VA TECH ELIN EBG Elektronik GmbH & Co, Vienna, Austria
Sven Erik Mattsson, Dynasim, Lund, Sweden
Jakob Mauss, Qtronic GmbH, Berlin, Germany
Chuck Newman, Ford Motor Company, Dearborn, MI, U.S.A.
Kaj Nyström, Linköping University, Linköping, Sweden
Hans Olsson, Dynasim, Lund, Sweden
Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany
Markus Plainer, Arsenal Research, Vienna, Austria
Adrian Pop, Linköping University, Linköping, Sweden
Katrín Pröhl, Technical University Hamburg-Harburg, Germany
Christoph Richter, Technical University of Braunschweig, Braunschweig, Germany
Anders Sandholm, Linköping University, Linköping, Sweden
Christian Schweiger, German Aerospace Center, Oberpfaffenhofen, Germany
Michael Tiller, Ford Motor Company/Emmeskay, Dearborn, MI, U.S.A
Hubertus Tummescheit, Modelon AB, Lund, Sweden
Hans-Jürg Wiesmann, ABB Switzerland Ltd., Corporate Research, Baden, Switzerland

E.8.2 Main Changes in Modelica 3.0

Modelica 3.0 is a “clean-up” version of the Modelica language. For example, the specification is newly written to define the language in a better way, errors in the language are fixed, unclear or undefined items are more precisely described, and mild restrictions are introduced into the language in order that a Modelica translator can much earlier and more precisely detect modeling errors. Automated conversion of models to Modelica 3.0 is possible. Furthermore, a tool can potentially handle much larger models.

This Modelica version is for the first time (slightly) not backward compatible to previous versions (all previous versions have been backward compatible with exception of tiny issues as newly introduced keywords). As a result, e.g., “unsafe” models of previous Modelica versions are no longer valid. It is expected that Modelica tool vendors provide (semi-) automatic mechanisms for conversion of models and libraries.

The following main changes in Modelica 3.0 are *not* backwards compatible:

- Restrictions to connectors (see section 9.3): For each non-partial connector class the number of flow variables shall be equal to the number of variables that are neither parameter, constant, input, output, nor flow. For example, the following connector is illegal in Modelica 3:

```
connector notValid // illegal connector
  Real r1;
  Real r2;
  flow Real r3;
end notValid;
```

- In a non-partial `model` or `block`, all non-connector inputs of `model` or `block` components must have binding equations.
- A component declared with the `inner` or `outer` prefix shall not be of a class having top-level public connectors containing inputs.
- Modifiers for components shall only contain redeclarations of replaceable elements and binding equations for parameters, constants, inputs and variables having a default binding equation.
- All non-partial `model` and `block` classes must be locally balanced (see section 4.7). This means that the local number of unknowns equals the local equation size. Together with other restrictions, this leads to the strong property that a simulation model is always globally balanced (i.e., the number of unknowns is equal to the number of equations).
- Prefixes `input`, `output`, `inner`, `outer`, `flow` are not allowed in a record
- The built-in operators `abs` and `sign` do no longer generate events but are implicitly defined with a `noEvent` operator.
- The constraining clause of a replaceable class or component is changed from keyword `extends` to the new keyword `constrainedby` (since the `extends` keyword could lead to the wrong impression that the redeclared model must inherit from the constraining class, but this is not the case).
- The `isPresent` construct, which was not implemented in tools, was removed.

The following changes in the Modelica 3.0 graphical annotations are also *not* backwards compatible:

- Changed the definition of icon placement (record Transformation), so that the actual coordinates of the icon of a model instance are defined in the class where the instance is defined and no longer in the class where the icon is defined (this was a flaw in the Modelica 2.0 graphical annotations).
- Improved the definition of the rotation of a graphical primitive.
- Change `fontSize` unit from `DrawingUnit` to `pt` (since this is the usual unit for fonts).

The following main changes in Modelica 3.0 are *backwards compatible*:

- New element-wise operators: `.+`, `.-`, `.*`, `./`, `.^`.
- A third argument `AssertionLevel` to built-in function `assert(...)` in order that warnings can optionally be defined.
- New annotations section 18.1 *Vendor-Specific Annotations*:
In this section it is precisely defined how vendor-specific annotations should be marked. Any tool shall save files with all standard annotations (defined in chapter 18) and all vendor-specific annotations intact. The advantage is that a typo in non-vendor annotations can now be detected and marked as an error, whereas in previous versions this had to be ignored.
- New annotation in section 18.2 *Annotations for Documentation*:
`preferredView = info, diagram` or `text`
- New annotations section 18.3 *Annotations for Code Generation*:
`Evaluate`, `HideResult`, `Inline`, `LateInline`, `smoothOrder`
- New annotation section 18.4.1 *Annotations for Simulation Experiments*:
`StartTime`, `StopTime`, `Tolerance` to define important parameters of an experiment setup.

- New annotations for graphical annotations in section 18.6:
 New attribute `Smooth = enumeration(None, Bezier)` for graphical objects and connection lines (Bezier defines a Bezier spline).
 New attribute `visible` in `record Placement` allows making a graphical annotation invisible (e.g. after inheritance).
 New attributes `startAngle`, `endAngle` in record ellipse to define part of an ellipse.
 New layer specific annotations `IconMap` and `DiagramMap` for `extends`.
 New attribute `horizontalAlignment` to the `Text` record to define the horizontal alignment of text.
- New annotations for schematic animation and interactive user input in section 18.6:
`DynamicSelect` to modify annotation literals by the actual values of variables.
`OnMouseDownSetBoolean`, `OnMouseUpSetBoolean`, `OnMouseMoveXSetReal`, `OnMouseMoveYSetReal`, `OnMouseDownEditReal`, `OnMouseDownEditString` to interactively set the variable of a class during simulation.

The following errors have been fixed in the Modelica specification:

- Syntax rule for a function call (e.g., according to the grammar in Modelica 2.2 a function call of the form `Modelica.Math.Matrices.eig(...)` was invalid because the function name could not have `'.'`. However, all Modelica tools supported the desired “full Modelica name” also for function calls).

E.9 Modelica 2.2

Modelica 2.2 was released February 2, 2005. The Modelica 2.2 specification was edited by Hans Olsson, Michael Tiller and Martin Otter.

E.9.1 Contributors to the Modelica Language, Version 2.2

Bernhard Bachmann , University of Applied Sciences, Bielefeld, Germany

John Batteh, Ford Motor Company, Dearborn, MI, U.S.A.

Dag Brück, Dynasim, Lund, Sweden

Francesco Casella, Politecnico di Milano, Milano, Italy

Christoph Clauß, Fraunhofer Institute for Integrated Circuits, Dresden, Germany

Jonas Eborn, Modelon AB, Lund, Sweden

Hilding Elmqvist, Dynasim, Lund, Sweden

Rüdiger Franke, ABB Corporate Research, Ladenburg, Germany

Peter Fritzson, Linköping University, Sweden

Anton Haumer, Technical Consulting & Electrical Engineering, St.Andrae-Woerdern, Austria

Christian Kral, arsenal research, Vienna, Austria

Sven Erik Mattsson, Dynasim, Lund, Sweden

Chuck Newman, Ford Motor Company, Dearborn, MI, U.S.A.

Hans Olsson, Dynasim, Lund, Sweden

Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany

Markus Plainer, Arsenal Research, Vienna, Austria

Adrian Pop, Linköping University, Sweden

Katrin Pröll, Technical University Hamburg-Harburg, Germany

André Schneider, Fraunhofer Institute for Integrated Circuits, Dresden, Germany

Christian Schweiger, German Aerospace Center, Oberpfaffenhofen, Germany

Michael Tiller, Ford Motor Company, Dearborn, MI, U.S.A.

Hubertus Tummescheit, Modelon AB, Lund, Sweden

Hans-Jürg Wiesmann, ABB Switzerland Ltd., Corporate Research, Baden, Switzerland

E.9.2 Main Changes in Modelica 2.2

The main changes in Modelica 2.2 are:

- Conditional component declarations to ignore component declarations depending on a parameter expression. Connection equations that reference a component that is no longer present, are ignored.
- In redeclarations some parts of the original declaration are automatically inherited by the new declaration. This is intended to make it easier to write declarations by not having to repeat common parts of the declarations, and does in particular apply to attributes that must be identical.
- Recursive **inner/outer** definitions to define hierarchically structured **inner/outer** declarations that can communicate with each other: An element declared with both the prefixes **inner** and **outer** conceptually introduces two declarations with the same name, one that follows the rules for **inner** and another that follows the rules for **outer**.
- A non-input array component declared in a function with a dimension size specified by colon(:) and no declaration assignment, can change size in the function in a simple and convenient way.
- A new type of connector, called **expandable connector** was introduced. This connector has less strict requirements about name matching of connected connectors and can be used conveniently in situations that required replaceable connectors previously. One main application area is to construct signal buses of complex systems.
- The derivative operator **der(expr)** may have an expression as argument and not only a variable name as previously, e.g., **der(m*h)** is interpreted as **der(m)*h + m*der(h)**.
- A function can be defined as partial derivative of another function, e.g.:
- **"function Gibbs_T = der(Gibbs,T)"** is a function that computes the partial derivative of function Gibbs with respect to its input argument T.
- External functions may have the new attribute **"builtin"**, additionally to **"C"** or **"FORTRAN 77"**. The **"builtin"** specification is only used for functions that are defined to be built-in in the Modelica language. The external function call mechanism for **"builtin"** functions is implementation-defined.

The language changes are backward compatible.

E.10 Modelica 2.1

Modelica 2.1 was released January 30, 2004. The Modelica 2.1 specification was edited by Hans Olsson and Martin Otter.

E.10.1 Contributors to the Modelica Language, Version 2.1

Mikael Adlers, MathCore, Linköping, Sweden

Peter Aronsson, Linköping University, Sweden

Bernhard Bachmann, University of Applied Sciences, Bielefeld, Germany

Peter Bunus, Linköping University, Sweden

Jonas Eborn, United Technologies Research Center, Hartford, U.S.A.

Hilding Elmqvist, Dynasim, Lund, Sweden

Rüdiger Franke, ABB Corporate Research, Ladenburg, Germany

Peter Fritzson, Linköping University, Sweden

Anton Haumer, Technical Consulting & Electrical Engineering, St.Andrae-Woerden, Austria

Olof Johansson, Linköping University, Sweden

Karin Lunde, R.O.S.E. Informatik GmbH, Heidenheim, Germany

Sven Erik Mattsson, Dynasim, Lund, Sweden

Hans Olsson, Dynasim, Lund, Sweden

Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany

Levon Saldamli, Linköping University, Sweden

Christian Schweiger, German Aerospace Center, Oberpfaffenhofen, Germany

Michael Tiller, Ford Motor Company, Dearborn, MI, U.S.A.

Hubertus Tummescheit, United Technologies Research Center, Hartford, U.S.A.

Hans-Jürg Wiesmann, ABB Switzerland Ltd., Corporate Research, Baden, Switzerland

E.10.2 Main Changes in Modelica 2.1

The main changes in Modelica 2.1 are:

- Arrays and array indices of Enumerations (needed, e.g., in the `Electrical.Digital` library currently under development).
- Connections into hierarchical connectors (needed, e.g., for convenient implementation of buses).
- Optional output arguments of Modelica functions. The presence of actual input and/or output arguments can be inquired with the new built-in function `isPresent`. The previous built-in function and attribute `enable` was removed.
- Making the default constraining type more useful by inheriting the base constraining type automatically to modifications.
- Enhanced redeclaration as needed, e.g., in the Modelica.Media library under development (e.g. `redeclare model name` or `model extends name (<modifications>)`).
- Handling of overdetermined connectors (needed, e.g., for multi-body systems and electrical power systems) including the new built-in package `Connections` with operators `Connections.branch`, `Connections.root`, `Connections.potentialRoot`, `Connections.isRoot`.
- Statement `break` in the while-loop of an algorithm section.
- Statement `return` in a Modelica function.
- Built-in function `String` to provide a string representation of `Boolean`, `Integer`, `Real` and enumeration types.
- Built-in function `Integer` to provide the `Integer` representation of an enumeration type.
- Built-in function `semiLinear` to define a characteristics with two slopes and a set of rules for symbolic transformations, especially when the function becomes underdetermined (this function is used in the Modelica Fluid library under development to define reversing flow in a mathematically clean way).
- More general identifiers by having any character in single quotes, e.g. `'+'` or `'123.456#1'` are valid identifiers. `'x'` and `x` are different identifiers. This is useful for a direct mapping of product identifiers to model names and for having the usual symbols for digital electrical signals as enumerations (such as `'+'`, `'-'`, `'0'`, `'1'`).
- New annotations:
 - For version handling of libraries and models (`version`, `uses`, `conversion`),
 - for revision logging (`revisions`),
 - for using a Modelica name as link in a HTML documentation text,

- for convenient `inner` declaration in a GUI (`defaultComponentName`, `defaultComponentPrefixes`),
- for parameter menu structuring (`Dialog`, `enable`, `tab`, `group`), and
- for library specific error messages (`missingInnerMessage`, `unassignedMessage`).

Fixing some minor errors in the grammar and semantic specification.

The language changes are backward compatible, except for the introduction of the new keywords `break` and `return`, the new built-in package `Connections` and the removing of built-in function and attribute `enable`.

E.11 Modelica 2.0

Modelica 2.0 was released January, 30 2002, and the draft was released on December 18 in 2001. The Modelica 2.0 specification was edited by Hans Olsson. Modelica is a registered trademark owned by the Modelica Association since November 2001.

E.11.1 Contributors to the Modelica Language, Version 2.0

Peter Aronsson, Linköping University, Sweden

Bernhard Bachmann , University of Applied Sciences, Bielefeld

Peter Beater, University of Paderborn, Germany

Dag Brück, Dynasim, Lund, Sweden

Peter Bunus, Linköping University, Sweden

Hilding Elmqvist, Dynasim, Lund, Sweden

Vadim Engelson, Linköping University, Sweden

Peter Fritzson, Linköping University, Sweden

Rüdiger Franke, ABB Corporate Research, Ladenburg

Pavel Grozman, Equa, Stockholm, Sweden

Johan Gunnarsson, MathCore, Linköping

Mats Jirstrand, MathCore, Linköping

Sven Erik Mattsson, Dynasim, Lund, Sweden

Hans Olsson, Dynasim, Lund, Sweden

Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany

Levon Saldamli, Linköping University, Sweden

Michael Tiller, Ford Motor Company, Dearborn, MI, U.S.A.

Hubertus Tummescheit, Lund Institute of Technology, Sweden

Hans-Jürg Wiesmann, ABB Switzerland Ltd.,Corporate Research, Baden, Switzerland

E.11.2 Main Changes in Modelica 2.0

A detailed description of the enhancements introduced by Modelica 2.0 is given in the papers

- M. Otter, H. Olsson: New Features in Modelica 2.0. 2nd International Modelica Conference, March 18-19, DLR Oberpfaffenhofen, Proceedings, pp. 7.1–7.12, 2002. This paper can be downloaded from https://www.modelica.org/events/Conference2002/papers/p01_Otter.pdf

- Mattsson S. E., Elmqvist H., Otter M., and Olsson H.: Initialization of Hybrid Differential-Algebraic Equations in Modelica 2.0. 2nd International Modelica Conference, March 18-19, DLR Oberpfaffenhofen, Proceedings, pp. 9–15, 2002. This paper can be downloaded from https://www.modelica.org/events/Conference2002/papers/p02_Mattsson.pdf

The main changes in Modelica 2.0 are:

- Full specification of initialization in order to compute consistent initial values of all variables appearing in a model before performing an operation, such as simulation or linearization.
- Specified the graphical appearance of Modelica object diagrams, thereby ensuring portability of model topology information and improving the previous informal graphical description, e.g., with separate icon and diagram positions.
- Enumeration types to allow the definition of options and properties in an understandable, safe and efficient way.
- Support for (optional) explicit preference in state-selection in order that a modeler can incorporate application specific knowledge to guide the solution process, e.g., for real-time simulation.
- Iterators in array constructors and reduction operators, to support more powerful expressions, especially in declarations, in order to avoid inconvenient and less efficient local function definitions.
- Support for generic formulation of blocks applicable to both scalar and vector connectors, connection of (automatically) vectorized blocks, and simpler input/output connectors. This allows significant simplifications of the input/output block library of Modelica, e.g., since only scalar versions of all blocks have to be provided. Furthermore, new library components can be incorporated more easily.
- Record constructor to allow, e.g., the construction of data sheet libraries.
- Functions with mixed positional and named arguments. Optional results and default arguments make the same function fit for beginners and expert users.
- Additional utilities for external C functions that are interfaced to Modelica models, especially supporting external functions returning strings and external functions with internal memory (e.g., to interface user-defined tables, property databases, sparse matrix handling, hardware interfaces).
- Added an index, and specification of some basic constructs that had previously not formally been defined, such as while-clauses, if-clauses.

The language changes are backward compatible, except for the introduction of the new keyword enumeration and the removal of the operator `analysisType`. The library change of the block library which will become available soon requires changes in user-models.

E.12 Modelica 1.4

Modelica 1.4 was released December 15, 2000. The Modelica Association was formed in Feb. 5, 2000 and is now responsible for the design of the Modelica language. The Modelica 1.4 specification was edited by Hans Olsson and Dag Brück.

E.12.1 Contributors to the Modelica Language, Version 1.4

Bernhard Bachmann, Fachhochschule Bielefeld, Germany

Peter Bunus, MathCore, Linköping, Sweden

Dag Brück, Dynasim, Lund, Sweden

Hilding Elmqvist, Dynasim, Lund, Sweden

Vadim Engelson, Linköping University, Sweden

Jorge Ferreira, University of Aveiro, Portugal

Peter Fritzson, Linköping University, Linköping, Sweden

Pavel Grozman, Equa, Stockholm, Sweden

Johan Gunnarsson, MathCore, Linköping, Sweden

Mats Jirstrand, MathCore, Linköping, Sweden

Clemens Klein-Robbenhaar, Germany

Pontus Lidman, MathCore, Linköping, Sweden

Sven Erik Mattsson, Dynasim, Lund, Sweden

Hans Olsson, Dynasim, Lund, Sweden

Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany

Tommy Persson, Linköping University, Sweden

Levon Saldamli, Linköping University, Sweden

André Schneider, Fraunhofer Institute for Integrated Circuits, Dresden, Germany

Michael Tiller, Ford Motor Company, Dearborn, MI, U.S.A.

Hubertus Tummescheit, Lund Institute of Technology, Sweden

Hans-Jürg Wiesmann, ABB Corporate Research Ltd., Baden, Switzerland

E.12.2 Contributors to the Modelica Standard Library

Peter Beater, University of Paderborn, Germany

Christoph Clauß, Fraunhofer Institute for Integrated Circuits, Dresden, Germany

Martin Otter, German Aerospace Center, Oberpfaffenhofen, Germany

André Schneider, Fraunhofer Institute for Integrated Circuits, Dresden, Germany

Hubertus Tummescheit, Lund Institute of Technology, Sweden

E.12.3 Main Changes in Modelica 1.4

- Removed declare-before-use rule. This simplifies graphical user environments, because there exists no order of declarations when components are graphically composed together.
- Refined package concept by introducing encapsulated classes and import mechanism. Encapsulated classes can be seen as “self-contained units”: When copying or moving an encapsulated class, at most the import statements in this class have to be changed.
- Refined when-clause: The `nondiscrete` keyword is removed, equations in when-clauses must have a unique variable name on left hand side variable and the exact mapping of when-clauses to equations is defined. As a result, when-clauses are now precisely defined without referring to a sorting algorithm and it is possible to handle algebraic loops between when-clauses with different conditions and between when-clauses and the continuous-time part of a model. The `discrete` keyword is now optional, simplifying the library development because only one type of connector is needed and not several types which do contain or do not contain the discrete prefix on variables. Additionally, when-clauses in algorithm sections may have `elsewhen`-clauses which simplifies the definition of priorities between when-clauses.
- For replaceable declarations: allowed constraining clauses, and annotations listing suitable redeclarations. This allows a graphical user environment to automatically build menus with meaningful choices.
- Functions can specify their derivative. This allows, e.g., the application of the Pantelides algorithm to reduce the index of a DAE also for external functions.
- New built-in operator `rem` (remainder) and the built-in operators `div`, `mod`, `ceil`, `floor`, `integer`, previously only allowed to be used in when-clauses can now be used everywhere, because state events are automatically generated when the result value of one of these operator changes discontinuously.

- Quantity attribute also for base types `Boolean`, `Integer`, `String` (and not only for `Real`), in order to allow abstracted variables to refer to physical quantities (e.g. `Boolean i(quantity="Current")` is true if current is flowing and is false if no current is flowing).
- `final` keyword also allowed in declaration, to prevent modification. Example:

```

model A
  Real x[:];
  final Integer n=size(x,1);
end A;

```

- Several minor enhancements, such as usage of dot-notation in modifications (e.g.: `A x(B.C=1, B.D=2)` is the same as `A x(B(C=1, D=2))`).
- Internally restructured specification.

Modelica 1.4 is backwards compatible with Modelica 1.3, with the exception of (1) some exotic cases where different results are achieved with the removed “declare-before-use-rule” and the previous declaration order, (2) when-clauses in equations sections, which use the general form `expr1 = expr2` (now only `v = expr` is allowed + some special cases for functions), (3) some exotic cases where a when-clause may be no longer evaluated at the initial time, because the initialization of the when-condition is now defined in a more meaningful way (before Modelica 1.4, every condition in a when-clause has a “previous” value of false), and (4) models containing the `nondiscrete` keyword which was removed.

E.13 Modelica 1.3 and Older Versions.

Modelica 1.3 was released December 15, 1999.

E.13.1 Contributors up to Modelica 1.3

The following list contributors and their affiliations at the time when Modelica 1.3 was released.

Hilding Elmqvist, Dynasim AB, Lund, Sweden

Bernhard Bachmann, ABB Corporate Research Center Heidelberg

Francois Boudaud, Gaz de France, Paris, France

Jan Broenink, University of Twente, Enschede, Netherlands

Dag Brück, Dynasim AB, Lund, Sweden

Thilo Ernst, GMD FIRST, Berlin, Germany

Ruediger Franke, ABB Network Partner Ltd. Baden, Switzerland

Peter Fritzson, Linköping University, Sweden

Alexandre Jeandel, Gaz de France, Paris, France

Pavel Grozman, Bris Data AB, Stockholm, Sweden

Kaj Juslin, VTT, Espoo, Finland

David Kågedal, Linköping University, Sweden

Mattias Klose, Technical University of Berlin, Germany

Nathalie Loubere, Gaz de France, Paris, France

Sven-Erik Mattsson, Dynasim AB, Lund, Sweden

Peter J. Mosterman, DLR Oberpfaffenhofen, Germany

Henrik Nilsson, Linköping University, Sweden

Hans Olsson, , Dynasim AB, Lund, Sweden

Martin Otter, DLR Oberpfaffenhofen, Germany

Per Sahlin, Bris Data AB, Stockholm, Sweden

Andrée Schneider, Fraunhofer Institute for Integrated Circuits, Dresden, Germany

Michael Tiller, Ford Motor Company, Dearborn, MI, U.S.A.

Hubertus Tummescheit, Lund Institute of Technology, Sweden

Hans Vangheluwe, University of Gent, Belgium

E.13.2 Main Changes in Modelica 1.3

Modelica 1.3 was released December 15, 1999.

- Defined connection semantics for `inner/outer` connectors.
- Defined semantics for protected element.
- Defined that least variable variability prefix wins.
- Improved semantic definition of array expressions.
- Defined scope of for-loop variables.

E.13.3 Main Changes in Modelica 1.2

Modelica 1.2 was released June 15, 1999.

- Changed the external function interface to give greater flexibility.
- Introduced `inner/outer` for dynamic types.
- Redefined `final` keyword to only restrict further modification.
- Restricted redeclaration to replaceable elements.
- Defined semantics for if-clauses.
- Defined allowed code optimizations.
- Refined the semantics of event-handling.
- Introduced `fixed` and `nominal` attributes.
- Introduced `terminate` and `analysisType`.

E.13.4 Main Changes in Modelica 1.1

Modelica 1.1 was released in December 1998.

Major changes:

- Specification as a separate document from the rationale.
- Introduced prefixes `discrete` and `nondiscrete`.
- Introduced `pre` and `when`.
- Defined semantics for array expressions.
- Introduced built-in functions and operators (only `connect` was present in Modelica 1.0).

E.13.5 Modelica 1.0

Modelica 1, the first version of Modelica, was released in September 1997, and had the language specification as a short appendix to the rationale.

Bibliography

- Benveniste, Albert, Paul Caspi, Stephen A. Edwards, Nicolas Halbwachs, Paul Le Guernic, and Robert de Simone (2003). “The Synchronous Languages Twelve Years Later”. In: *Proceedings of the IEEE* 91.1. URL: <https://doi.org/10.1109/JPROC.2002.805826> (cit. on p. 207).
- Colaço, Jean-Louis and Marc Pouzet (Oct. 2003). “Clocks as First Class Abstract Types”. In: *Third International Workshop on Embedded Software, EMSOFT 2003*. Philadelphia, Pennsylvania, USA. URL: <http://www.di.ens.fr/~pouzet/lucid-synchrone/papers/emsoft03.ps.gz> (cit. on p. 207).
- Elmqvist, Hilding, Martin Otter, and Francois E. Cellier (June 1995). “Inline Integration: A New Mixed Symbolic/Numeric Approach for Solving Differential-Algebraic Equation Systems”. In: *Proceedings of ESM’95, European Simulation Multiconference*. Prague, Czech Republic, pp. xxiii–xxxiv. URL: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.127.3787%5C&rep=rep1%5C&type=pdf> (cit. on p. 224).
- Forget, Julien, Frédéric Boniol, David Lesens, and Claire Pagetti (Dec. 2008). “A Multi-Periodic Synchronous Data-Flow Language”. In: *11th IEEE High Assurance Systems Engineering Symposium (HASE’08)*. Nanjing, China, pp. 251–260. URL: <https://doi.org/10.1109/HASE.2008.47> (cit. on p. 207).
- Harel, David (1987). “Statecharts: A Visual Formalism for Complex Systems”. In: *Science of Computer Programming* 8, pp. 231–274. URL: http://www.inf.ed.ac.uk/teaching/courses/seoc1/2005_2006/resources/statecharts.pdf (cit. on p. 229).
- Pantelides, Constantinos C. (Mar. 1988). “The Consistent Initialization of Differential-Algebraic Systems”. In: *SIAM Journal on Scientific and Statistical Computing* 9.2, pp. 213–231. URL: <http://epubs.siam.org/doi/abs/10.1137/0909014> (cit. on p. 259).
- Pouzet, Marc (2006). *Lucid Synchrone, Version 3.0, Tutorial and Reference Manual*. URL: <http://www.di.ens.fr/~pouzet/lucid-synchrone/> (cit. on pp. 207, 229).
- Thümmel, Michael, Gertjan Looye, Matthias Kurze, Martin Otter, and Johann Bals (Mar. 2005). “Non-linear Inverse Models for Control”. In: *Proceedings of 4th International Modelica Conference, ed. G. Schmitz*. Hamburg, Germany. URL: https://www.modelica.org/events/Conference2005/online_proceedings/Session3/Session3c3.pdf (cit. on pp. 208, 221).