
Systems Biology Markup Language (SBML) Level 2: Structures and Facilities for Model Definitions

Michael Hucka
Andrew M. Finney
Stefan Hoops
Sarah M. Keating
Nicolas Le Novère

California Institute of Technology, USA
Oxford Computer Consultants Ltd, UK
Virginia Bioinformatics Institute, USA
University of Hertfordshire, UK
EMBL European Bioinformatics Institute, UK

sbml-editors@sbml.org

SBML Level 2 Version 3

Release 1

16 June 2007

Corrections and other changes to this SBML language specification may appear over time.
Notifications of new releases are broadcast on the mailing list sbml-announce@sbml.org

The latest release of the SBML Level 2 Version 3 specification is available at
<http://sbml.org/specifications/sbml-level-2/version-3/>

This release of the specification is available at
<http://sbml.org/specifications/sbml-level-2/version-3/release-1/>

The list of known issues in all releases of SBML Level 2 Version 3 is available at
<http://sbml.org/specifications/sbml-level-2/version-3/errata/>

The XML Schema for SBML Level 2 Version 3 is available at
<http://sbml.org/xml-schemas/>

Contents

| | | |
|------------------------|--|------------|
| 1 | Introduction | 3 |
| 1.1 | Developments, discussions, and notifications of updates | 3 |
| 1.2 | SBML Levels, Versions, and Releases | 3 |
| 1.3 | Language features and backward compatibility | 4 |
| 1.4 | Document conventions | 5 |
| 2 | Overview of SBML | 9 |
| 3 | Preliminary definitions and principles | 11 |
| 3.1 | Primitive data types | 11 |
| 3.2 | Type <i>SBase</i> | 13 |
| 3.3 | The <i>id</i> and <i>name</i> attributes on SBML components | 18 |
| 3.4 | Mathematical formulas in SBML Level 2 | 20 |
| 4 | SBML components | 30 |
| 4.1 | The SBML container | 30 |
| 4.2 | Model | 31 |
| 4.3 | Function definitions | 33 |
| 4.4 | Unit definitions | 35 |
| 4.5 | Compartment types | 42 |
| 4.6 | Species types | 43 |
| 4.7 | Compartments | 44 |
| 4.8 | Species | 47 |
| 4.9 | Parameters | 51 |
| 4.10 | Initial assignments | 53 |
| 4.11 | Rules | 56 |
| 4.12 | Constraints | 61 |
| 4.13 | Reactions | 63 |
| 4.14 | Events | 77 |
| 5 | The Systems Biology Ontology and the <i>sboTerm</i> attribute | 83 |
| 5.1 | Principles | 83 |
| 5.2 | Using SBO and <i>sboTerm</i> | 84 |
| 5.3 | Relationships to the SBML annotation element | 89 |
| 5.4 | Discussion | 89 |
| 6 | A standard format for the annotation element | 91 |
| 6.1 | Motivation | 91 |
| 6.2 | XML namespaces in the standard annotation | 91 |
| 6.3 | General syntax for the standard annotation | 92 |
| 6.4 | Use of URIs | 93 |
| 6.5 | Relation elements | 94 |
| 6.6 | Model history | 94 |
| 6.7 | Examples | 96 |
| 7 | Example models expressed in XML using SBML | 102 |
| 7.1 | A simple example application of SBML | 102 |
| 7.2 | Example involving units | 103 |
| 7.3 | Example of a discrete version of a simple dimerization reaction | 105 |
| 7.4 | Example involving assignment rules | 107 |
| 7.5 | Example involving algebraic rules | 109 |
| 7.6 | Example with combinations of <i>boundaryCondition</i> and constant values on <i>Species</i> with <i>RateRule</i> objects | 111 |
| 7.7 | Example of translation from a multi-compartmental model to ODEs | 112 |
| 7.8 | Example involving function definitions | 115 |
| 7.9 | Example involving <i>delay</i> functions | 116 |
| 7.10 | Example involving events | 117 |
| 7.11 | Example involving two-dimensional compartments | 119 |
| 8 | Discussion | 124 |
| 8.1 | Future enhancements: SBML Level 3 and beyond | 124 |
| Acknowledgments | | 126 |
| A | XML Schema for SBML | 127 |
| B | XML Schema for MathML subset | 136 |
| C | Validation rules for SBML | 140 |
| D | A method for assessing whether an SBML model is overdetermined | 153 |
| E | Mathematical consequences of the <i>fast</i> attribute on Reaction | 156 |
| F | Processing and validating <i>SBase</i> notes and Constraint message content | 158 |
| G | Major changes between versions of SBML Level 2 and implications for backward compatibility | 159 |
| G.1 | Between Version 2 and Version 1 | 159 |
| G.2 | Between Version 3 and Version 2 | 160 |
| References | | 162 |

1 Introduction

We present the Systems Biology Markup Language (SBML) Level 2 Version 3 Release 1, a model representation format for systems biology. SBML is oriented towards describing systems of biochemical reactions of the sort common in research on a number of topics, including cell signaling pathways, metabolic pathways, biochemical reactions, gene regulation, and many others. SBML is defined in a neutral fashion with respect to programming languages and software encoding; however, it is primarily oriented towards allowing models to be encoded using XML, the eXtensible Markup Language (Bosak and Bray, 1999; Bray et al., 2000). This document contains many examples of SBML models written in XML, as well as the text of an XML Schema (Biron and Malhotra, 2000; Fallside, 2000; Thompson et al., 2000) that defines SBML Level 2 Version 3. A downloadable copy of the XML Schema and other related documents and software are also available from the SBML project web site, <http://sbml.org/>.

The SBML project is not an attempt to define a universal language for representing quantitative models. The rapidly evolving views of biological function, coupled with the vigorous rates at which new computational techniques and individual tools are being developed today, are incompatible with a one-size-fits-all idea of a universal language. A more realistic alternative is to acknowledge the diversity of approaches and methods being explored by different software tool developers, and seek a common intermediate format—a *lingua franca*—enabling communication of the most essential aspects of the models.

The definition of the model description language presented here does not specify *how* programs should communicate or read/write SBML. We assume that for a simulation program to communicate a model encoded in SBML, the program will have to translate its internal data structures to and from SBML, use a suitable transmission medium and protocol, etc., but these issues are outside of the scope of this document.

1.1 Developments, discussions, and notifications of updates

SBML has been, and continues to be, developed in collaboration with an international community of researchers and software developers. As in many projects, the primary mode of interaction between members is electronic mail. Discussions about SBML take place on the mailing list sbml-discuss@caltech.edu. The mailing list archives and a web browser-based interface to the list are available at <http://sbml.org/forums/>.

A low-volume, broadcast-only mailing list is available where notifications of revisions to the SBML specification, notices of votes on SBML technical issues, and other critical matters are announced. This list is sbml-announce@caltech.edu and anyone may subscribe to it freely. This list will never be used for advertising and its membership list will never be disclosed. *It is vitally important that all users of SBML stay informed about new releases and other developments by subscribing to this list*, even if they do not wish to participate in discussions on the sbml-discuss@caltech.edu list. Please visit the SBML project web site, <http://sbml.org/>, for information about how to subscribe to sbml-announce@caltech.edu as well as for access to the list archives.

In Section 8.1, we attempt to acknowledge as many contributors to SBML's development as we can, but as SBML evolves, it becomes increasingly difficult to detail the individual contributions on a project that has truly become an international community effort.

1.2 SBML Levels, Versions, and Releases

Major editions of SBML are termed *levels* and represent substantial changes to the composition and structure of the language. The edition of SBML defined in this document, SBML Level 2, represents an evolution of the language resulting from the practical experiences of many users and developers working with SBML Level 1 since its introduction in the year 2001 (Hucka et al., 2001, 2003). All of the constructs of Level 1 can be mapped to Level 2. In addition, a subset of Level 2 constructs can be mapped to Level 1. However, the levels remain distinct; a valid SBML Level 1 document is not a valid SBML Level 2 document, and likewise, a valid SBML Level 2 document is not a valid SBML Level 1 document.

Minor revisions of SBML are termed *versions* and constitute changes within a Level to correct, adjust, and refine language features. The present document defines SBML Level 2 Version 3. In Section 1.4.1 below

1 explains how color is used in this document to indicate changes; a separate document provides a detailed
2 listing of the changes between versions of SBML Level 2 as well as between SBML Level 2 Version 3 and
3 SBML Level 1 Version 2.

4 Specification documents inevitably require minor editorial changes as its users discover errors and ambi-
5 guities. As a practical reality, these discoveries occur over time. In the context of SBML, such problems
6 are formally announced publicly as *errata* in a given specification document. Borrowing concepts from the
7 World Wide Web Consortium (Jacobs, 2004), we define SBML errata as changes of the following types:
8 (a) formatting changes that do not result in changes to textual content; (b) corrections that do not affect
9 conformance of software implementing support for a given combination of SBML Level and Version; and (c)
10 corrections that *may* affect such software conformance, but add no new language features. A change that
11 affects conformance is one that either turns conforming data, processors, or other conforming software into
12 non-conforming software, or turns non-conforming software into conforming software, or clears up an ambi-
13 guity or insufficiently documented part of the specification in such a way that software whose conformance
14 was once unclear now becomes clearly conforming or non-conforming (Jacobs, 2004). In short, errata do not
15 change the fundamental semantics or syntax of SBML; they clarify and disambiguate the specification and
16 correct errors. (New syntax and semantics are only introduced in SBML Versions and Levels.) An electronic
17 tracking system for reporting and monitoring such issues is available at <http://sbml.org/issue-tracker>.

18 SBML errata result in new *Releases* of the SBML specification. Each release is numbered with an integer,
19 with the first release of the specification being called release number 1. Subsequent releases of an SBML
20 specification document contain a section listing the accumulated errata reported and corrected since the
21 first release. A complete list of the errata for SBML Level 2 Version 3 since the publication of Release 1
22 is also made publicly available at <http://sbml.org/specifications/sbml-level-2/version-3/errata/>.
23 Announcements of errata, releases of the SBML specification and other major changes are made on the
24 sbml-announce@caltech.edu mailing list.

25 1.3 Language features and backward compatibility

26 Some language features of previous SBML Levels and Versions have been either deprecated or removed
27 entirely in SBML Level 2 Version 3. For the purposes of SBML specifications, the following are the definitions
28 of *deprecated feature* and *removed feature*:

29 *Removed language feature*: A syntactic construct that was present in previous SBML Levels and/or Versions
30 within a Level, and has been removed beginning with a specific SBML Level and Version. Models
31 containing such constructs do not conform to the specification of that SBML Level and Version.

32 *Deprecated language feature*: A syntactic construct that was present in previous SBML Levels and/or
33 Versions within a Level, and while still present in the language definition, has been identified as non-
34 essential and planned for future removal. Beginning with the Level and Version in which a given
35 feature is deprecated, software tools should not generate SBML models containing the deprecated
36 feature; however, for backward compatibility, software tools reading SBML should support the feature
37 until it is actually removed.

38 As a matter of SBML design philosophy, the preferred approach to removing features is by deprecating them
39 if possible. Immediate removal of SBML features is not done unless serious problems have been discovered
40 involving those features, and keeping them would create logical inconsistencies or extremely difficult-to-
41 resolve problems. The deprecation or outright removal of features in a language, whether SBML or other,
42 can have significant impact on backwards compatibility. Such changes are also inevitable over the course
43 of a language's evolution. SBML must by necessity continue evolving through the experiences of its users
44 and implementors. Eventually, some features will be deemed unhelpful despite the best intentions of the
45 language editors to design a timeless language.

46 Throughout the SBML specification, removed and deprecated features are discussed in the text of the sections
47 where the features previously appeared. Appendix G lists the changes and describes their motivations in
48 more detail.

1.4 Document conventions

In this section, we describe the conventions we use in this specification document in an effort to communicate information more effectively and consistently.

1.4.1 Color conventions

Throughout this document, we use coloring to carry additional information for the benefit of those viewing the document on media that can display color:

- We use red color in text and figures to indicate changes between this version of the specification (Level 2 Version 3 Release 1) and the *most recent previous version* of the specification (which, for the present case, is SBML Level 2 Version 2 Release 1). The changes may be either additions or deletions of text; in the case of deletions, entire sentences, paragraphs or sections are colored to indicate a change has occurred inside them.
- We use blue color in text to indicate a hyperlink from one point in this document to another. Clicking your computer's pointing device on blue-colored text will cause a jump to the section, figure, table or page to which the link refers. (Of course, this capability is only available when using electronic formats that support hyperlinking, such as PDF and HTML.)

1.4.2 Typographical conventions for names

The following typographical notations are used in this document to distinguish objects and data types from other kinds of entities:

AbstractClass: Abstract classes are classes that are never instantiated directly, but rather serve as parents of other classes. Their names begin with a capital letter and they are printed in a slanted, bold, sans-serif typeface. In electronic document formats, the class names are also hyperlinked to their definitions in the specification. For example, in the PDF and HTML versions of this document, clicking on the word **Sbase** will send the reader to the section containing the definition of this class.

Class: Names of ordinary (concrete) classes begin with a capital letter and are printed in an upright, bold, sans-serif typeface. In electronic document formats, the class names are also hyperlinked to their definitions in the specification. For example, in the PDF and HTML versions of this document, clicking on the word **Species** will send the reader to the section containing the definition of this class.

Something, otherThing: **Attributes** of classes, data type names, literal XML, and generally all tokens *other* than SBML UML class names, are printed in an upright typewriter typeface. Primitive types defined by SBML begin with a capital letter, but unfortunately, XML Schema 1.0 does not follow any convention and primitive XML types may either start with a capital letter (e.g., **ID**) or not (e.g., **double**).

1.4.3 UML notation

Previous specifications of SBML used a notation that was at one time (in the days of SBML Level 1) fairly close to UML, the Unified Modeling Language (Eriksson and Penker, 1998; Oestereich, 1999), though many details were omitted from the UML diagrams themselves. Over the years, the notation used in successive specifications of SBML grew increasingly less UML-like. Beginning with SBML Level 2 Version 3, we have completely overhauled the specification's use of UML and once again define the XML syntax of SBML using, as much as possible, proper and complete UML 1.0. We then systematically map this UML notation to XML, using XML Schema 1.0 (Biron and Malhotra, 2000; Fallside, 2000; Thompson et al., 2000) to express the overall syntax of SBML. In the rest of this section, we summarize the UML notation used in this document and explain the few embellishments needed to support transformation to XML form. A complete Schema for SBML is given in Appendix A.

We see three main advantages to using UML as a basis for defining SBML data objects. First, compared to using other notations or a programming language, the UML visual representations are generally easier to grasp by readers who are not computer scientists. Second, the notation is implementation-neutral: the

objects can be encoded in any concrete implementation language—not just XML, but C, Java and other languages as well. Third, UML is a de facto industry standard that is documented in many resources. Readers are therefore more likely to be familiar with it than other notations.

Object class definitions

Object classes in UML diagrams are drawn as simple tripartite boxes, as shown in Figure 1 (left). UML allows for operations as well as data attributes to be defined, but SBML only uses data attributes, so all SBML class diagrams use only the top two portions of a UML class box (see the right-hand diagram of Figure 1).



Figure 1: (Left) The general form of a UML class diagram. (Right) Example of a class diagram of the sort seen in SBML. SBML classes never use operators, so SBML class diagrams only show the top two parts.

As mentioned above, the names of ordinary (concrete) classes begin with a capital letter and are printed in an upright, bold, sans-serif typeface. The names of attributes begin with a lower-case letter and generally use a mixed case (sometimes called “camel case”) style when the name consists of multiple words. Attributes and their data types appear in the part below the class name, with one attribute defined per line. The colon character on each line separates the name of the attribute (on the left) from the type of data that it stores (on the right). The subset of data types permitted for SBML attributes is given in Section 3.1.

In the right-hand diagram of Figure 1, the symbols `attribute` and `anotherAttribute` represent attributes of the object class `ExampleClass`. The data type of `attribute` is `int`, and the data type of `anotherAttribute` is `double`. In the scheme used by SBML for translating UML to XML, object attributes map directly to XML attributes. Thus, in XML, `ExampleClass` would yield an element of the form `<element attribute="42" anotherAttribute="10.0">`.

Notice that the element name is not `<ExampleClass ...>`. Somewhat paradoxically, the name of the element is *not* the name of the UML class defining its structure. The reason for this may be subtle at first, but quickly becomes obvious: object classes define the form of an object’s *content*, but a class definition by itself does not define the *label* or symbol referring to an instance of that content. It is this label that becomes the name of the XML element. In XML, this symbol is most naturally equated with an element name. This point will hopefully become more clear with additional examples below.

Subelements

We use UML composite aggregation to indicate a class object can have other class objects as parts. Such containment hierarchies map directly to element-subelement relationships in XML. Figure 2 gives an example.



Figure 2: Example illustrating composite aggregation: the definition of one class of objects employing another class of objects in a part-whole relationship. In this particular example, an instance of a **Whole** class object must contain exactly one instance of a **Part** class object, and the symbol referring to the **Part** class object is `inside`. In XML, this symbol becomes the name of a subelement and the content of the subelement follows the definition of **Part**.

The line with the black diamond indicates composite aggregation, with the diamond located on the “container” side and the other end located at the object class being contained. The label on the line is the symbol used to refer to instances of the contained object, which in XML, maps directly to the name of an

XML element. The class pointed to by the aggregation relationship (**Part** in Figure 2) defines the *contents* of that element. Thus, if we are told that some element named **barney** is of class **Whole**, the following is an example XML fragment consistent with the class definition of Figure 2:

```
<barney A="110" B="some string">
  <inside C="444.4">
</barney>
```

Sometimes numbers are placed above the line near the “contained” side of an aggregation to indicate how many instances can be contained. The common cases in SBML are the following: [0..*] to signify a list containing zero or more; [1..*] to signify a list containing at least one; and [0..1] to signify exactly zero or one. The absence of a numerical label means “exactly 1”. This notation appears throughout this specification document.

Inheritance

Classes can inherit properties from other classes. Since SBML only uses data attributes and not operations, inheritance in SBML simply involves data attributes from a parent class being inherited by child classes. Inheritance is indicated by a line between two classes, with an open triangle next to the parent class; Figure 3 illustrates this. In this example, the instances of object class **Child** would have not only attributes C and D, but also attributes A and B. All of these attributes would be required (not optional) on instances of class **Child** because they are mandatory on both **Parent** and **Child**.

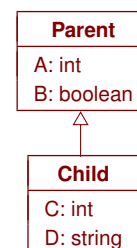


Figure 3: Inheritance.

Additional notations for XML purposes

Not everything is easily expressed in plain UML. For example, it is often necessary to indicate some constraints placed on the values of an attribute. In computer programming uses of UML, such constraints are often expressed using Object Constraint Language (OCL), but since we are most interested in the XML rendition of SBML, in this specification we use XML Schema 1.0 (when possible) as the language for expressing value constraints. Constraints on the values of attributes are written as expressions surrounded by braces ({ }) after the data type declaration, as in the example of Figure 4.



Figure 4: A more complex example definition drawing on the concepts introduced so far in this section. Both **Sbml** and **Model** are derived from **SBase**; further, **Sbml** contains a single **Model** object named `model`. Note the constraints on the values of the attributes in **Sbml**; they are enclosed in braces and written in XML Schema language. The particular constraints here state that both the `level` and `version` attributes must be present, and that the values are fixed as indicated.

In other situations, when something cannot be concisely expressed using a few words of XML Schema, we write constraints using English language descriptions surrounded by braces ({ }). To help distinguish these from literal XML Schema, we set the English text in a slanted typeface. The text accompanying all SBML component definitions provides explanations of the constraints and any other conditions applicable to the use of the components.

1 *Compatibility issues and warnings*

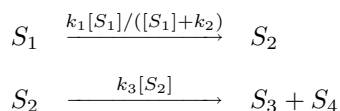
2 One important and confusing point that goes against the grain of XML must be highlighted: the order
3 in which subelements appear within SBML elements *is* significant and *must* follow the order given in the
4 corresponding object definition. This ordering is also difficult to express in plain UML, so we resort to
5 using the approach of stating ordering requirements as constraints written in English and (again) enclosed
6 in braces ({ }). Figure 8 on page 14 gives an example of this.

7 The ordering restriction also holds true when a subclass inherits attributes and elements from a base class:
8 the base class attributes and elements must occur before those introduced by the subclass.

9 This ordering constraint stems from aspects of XML Schema beyond our control (specifically, the need to
10 use XML Schema's **sequence** construct to define the object classes). It is an occasional source of software
11 compatibility problems, because validating XML parsers will generate errors if the ordering within an XML
12 element does not correspond to the SBML object class definition.

2 Overview of SBML

The following is an example of a simple network of biochemical reactions that can be represented in SBML:



In this particular set of chemical equations above, the symbols in square brackets (e.g., “[S_1]”) represent concentrations of molecular species, the arrows represent reactions, and the formulas above the arrows represent the rates at which the reactions take place. (And while this example uses concentrations, it could equally have used other measures such as molecular counts.) Broken down into its constituents, this model contains a number of components: reactant species, product species, reactions, reaction rates, and parameters in the rate expressions. To analyze or simulate this network, additional components must be made explicit, including compartments for the species, and units on the various quantities.

SBML allows models of arbitrary complexity to be represented. Each type of component in a model is described using a specific type of data **object** that organizes the relevant information. The top level of an SBML model definition consists of lists of these components, with every list being optional:

| | |
|--|----------------|
| <i>beginning of model definition</i> | |
| <i>list of function definitions (optional)</i> | (Section 4.3) |
| <i>list of unit definitions (optional)</i> | (Section 4.4) |
| <i>list of compartment types (optional)</i> | (Section 4.5) |
| <i>list of species types (optional)</i> | (Section 4.6) |
| <i>list of compartments (optional)</i> | (Section 4.7) |
| <i>list of species (optional)</i> | (Section 4.8) |
| <i>list of parameters (optional)</i> | (Section 4.9) |
| <i>list of initial assignments (optional)</i> | (Section 4.10) |
| <i>list of rules (optional)</i> | (Section 4.11) |
| <i>list of constraints (optional)</i> | (Section 4.12) |
| <i>list of reactions (optional)</i> | (Section 4.13) |
| <i>list of events (optional)</i> | (Section 4.14) |
| <i>end of model definition</i> | |

The meaning of each component is as follows:

Function definition: A named mathematical function that may be used throughout the rest of a model.

Unit definition: A named definition of a new unit of **measurement**, or a redefinition of an existing SBML default unit. Named units can be used in the expression of quantities in a model.

Compartment Type: A type of location where reacting entities such as chemical substances may be located.

Species type: A type of entity that can participate in reactions. **Typical** examples of species types include ions such as Ca^{2+} , molecules such as glucose or ATP, and more.

Compartment: A well-stirred container of a particular type and finite size where species may be located. A model may contain multiple compartments of the same compartment type. Every species in a model must be located in a compartment.

Species: A pool of entities of the same *species type* located in a specific *compartment*.

Parameter: A quantity with a symbolic name. In SBML, the term *parameter* is used in a generic sense to refer to named quantities regardless of whether they are constants or variables in a model. SBML Level 2 provides the ability to define parameters that are global to a model as well as parameters that are local to a single reaction.

1 *Initial Assignment:* A mathematical expression used to determine the initial conditions of a model. This
2 type of **object** can only be used to define how the value of a variable can be calculated from other
3 values and variables at the start of simulated time.

4 *Rule:* A mathematical expression added to the set of equations constructed based on the reactions defined
5 in a model. Rules can be used to define how a variable's value can be calculated from other variables,
6 or used to define the rate of change of a variable. The set of rules in a model can be used with the
7 reaction rate equations to determine the behavior of the model with respect to time. The set of rules
8 constrains the model for the entire duration of simulated time.

9 *Constraint:* A means of detecting out-of-bounds conditions during a dynamical simulation and optionally
10 issuing diagnostic messages. Constraints are defined by an arbitrary mathematical expression comput-
11 ing a true/false value from model variables, parameters and constants. An SBML constraint applies at
12 all instants of simulated time; however, the set of constraints in model should not be used to *determine*
13 the behavior of the model with respect to time.

14 *Reaction:* A statement describing some transformation, transport or binding process that can change the
15 amount of one or more species. For example, a reaction may describe how certain entities (reactants) are
16 transformed into certain other entities (products). Reactions have associated kinetic rate expressions
17 describing how quickly they take place.

18 *Event:* A statement describing an instantaneous, discontinuous change in a set of variables of any type
19 (species quantity, compartment size or parameter value) when a triggering condition is satisfied.

20 A software package can read an SBML model description and translate it into its own internal format for
21 model analysis. For example, a package might provide the ability to simulate the model by constructing
22 differential equations representing the network and then perform numerical time integration on the equations
23 to explore the model's dynamic behavior. By supporting SBML as an input and output format, different
24 software tools can all operate on an identical external representation of a model, removing opportunities for
25 errors in translation and assuring a common starting point for analyses and simulations.

3 Preliminary definitions and principles

This section covers certain concepts and constructs that are used repeatedly in the rest of SBML Level 2.

3.1 Primitive data types

Most primitive types in SBML are taken from the data types defined in XML Schema 1.0 (Biron and Malhotra, 2000; Fallside, 2000; Thompson et al., 2000). A few other primitive types are defined by SBML itself. What follows is a summary of the XML Schema types and the definitions of the SBML-specific types. Note that while we have tried to provide accurate and complete summaries of the XML Schema types, the following should not be taken to be normative definitions of these types. Readers should consult the XML Schema 1.0 specification for the normative definitions of the XML types used by SBML.

3.1.1 Type string

The XML Schema 1.0 type `string` is used to represent finite-length strings of characters. The characters permitted to appear in XML Schema `string` include all Unicode characters (Unicode Consortium, 1996) except for two delimiter characters, 0xFFFE and 0xFFFF (Biron and Malhotra, 2000). In addition, the following quoting rules specified by XML for character data (Bray et al., 2000) must be obeyed:

- The ampersand (&) character must be escaped using the entity `&`;
- The apostrophe (') and quotation mark (") characters must be escaped using the entities `'` and `"`; , respectively, when those characters are used to delimit a string attribute value.

Other XML built-in character or entity references, e.g., `<` and `&x1A;`, are permitted in strings.

3.1.2 Type boolean

The XML Schema 1.0 type `boolean` is used as the data type for SBML object attributes that represent binary true/false values. XML Schema 1.0 defines the possible literal values of `boolean` as the following: “true”, “false”, “1”, and “0”. The value “1” maps to “true” and the value “0” maps to “false”.

Note that there is a discrepancy between the value spaces of type `boolean` as defined by XML Schema 1.0 and MathML: the latter uses only “true” and “false” to represent boolean values and “0” and “1” are interpreted as numbers. Software tools should take care to not to use “0” and “1” as boolean values in MathML expressions. See further discussion in Section 3.4.4.

3.1.3 Type int

The XML Schema 1.0 type `int` is used to represent decimal integer numbers in SBML. The literal representation of an `int` is a finite-length sequence of decimal digit characters with an optional leading sign (“+” or “-”). If the sign is omitted, “+” is assumed. The value space of `int` is the same as a standard 32-bit signed integer in programming languages such as C, i.e., 2147483647 to -2147483648.

3.1.4 Type positiveInteger

The XML Schema 1.0 type `positiveInteger` is used to represent nonzero, nonnegative, decimal integers: i.e., 1, 2, 3, The literal representation of an integer is a finite-length sequence of decimal digit characters, optionally preceded by a positive sign (“+”). There is no restriction on the absolute size of `positiveInteger` values in XML Schema; however, the only situations where this type is used in SBML involve very low-numbered integers. Consequently, applications may safely treat `positiveInteger` as unsigned 32-bit integers.

3.1.5 Type double

The XML Schema 1.0 type `double` is the data type of floating point numerical quantities in SBML. It is restricted to IEEE double-precision 64-bit floating point type IEEE 754-1985. The value space of `double` consists of (a) the numerical values $m \times 2^x$, where m is an integer whose absolute value is less than 2^{53} ,

and x is an integer between -1075 and 970, inclusive, (b) the special value positive infinity (INF), (c) the special value negative infinity (-INF), and (d) the special value not-a-number (NaN). The order relation on the values is the following: $x < y$ if and only if $y - x$ is positive for values of x and y in the value space of **double**. Positive infinity is greater than all other values other than NaN. NaN is equal to itself but is neither greater nor less than any other value in the value space. (Software implementors should consult the XML Schema 1.0 definition of **double** for additional details about equality and relationships to IEEE 754-1985.)

The general form of **double** numbers is “ xey ”, where x is a decimal number (the mantissa), “e” is a separator character, and y is an exponent; the meaning of this is “ x multiplied by 10 raised to the power of y ”, i.e., $x \times 10^y$. More precisely, a **double** value consists of a mantissa with an optional leading sign (“+” or “-”), optionally followed by the character E or e followed by an integer (the exponent). The mantissa must be a decimal number: an integer optionally followed by a period (.) optionally followed by another integer. If the leading sign is omitted, “+” is assumed. An omitted E or e and exponent means that a value of 0 is assumed for the exponent. If the E or e is present, it must be followed by an integer or an error results. The integer acting as an exponent must consist of a decimal number optionally preceded by a leading sign (“+” or “-”). If the sign is omitted, “+” is assumed. The following are examples of legal literal **double** values:

-1E4, +4, 234.234e3, 6.02E-23, 0.3e+11, 2, 0, -0, INF, -INF, NaN

As described in Section 3.4, SBML uses a subset of the MathML 2.0 standard (W3C, 2000b) for expressing mathematical formulas in XML. This is done by stipulating that the MathML language be used whenever a mathematical formula must be written into an SBML model. Doing this, however, requires facing two problems: first, the syntax of numbers in scientific notation (“e-notation”) is different in MathML from that just described for **double**, and second, the value space of integers and floating-point numbers in MathML is not defined in the same way as in XML Schema 1.0. We elaborate on these issues in Section 3.4.2; here we summarize the solution taken in SBML. First, within MathML, the mantissa and exponent of numbers in “e-notation” format must be separated by one `<sep/>` element. This leads to numbers of the form `<cn type="e-notation"> 2 <sep/> -5 </cn>`. Second, SBML stipulates that the representation of numbers in MathML expressions obey the same restrictions on values as defined for types **double** and **int** (Section 3.1.3).

3.1.6 Type ID

The XML Schema 1.0 type ID is identical to the XML 1.0 type ID. The literal representation of this type consists of strings of characters restricted as summarized in Figure 5.

```

NCNameChar ::= letter | digit | '.' | '-' | '_' | ':' | CombiningChar | Extender
ID          ::= ( letter | '-' | ':' ) NCNameChar*

```

Figure 5: Type ID expressed in the variant of BNF used by the XML 1.0 specification (Bray et al., 2004). The characters (and) are used for grouping, the character * indicates “zero or more times”, and the character | indicates “or”. The production letter consists of the basic upper and lower case alphabetic characters of the Latin alphabet along with a large number of related characters defined by Unicode 2.0; similarly, the production digit consists of the numerals 0 . . 9 along with related Unicode 2.0 characters. The CombiningChar production is a list of characters that add such things as accents to the preceding character. (For example, the Unicode character #x030A when combined with ‘a’ produces ‘á’.) The Extender production is a list of characters that extend the shape of the preceding character. Please consult the XML 1.0 specification (Bray et al., 2004) for the complete definitions of letter, digit, CombiningChar, and Extender.

In SBML, type ID is the data type of the **metaid** attribute on **Sbase**, described in Section 3.2. An important aspect of ID is the XML requirement that a given value of ID must be unique throughout an XML document. All data values of type ID are considered to reside in a single common global namespace spanning the entire XML document, regardless of the **attribute** where type ID is used and regardless of the level of nesting of the **objects** (or XML elements).

3.1.7 Type SId

The type SId is the type of the **id** attribute found on the majority of SBML components. SId is a data type derived from the basic XML type **string**, but with restrictions about the characters permitted and the sequences in which those characters may appear. The definition is shown in Figure 6 on the next page.

```

1         letter ::= 'a'..'z','A'..'Z'
2         digit  ::= '0'..'9'
3         idChar ::= letter | digit | '_'
4         SId    ::= ( letter | '_' ) idChar*

```

Figure 6: The definition of the type SId. (Please see the caption of Figure 5 for an explanation of the notation.)

The equality of SId values is determined by an exact character sequence match; i.e., comparisons of these identifiers must be performed in a case-sensitive manner. This applies to all uses of SId.

The SId is purposefully not derived from the XML ID type (Section 3.1.6). Using XML’s ID would force all SBML identifiers to exist in a single global namespace, which would affect not only the form of local parameter definitions but also future SBML extensions for supporting model/submodel composition. Further, the use of the ID type for SBML identifiers would have limited utility because MathML 2.0 ci elements are not of the type IDREF (see Section 3.4). Since the IDREF/ID linkage cannot be exploited in MathML constructs, the utility of the XML ID type is greatly reduced. Finally, unlike ID, SId does not include Unicode character codes; the identifiers are plain text.

3.1.8 Type UnitSId

The type UnitSId is derived from SId (Section 3.1.7) and has identical syntax. The UnitSId type is used as the data type for the identifiers of units (Section 4.4.1) and for references to unit identifiers in SBML objects. The purpose of having a separate data type for such identifiers is enable the space of possible unit identifier values to be separated from the space of all other identifier values in SBML. The equality of UnitSId values is determined by an exact character sequence match; i.e., comparisons of these identifiers must be performed in a case-sensitive manner.

A number of reserved symbols are defined in the space of values of UnitSId. These reserved symbols are the list of base unit names defined in Table 1 on page 37, and the SBML built-in units “substance”, “volume”, “area”, “length”, and “time” listed in Table 2 on page 41. These symbols and their use is described in Section 4.4.

3.1.9 Type SBOTerm

The type SBOTerm is used as the data type of the attribute sboTerm on Sbase. The type consists of strings of characters matching the restricted pattern described in Figure 7.

```

28         digit    ::= '0'..'9'
29         SBOTerm  ::= 'SBO:' digit digit digit digit digit digit digit

```

Figure 7: The definition of SBOTerm. The SBOTerm type consists of strings beginning with SBO: and followed by seven decimal digits. (Please see the caption of Figure 5 for an explanation of the notation.)

Examples of valid string values of type SBOTerm are “SBO:0000014” and “SBO:0003204”. These values are meant to be the identifiers of terms from an ontology whose vocabulary describes entities and processes in computational models. Section 5 provides more information about the ontology and principles for the use of these terms in SBML models.

3.2 Type SBase

Nearly every object composing an SBML Level 2 model definition has a specific data type that is derived directly or indirectly from a single abstract type called Sbase. In addition to serving as the parent class for most other classes of objects in SBML, this base type is designed to allow a modeler or a software package to attach arbitrary information to each major element or list in an SBML model. The definition of Sbase is presented in Figure 8 on the following page.

Sbase contains two attributes and two subelements, all of which are optional: metaid, sboTerm, notes and annotation. These are discussed separately in the following subsections.

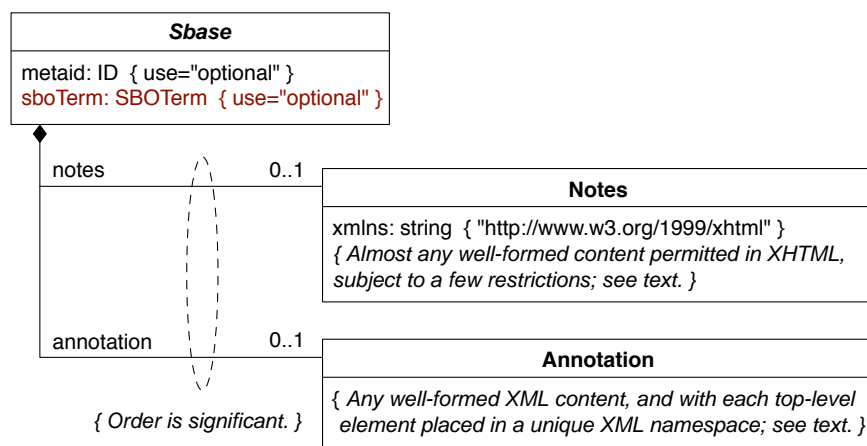


Figure 8: The definition of *Sbase*. Please refer to Section 1.4 for a summary of the UML notation used here. Note that the order of appearance of subelements *notes* and *annotation* is significant in instances of objects derived from *SBase*: *notes* must always come before *annotation*. (This requirement arises from XML Schema 1.0.)

3.2.1 The metaid attribute

The **metaid attribute** is present for supporting metadata annotations using RDF (Resource Description Format; Lassila and Swick, 1999). It has a data type of XML ID (the XML identifier type; see Section 3.1.6), which means each **metaid** value must be globally unique within an SBML file. The **metaid** value serves to identify a model component for purposes such as referencing that component from metadata placed within **annotation elements** (see Section 3.2.4). Such metadata can use RDF **description** elements, in which an RDF attribute called “**rdf:about**” points to the **metaid** identifier of an object defined in the SBML model. This topic is discussed in greater detail in Section 6.

3.2.2 The sboTerm attribute

The **attribute** called **sboTerm** is provided on *Sbase* to support the use of the Systems Biology Ontology (SBO; see Section 5). When a value is given to this attribute, it must conform to the data type **SBOTerm** (Sections 3.1.9). SBO terms are a type of optional annotation, and each different class of SBML object derived from *Sbase* imposes its own requirements about the values permitted for **sboTerm**. Specific details on the permitted values are provided with the definitions of SBML classes throughout this specification document, and a broader discussion is provided in Section 5.

3.2.3 The notes element

The **element** **notes** in *Sbase* is a container for XHTML 1.0 (Pemberton et al., 2002) content. It is intended to serve as a place for storing optional information intended to be seen by humans. An example use of the **notes element** would be to contain formatted user comments about the model **element** in which the **notes element** is enclosed. Every object derived directly or indirectly from type *Sbase* can have a separate value for **notes**, allowing users considerable freedom when adding comments to their models.

XHTML 1.0 is simply a formulation of HTML 4 in XML 1.0. This means the full power of HTML formatting is available for use in **notes** content. The intention behind requiring XHTML (rather than, for example, plain HTML or plain text) for **notes** content is to balance several conflicting goals: (1) choosing a format for notes that is compatible with the XML form of SBML (plain HTML would not be); (2) supporting an international formatting standard so that users have more control over the appearance of notes and can predict to some degree how their notes will be displayed in different tools and environments (which argues against using plain-text notes); and (3) achieving these goals using an approach that is hopefully easy enough for software developers to support using off-the-shelf programming libraries. It is worth noting in passing that the requirement for XHTML does not *prevent* users from entering plain-text content with simple space/tab/newline formatting: it merely requires using the standard `<pre>...</pre>` element of (X)HTML.

1 Modern libraries for displaying and editing (X)HTML content are commonly available in contemporary
2 software programming environments, and software developers may wish to avail themselves of these facilities
3 rather than implementing their own XHTML support systems.

4 *XML namespace requirements for notes*

5 The XML content of **notes** elements must declare the use of the XHTML XML namespace. This can be
6 done in multiple ways. One way is to place a namespace declaration for the appropriate namespace URI
7 (which is <http://www.w3.org/1999/xhtml>) on the top-level **Sbml** object (see Section 4.1) and then reference
8 the namespace in the **notes** content using a prefix. The following example illustrates this approach:

```
9 <sbml xmlns="http://www.sbml.org/sbml/level2/version2" level="2" version="3"  
10     xmlns:xhtml="http://www.w3.org/1999/xhtml">  
11     ...  
12     <notes>  
13         <xhtml:body>  
14             <xhtml:center><xhtml:h2>A Simple Mitotic Oscillator</xhtml:h2></xhtml:center>  
15             <xhtml:p>A minimal cascade model for the mitotic oscillator  
16             involving cyclin and cdc2 kinase</xhtml:p>  
17         </xhtml:body>  
18     </notes>  
19     ...
```

20 Another approach is to declare the XHTML namespace within the **notes** content itself, as in the following
21 example:

```
22     ...  
23     <notes>  
24         <body xmlns="http://www.w3.org/1999/xhtml">  
25             <center><h2>A Simple Mitotic Oscillator</h2></center>  
26             <p>A minimal cascade model for the mitotic oscillator  
27             involving cyclin and cdc2 kinase</p>  
28         </body>  
29     </notes>  
30     ...
```

31 The `xmlns="http://www.w3.org/1999/xhtml"` declaration on **body** as shown above changes the default XML
32 namespace within it, such that all of its content is by default in the XHTML namespace. This is a particularly
33 convenient approach because it obviates the need to prefix every element with a namespace prefix (i.e., **xhtml**:
34 in the previous case). Other approaches are also possible.

35 *The content of notes*

36 SBML does not require the content of **notes** to be any particular XHTML element; the content can be
37 almost any well-formed XHTML content. There are only two simple restrictions. The first restriction comes
38 from the requirements of XML: the **notes** element must not contain an XML declaration nor a DOCTYPE
39 declaration. That is, **notes** must *not* contain

```
40 <?xml version="1.0" encoding="UTF-8"?>
```

41 nor (where the following is only one specific example of a DOCTYPE declaration)

```
42 <!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"  
43     "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
```

44 The second restriction is intended to balance freedom of content with the complexity of implementing software
45 that can interpret the content. The content of **notes** in SBML can consist only of the following possibilities:

- 46 1. A complete XHTML document (minus the XML and DOCTYPE declarations, of course), that is,
47 XHTML content beginning with the **html** tag. The following is an example skeleton:

```

1      <notes>
2          <html xmlns="http://www.w3.org/1999/xhtml">
3              ...
4          </html>
5      </notes>

```

2. The **body** element from an XHTML document. The following is an example skeleton:

```

7      <notes>
8          <body xmlns="http://www.w3.org/1999/xhtml">
9              ...
10         </body>
11     </notes>

```

3. Any XHTML content that would be permitted within a **body** element. If this consists of multiple elements, each one must declare the XML namespace separately. The following is an example **fragment**:

```

14     <notes>
15         <p xmlns="http://www.w3.org/1999/xhtml">
16             ...
17         </p>
18         <p xmlns="http://www.w3.org/1999/xhtml">
19             ...
20         </p>
21     </notes>

```

Another way to summarize the restrictions above is simply to say that the content of an SBML **notes** element can be only be a complete **html** element, a **body** element, or whatever is permitted inside a **body** element. In practice, this does not limit in any meaningful way what can be placed inside a **notes** element; for example, if an application or modeler wants to put a complete XHTML page, including a **head** element, it can be done by putting in everything starting with the **html** container. However, the restrictions above do make it somewhat simpler to write software that can read and write the **notes** content. Appendix F describes one possible approach to doing just that.

3.2.4 The annotation element

Whereas the **notes element** described above is a container for content to be shown directly to humans, the **annotation element** is a container for optional software-generated content *not* meant to be shown to humans. Every object derived from *Sbase* can have its own value for **annotation**. The **element's content** type is XML type **any**, allowing essentially arbitrary **well-formed XML** data content. SBML places only a few restrictions on the organization of the content; these are intended to help software tools read and write the data as well as help reduce conflicts between annotations added by different tools.

The use of XML namespaces in annotation

At the outset, software developers should keep in mind that multiple software tools may attempt to read and write annotation content. To reduce the potential for collisions between annotations written by different applications, SBML Level 2 Version 3 stipulates that tools must use XML namespaces (Bray et al., 1999) to specify the intended vocabulary of every annotation. The application's developers must choose a URI (*Universal Resource Identifier*; Harold and Means 2001; W3C 2000a) reference that uniquely identifies the vocabulary the application will use, and a prefix string for the annotations. Here is an example. Suppose an application uses the URI <http://www.mysim.org/ns> and the prefix **mysim** when writing annotations related to screen layout. The content of an annotation might look like the following:

```

45     <annotation>
46         <mysim:nodecolors xmlns:mysim="http://www.mysim.org/ns"
47             mysim:bgcolor="green" mysim:fgcolor="white"/>
48     </annotation>

```

In this particularly simple **example**, the content consists of a single XML element (**nodecolors**) with two attributes (**bgcolor**, **fgcolor**), all of which are prefixed by the string **mysim**. (Presumably this particular content would have meaning to the hypothetical application in question.) The content in this particular

example is small, but it should be clear that there could easily have been an arbitrarily large amount of data placed inside the `mysim:nodecolors` element.

The key point of the example above is that application-specific annotation data is entirely contained inside a single *top-level element* within the SBML **annotation** container. SBML Level 2 Version 3 places the following restrictions on annotations:

- Within a given SBML **annotation** element, there can only be one top-level element using a given namespace. An annotation element can contain multiple top-level elements but each must be in a different namespace.
- No top-level element in an **annotation** may use an SBML XML namespace, either explicitly by referencing one of the SBML XML namespace URIs or implicitly by failing to specify any namespace on the annotation. *As of SBML Level 2 Version 3, the defined SBML namespaces are the following URIs:*
 - <http://www.sbml.org/sbml/level1>
 - <http://www.sbml.org/sbml/level2>
 - <http://www.sbml.org/sbml/level2/version2>
 - <http://www.sbml.org/sbml/level2/version3>
- The ordering of top-level elements within a given **annotation** element is *not* significant. An application should not expect that its annotation content appears first in the **annotation** element, nor in any other particular location. *Moreover, the ordering of top-level annotation elements may be changed by different applications as they read and write the same SBML file.*

The use of XML namespaces in this manner is intended to improve the ability of multiple applications to place annotations on SBML model **elements** with reduced risks of interference or name collisions. Annotations stored by different simulation packages can therefore coexist in the same model definition. The rules governing the content of **annotation** elements are designed to enable applications to easily add, change, and remove their annotations from SBML elements while simultaneously preserving annotations inserted by other applications when mapping SBML from input to output.

*As a further simplification for developers of software and to improve software interoperability, applications are only required to preserve other annotations (i.e., annotations they do not recognize) when those annotations are self-contained entirely within **annotation**, complete with namespace declarations. The following is an example:*

```
<annotation>
  <topLevelElement xmlns:"URI">
    ... content in the namespace identified by "URI"...
  </topLevelElement>
</annotation>
```

Some more examples hopefully will make *these points* more clear. The next example is invalid because it contains a top-level element in the SBML XML namespace—this happens because no namespace is declared for the `<cytoplasm>` element, which means by default it falls into the **enclosing** SBML namespace:

```
<annotation>
  <cytoplasm/>
</annotation>
```

The following example is *also* invalid, *this time* because it contains two top-level elements using the same XML namespace. Note that it does not matter that these are two different top-level elements (`<nodecolors>` and `<textcolors>`); what matters is that these separate elements are both in the same namespace rather than having been collected and placed inside one overall container element for that namespace.

```
<annotation>
  <mysim:nodecolors xmlns:mysim="http://www.mysim.org/ns"
    mysim:bgcolor="green" mysim:fgcolor="white"/>
  <mysim:textcolors xmlns:mysim="http://www.mysim.org/ns"
    mysim:bgcolor="green" mysim:fgcolor="white"/>
</annotation>
```

1 On the other hand, the following example is valid:

```
2 <annotation>
3   <mysim:geometry xmlns:mysim="http://www.mysim.org/ns"
4     mysim:bgcolor="green" mysim:fgcolor="white">
5     <graph:node xmlns:graph="http://www.graph.org/ns"
6       graph:x="4" graph:y="5" />
7   </mysim:geometry>
8   <othersim:icon xmlns:othersim="http://www.othersim.com/">
9     WS2002
10  </othersim:icon>
11 </annotation>
```

12 For completeness, we note that annotations can legally be empty:

```
13 <annotation />
```

14 It is worth keeping in mind that although XML namespace names must be URIs, they are (like all XML
15 namespace names) *not required* to be directly usable in the sense of identifying an actual, retrieval document
16 or resource on the Internet (Bray et al., 1999). URIs such as <http://www.mysim.org/> may appear as though
17 they are (e.g.,) Internet addresses, but there are not the same thing. This style of URI strings, using a
18 domain name and other parts, is only a simple and commonly-used way of creating a unique name string.

19 Finally, note that the namespaces being referred to here are XML namespaces specifically in the context of
20 the **annotation element** on *Sbase*. The namespace issue here is unrelated to the namespaces discussed in
21 Section 3.3.1 in the context of component identifiers in SBML.

22 *Content of annotations and implications for software tools*

23 The **annotation element** in the definition of *Sbase* exists in order that software developers may attach
24 optional application-specific data to the **elements** in an SBML model. However, it is important that this
25 facility not be misused. In particular, it is *critical* that data essential to a model definition or that can be
26 encoded in existing SBML **elements** is *not* stored in **annotation**. Parameter values, functional dependencies
27 between model **elements**, etc., should not be recorded as annotations. It is crucial to keep in mind the fact
28 that data placed in annotations can be freely ignored by software applications. If such data affects the
29 interpretation of a model, then software interoperability is greatly impeded.

30 Here are examples of the kinds of data that may be appropriately stored in **annotation**: (a) information
31 about the graphical layout of model components; (b) application-specific processing instructions that do not
32 change the essential meaning of a model; (c) identification information for cross-referencing components in
33 a model with items in a data resource such as a database; and (d) information about the model that cannot
34 be readily encoded in existing SBML **elements**.

35 *Standardized format for certain classes of annotations*

36 For case (c) above (i.e., cross-references between model components and data resources, SBML Level 2
37 Version 3 recommends a standard format for use within **annotation** elements. It should be used in preference
38 to proprietary syntaxes to maximize the likelihood that multiple software tools will converge on the same
39 syntax for this kind of information. The recommended scheme is described in Section 6.

40 **3.3 The id and name attributes on SBML components**

41 As will become apparent below, most **objects** in SBML include two common **attributes**: **id** and **name**. These
42 **attributes** are not defined on *Sbase* (as explained in Section 3.3.3 below), but where they do appear, the
43 common rules of usage described below apply.

44 **3.3.1 The id attribute and identifier scoping**

45 The **id attribute** is mandatory on most **objects** in SBML. It is used to identify a component within the
46 model definition. Other SBML **objects** can refer to the component using this identifier. The data type of **id**
47 is always either **Sid** (Section 3.1.7) or **UnitSid** (Section 3.1.8), depending on the object in question.

1 A model can contain a large number of components representing different parts. This leads to a problem
2 in deciding the scope of an identifier: in what contexts does a given identifier X represent the same thing?
3 The approaches used in existing simulation packages tend to fall into two categories which we may call
4 global and local. The *global* approach places all identifiers into a single global space of identifiers, so that an
5 identifier X represents the same thing wherever it appears in a given model definition. The *local* approach
6 places symbols in separate identifier namespaces, depending on the context, where the context may be, for
7 example, individual reaction rate expressions. The latter approach means that a user may use the same
8 identifier X in different rate expressions and have each instance represent a different quantity.

9 The fact that different simulation programs may use different rules for identifier resolution poses a problem
10 for the exchange of models between simulation tools. Without careful consideration, a model written out in
11 SBML format by one program may be misinterpreted by another program. SBML Level 2 must therefore
12 include a specific set of rules for treating identifiers and their scopes.

13 The scoping rules in SBML Level 2 are relatively straightforward and are intended to avoid this problem
14 with a minimum of requirements on the implementation of software tools:

- 15 • The identifier (i.e., the value of the **attribute id**) of every **FunctionDefinition**, **CompartmentType**, **Speci-**
16 **esType**, **Compartment**, **Species**, **Parameter**, **Reaction**, **SpeciesReference**, **ModifierSpeciesReference**, **Event**,
17 and **Model**, must be unique across the set of all such identifiers in the model. This means, for example,
18 that a reaction and a species definition cannot both have the same identifier.
- 19 • The identifier of every **UnitDefinition** must be unique across the set of all such identifiers in the model.
20 However, unit identifiers live in a separate space of identifiers from other identifiers in the model, **by**
21 **virtue of the fact that the data type of unit identifiers is `UnitSid` (Section 3.1.8) and not `Sid`.**
- 22 • Each **Reaction** instance (see Section 4.13) establishes a separate private local space for local **Parameter**
23 identifiers. Within the definition of that reaction, local parameter identifiers override (shadow) identical
24 identifiers outside of that reaction. Of course, the corollary of this is that local parameters inside a
25 **Reaction** object instance are not visible to other objects outside of that reaction.

26 The set of rules above can enable software packages using either local or global identifier spaces for parame-
27 ters to exchange SBML model definitions. Software systems using local identifiers for parameters internally
28 should, in principle, be able to accept SBML model definitions without needing to change component identi-
29 fiers. Environments using a common global space of identifiers for parameters internally can perform
30 manipulations of the identifiers of local parameters within reaction definitions to avoid identifier collisions.

31 The guidelines described here will hopefully provide a clean transition path to future levels of SBML, when
32 submodels are introduced (Section 8.1). Submodels will provide the ability to compose one model from
33 a collection of other models. This capability will have to be built on top of SBML Level 2's namespace
34 organization. A straightforward approach to handling namespaces is to make each submodel's space be
35 private. The rules governing identifier scoping within a submodel can simply be the Level 2 namespace rule
36 described here, with each submodel having its own (to itself, global) namespace.

37 **3.3.2 The name attribute**

38 In contrast to the **id attribute**, the **name attribute** is optional and is not intended to be used for cross-
39 referencing purposes within a model. Its purpose instead is to provide a human-readable label for the
40 component. The data type of **name** is the type **string** defined in XML Schema (Biron and Malhotra, 2000;
41 Thompson et al., 2000) and discussed further in Section 3.1. SBML imposes no restrictions as to the content
42 of **name attributes** beyond those restrictions defined by the **string** type in XML Schema.

43 The recommended practice for handling **name** is as follows. If a software tool has the capability for displaying
44 the content of **name attributes**, it should display this content to the user as a component's label instead of
45 the component's **id**. If the user interface does not have this capability (e.g., because it cannot display or use
46 special characters in symbol names), or if the **name attribute** is missing on a given component, then the user
47 interface should display the value of the **id attribute** instead. (Script language interpreters are especially
48 likely to display **id instead of name**.)

1 As a consequence of the above, authors of systems that automatically generate the values of **id attributes**
2 should be aware some systems may display the **id**'s to the user. Authors therefore may wish to take some
3 care to have their software create **id** values that are: (a) reasonably easy for humans to type and read; and
4 (b) likely to be meaningful, e.g., the **id attribute** is an abbreviated form of the name **attribute** value.

5 An additional point worth mentioning is although there are restrictions on the uniqueness of **id** values (see
6 Section 3.3.1 above), there are no restrictions on the uniqueness of **name** values in a model. This allows
7 software packages leeway in assigning component identifiers.

8 **3.3.3 Why id and name are not defined on Sbase**

9 Although many SBML components **feature id** and **name**, these **attributes** are purposefully not defined on
10 **Sbase**. There are several reasons for this.

- 11 • The presence of an SBML identifier **attribute (id)** necessarily requires specifying scoping rules for the
12 corresponding identifiers. However, the **Sbase** abstract type is used as the basis for defining components
13 whose scoping rules are in some cases different from each other. (See Section 3.3.1 for more details).
14 If **Sbase** were to have an **id attribute**, then the specification of **Sbase** would need a default scoping
15 rule and this would then have to be overloaded on derived classes that needed different scoping. This
16 would make the SBML specification **even more** complex.
- 17 • Identifier are optional on some SBML components and required on most others. If **id** were defined as
18 optional on **Sbase**, most component classes would separately have to redefine **id** as being mandatory—
19 hardly an improvement over the current arrangement. Conversely, if **id** were defined as mandatory on
20 **Sbase**, it would prevent it from being optional on components where it *is* currently optional.
- 21 • The **Sbase** abstract type is used as the base type for certain **objects** such as **Sbml**, **AssignmentRule**,
22 etc., which do not have identifiers because these components do not need to be referenced by other
23 components. If **Sbase** had a mandatory **id attribute**, *all* objects of these other types in a model
24 would then need to be assigned unique identifiers. Similarly, **because Sbase is the base type of the**
25 **listOf_____ lists**, putting **id** on **Sbase** would require all of these lists in a model to be given
26 identifiers. This would be a needless burden on software developers, tools, and SBML users, requiring
27 them to generate and store additional identifiers for objects that never need them.
- 28 • **Sbase** does not have a **name** simply because such an **attribute** is always paired with an **id**. Without **id**
29 on **Sbase**, it does not make sense to have **name**.

30 **3.4 Mathematical formulas in SBML Level 2**

31 Mathematical expressions in SBML Level 2 are represented using MathML 2.0 (W3C, 2000b). MathML is
32 an international standard for encoding mathematical expressions using XML. There are two principal facets
33 of MathML, one for encoding content (i.e., the semantic interpretation of a mathematical expression), and
34 another for encoding presentation or display characteristics. SBML only makes direct use of a subset of
35 the content portion of MathML. By borrowing a separately-developed XML standard, we can avoid having
36 to define a specialized syntax for mathematical expressions in SBML and simultaneously leverage existing
37 intellectual and technological work already done in the MathML community. However, it is not possible
38 to produce a completely smooth and conflict-free interface between MathML and other standards used by
39 SBML (in particular, XML Schema). Two specific issues and their resolutions are discussed in Sections 3.4.2.

40 The XML namespace URI for all MathML elements is <http://www.w3.org/1998/Math/MathML>. Everywhere
41 MathML content is allowed in SBML, the MathML elements must be properly placed within the MathML 2.0
42 namespace. In XML, this can be accomplished in a number of ways, and the examples throughout this
43 specification illustrate the use of this namespace and MathML in SBML. Please refer to the W3C document
44 by Bray et al. (1999) for more technical information about using XML namespaces.

3.4.1 Subset of MathML used in SBML Level 2

The subset of MathML 2.0 elements used in SBML Level 2 is similar to that used by CellML (Hedley et al., 2001), another model definition language with similar goals as SBML. The subset of MathML elements used in SBML is listed below:

- *token*: `cn`, `ci`, `csymbol`, `sep`
- *general*: `apply`, `piecewise`, `piece`, `otherwise`, `lambda` (the last is restricted to use in [FunctionDefinition](#))
- *relational operators*: `eq`, `neq`, `gt`, `lt`, `geq`, `leq`
- *arithmetic operators*: `plus`, `minus`, `times`, `divide`, `power`, `root`, `abs`, `exp`, `ln`, `log`, `floor`, `ceiling`, `factorial`
- *logical operators*: `and`, `or`, `xor`, `not`
- *qualifiers*: `degree`, `bvar`, `logbase`
- *trigonometric operators*: `sin`, `cos`, `tan`, `sec`, `csc`, `cot`, `sinh`, `cosh`, `tanh`, `sech`, `csch`, `coth`, `arcsin`, `arccos`, `arctan`, `arcsec`, `arccsc`, `arccot`, `arcsinh`, `arccosh`, `arctanh`, `arcsech`, `arccsch`, `arccoth`
- *constants*: `true`, `false`, `notanumber`, `pi`, `infinity`, `exponentiale`
- *annotation*: `semantics`, `annotation`, `annotation-xml`

The inclusion of logical operators, relational operators, `piecewise`, `piece`, and `otherwise` elements facilitates the encoding of discontinuous expressions. Note that MathML elements for representing partial differential calculus are not included. We anticipate that the requirements for partial differential calculus will be addressed in proposals for future SBML geometry representations (see Section 8.1).

As defined by MathML 2.0, the semantic interpretation of the mathematical functions listed above follows the definitions of the functions laid out by Abramowitz and Stegun (1977) and Zwillinger (1996). Readers are directed to these sources and the MathML specification for information about such things as which principal values of the inverse trigonometric functions to use.

Software authors should take particular note of the MathML semantics of the N-ary operators `plus`, `times`, `and`, `or` and `xor`, when they are used with different numbers of arguments. The MathML specification (W3C, 2000b) appendix C.2.3 describes the semantics for these operators with zero, one, and more arguments.

The following are the only attributes permitted on MathML elements in SBML (in addition to the `xmlns` attribute on `math` elements):

- `style`, `class` and `id` on any element;
- `encoding` and `definitionURL` on `csymbol` elements; and
- `type` on `cn` elements.

Missing values for these attributes are to be treated in the same way as defined by MathML. These restrictions on attributes are designed to confine the MathML elements to their default semantics and to avoid conflicts in the interpretation of the type of token elements.

3.4.2 Numbers and `cn` elements

In MathML, literal numbers are written as the content portion of a particular element called `cn`. This element takes an optional attribute, `type`, used to indicate the *type* of the number (such as whether it is meant to be an integer or a floating-point quantity). Here is an example of its use:

```
<math xmlns="http://www.w3.org/1998/Math/MathML">
  <apply>
    <times/> <cn type="integer"> 42 </cn> <cn type="real"> 3.3 </cn>
  </apply>
</math>
```

1 The content of a **cn** element must be a number. The number can be preceded and succeeded by whitespace
2 (see Section 3.4.5). The following are the only permissible values for the **type** attribute on MathML **cn**
3 elements: “**e-notation**”, “**real**”, “**integer**”, and “**rational**”. The value of the **type** attribute defaults to
4 “**real**” if it is not specified on a given **cn** element.

5 *Value space restrictions on cn content*

6 SBML imposes certain restrictions on the value space of numbers allowed in MathML expressions. According
7 to the MathML 2.0 specification, the values of the content of **cn** elements do not necessarily have to conform
8 to any specific floating point or integer representations designed for CPU implementation. For example, in
9 strict MathML, the value of a **cn** element could exceed the maximum value that can be stored in a IEEE
10 64 bit floating point number (IEEE 754). This is different from the XML Schema type **double** that is used
11 in the definition of floating point **attributes** of **objects** in SBML; the XML Schema **double** is restricted
12 to IEEE double-precision 64-bit floating point type IEEE 754-1985. To avoid an inconsistency that would
13 result between numbers elsewhere in SBML and numbers in MathML expressions, SBML Level 2 Version 3
14 imposes the following restriction on MathML content appearing in SBML:

- 15 • Integer values (i.e., the **values** of **cn** elements having **type**="integer" and both values in **cn** elements
16 having **type**="rational") must conform to the **int** type used elsewhere in SBML (Section 3.1.3)
- 17 • Floating-point values (i.e., the content of **cn** elements having **type**="real" or **type**="e-notation")
18 must conform to the **double** type used elsewhere in SBML (Section 3.1.5)

19 *Syntactic differences in the representation of numbers in scientific notation*

20 It is important to note that MathML uses a style of scientific notation that differs from what is defined in XML
21 Schema, and consequently what is used in SBML attribute values. The MathML 2.0 type “**e-notation**” (as
22 well as the type “**rational**”) requires the mantissa and exponent to be separated by one **<sep/>** element. The
23 mantissa must be a real number and the exponent part must be a signed integer. This leads to expressions
24 such as

```
25 <cn type="e-notation"> 2 <sep/> -5 </cn>
```

26 for the number 2×10^{-5} . It is especially important to note that the expression

```
27 <cn type="e-notation"> 2e-5 </cn>
```

28 is *not valid* in MathML 2.0 and therefore cannot be used in MathML content in SBML. However, elsewhere in
29 SBML, when an attribute value is declared to have the data type **double** (a type taken from XML Schema),
30 the compact notation “**2e-5**” is in fact allowed. In other words, within MathML expressions contained in
31 SBML (and *only* within such MathML expressions), numbers in scientific notation must take the form **<cn**
32 **type="e-notation"> 2 <sep/> -5 </cn>**, and everywhere else they must take the form “**2e-5**”.

33 This is a regrettable difference between two standards that SBML relies upon, but it is not feasible to
34 redefine these types within SBML because the result would be incompatible with parser libraries written to
35 conform with the MathML and XML Schema standards. It is also not possible to use XML Schema to define
36 a data type for SBML attribute values permitting the use of the **<sep/>** notation, because XML attribute
37 values cannot contain XML elements—that is, **<sep/>** cannot appear in an XML attribute value.

38 *Units of numbers in MathML cn expressions*

39 What units should be attributed to values appearing inside MathML **cn** elements? One answer is to assume
40 that the units should be “whatever units appropriate in the context where the number appears”. This
41 implies that units can always be assigned unambiguously to any number by inspecting the expression in
42 which it appears, and this turns out to be false. Another answer is that numbers should be considered
43 “dimensionless”. Many people argue that this is the correct interpretation, but even if it is, there is an
44 overriding practical reason why it cannot be adopted for SBML’s domain of application: when numbers
45 appear in expressions in SBML, they are *rarely intended* by the modeler to have the unit “**dimensionless**”

1 even if the unit is not declared—the numbers are *supposed* to have specific units, but the units are usually
2 undeclared. (Being “dimensionless” is not the same as having *undeclared* units!) If SBML defined numbers
3 as being *by default* dimensionless, it would result in many models being technically incorrect without the
4 modeler being aware of it unless their software tools performed dimensional analysis. Most software tools
5 today still do not perform dimensional analysis, and so the inconsistency of units (and potential errors in
6 the model) would not be detected until other researchers and database curators attempted to use the model
7 in software packages that *did* check units. We believe the negative impact on interoperability and people’s
8 confidence in SBML as a reliable medium would be too high.

9 As a result, the current approach in SBML is to leave the default units of literal numbers in MathML
10 content undefined. Software packages and modelers are encouraged to explicitly add unit declarations to
11 numbers. There is a simple mechanism in SBML for associating units with numbers: do not use literal
12 numbers at all; instead, define **Parameter** objects (Section 4.9) for every quantity, declare units for each such
13 parameter value in its definition, and then insert the parameters in place of numbers in expressions. This
14 leads to mathematical formulas whose units can be fully determined, permitting software tools to perform
15 dimensional analysis and potentially report problems with a model.

16 In summary: literal numbers appearing within MathML content in SBML have no declared units.

17 **3.4.3 Use of `ci` elements in MathML expressions in SBML**

18 The content of a `ci` element must be an SBML identifier that is declared elsewhere in the model. The
19 identifier can be preceded and succeeded by whitespace. The set of possible identifiers that can appear in a
20 `ci` element depends on the containing `element` in which the `ci` is used:

- 21 • If a `ci` element appears in the body of a **FunctionDefinition** object (Section 4.3), the referenced identifier
22 must be either (i) one of the declared arguments to that function, or (ii) the identifier of a previously
23 defined **FunctionDefinition** object in the model.
- 24 • Otherwise, the referenced identifier must be that of a **Species**, **Compartment**, **Parameter**, **FunctionDefi-**
25 **nition**, or **Reaction** object defined in the model. The following are the only possible interpretations of
26 using such an identifier in SBML:
 - 27 – *Species identifier*: When a **Species** identifier occurs in a `ci` element, it represents the quantity of
28 that species in units of either *amount of substance* or units of *concentration*, depending on the
29 species’ definition; see Section 4.8.5.
 - 30 – *Compartment identifier*: When a **Compartment** identifier occurs in a `ci` element, it represents the
31 size of the compartment. The units of measurement associated with the size of the compartment
32 are those given by the **Compartment** instance’s **units** attribute value; see Section 4.7.5.
 - 33 – *Parameter identifier*: When a **Parameter** identifier occurs in a `ci` element, it represents the nu-
34 merical value assigned to that parameter. The units associated with the parameter’s value are
35 those given by the **Parameter** instance’s **units** attribute; see Section 4.9.3.
 - 36 – *Function identifier*: When a **FunctionDefinition** identifier occurs in a `ci` element, it represents a
37 call to that function. Function references in MathML occur in the context of using MathML’s
38 **apply** and often involve supplying arguments to the function; see Section 4.3. The units associated
39 with the value returned by the function call are the overall units of the mathematical expression
40 contained in the function definition.
 - 41 – *Reaction identifier*: When a **Reaction** identifier occurs in a `ci` element, it represents the rate of
42 that reaction as defined by the **math** expression in the **KineticLaw** object within the **Reaction**. The
43 units associated with that rate are *substance/time*, where the *substance* and *time* units established
44 by the values of the SBML built-in units “**substance**” and “**time**”, respectively. These units may
45 be redefined globally in the model; see Section 4.4.3. If a **Reaction** instance has no **KineticLaw**,
46 using its reaction identifier this way is an error (perhaps indicating that the model is incomplete).

47 The content of `ci` elements in MathML formulas outside of a **KineticLaw** or **FunctionDefinition** must always
48 refer to objects declared in the top level global namespace; i.e., SBML uses “early binding” semantics. Inside

of **KineticLaw**, a **ci** element can additionally refer to local parameters defined within that **KineticLaw** instance; see Section 4.13.5 for more information.

3.4.4 Interpretation of boolean values

As noted already in Section 3.1.2, there is another unfortunate difference between the XML Schema 1.0 and MathML 2.0 standards that impacts mathematical expressions in SBML: in XML Schema, the value space of type **boolean** includes “**true**”, “**false**”, “**1**”, and “**0**”, whereas in MathML, only “**true**” and “**false**” count as boolean values.

The impact of this difference thankfully is minimal because the XML Schema definition is only used for attribute values on SBML objects, and those values turn out never to be accessible from MathML content in SBML—values of boolean attributes on SBML objects can never enter into MathML expressions. Nevertheless, software authors and users should be aware of the difference and in particular that “**0**” and “**1**” are interpreted as numerical quantities in mathematical expressions. There is no automatic conversion of “**0**” or “**1**” to boolean values in contexts where booleans are expected. This allows stricter type checking and unit verification during the validation of mathematical expressions.

3.4.5 Handling of whitespace

MathML 2.0 defines “whitespace” in the same way as XML does, i.e., the space character (Unicode hexadecimal code 0020), horizontal tab (code 0009), newline or line feed (code 000A), and carriage return (code 000D). In MathML, the content of elements such as **cn** and **ci** can be surrounded by whitespace characters. Prior to using the content, this whitespace is “trimmed” from both ends: all whitespace at the beginning and end of the content is removed (Ausbrooks et al., 2003). For example, in `<cn> 42 </cn>`, the amount of white space on either side of the “42” inside the `<cn> ... </cn>` container does not matter. Prior to interpreting the content, the whitespace is removed altogether.

3.4.6 Use of **csymbol** elements in MathML expressions in SBML

SBML Level 2 uses the MathML **csymbol** element to denote certain built-in mathematical entities without introducing reserved names into the component identifier namespace. The **encoding** attribute of **csymbol** must be set to “**text**”. The **definitionURL** should be set to one of the following URIs defined by SBML:

- <http://www.sbml.org/sbml/symbols/time>. This represents the current simulation time. See Section 3.4.7 for more information. The units of the current time entity is determined from the built-in **time** of Table 2 on page 41.
- <http://www.sbml.org/sbml/symbols/delay>. This represents a delay function. The delay function has the form $delay(x, d)$, taking two MathML expressions as arguments. Its value is the value of argument x at d time units before the current time. There are no restrictions on the form of x . The units of the d parameter are determined from the built-in **time**. The value of the d parameter, when evaluated, must be numerical (i.e., a number in MathML real, integer, or “e-notation” format) and be greater than or equal to 0. The *delay* function is useful for representing biological processes having a delayed response, but where the detail of the processes and delay mechanism is not relevant to the operation of a given model. See Section 3.4.7 below for additional considerations surrounding the use of this **csymbol**.

The following examples demonstrate these concepts. The XML fragment below encodes the formula $x + t$, where t stands for time.

```
<math xmlns="http://www.w3.org/1998/Math/MathML">
  <apply>
    <plus/>
    <ci> x </ci>
    <csymbol encoding="text" definitionURL="http://www.sbml.org/sbml/symbols/time">
      t
    </csymbol>
  </apply>
</math>
```


In the fragment above, the use of the token `t` is mostly a convenience for human readers—the string inside the `csymbol` could have been almost anything, because it is essentially ignored by MathML parsers and SBML. Some MathML and SBML processors will take note of the token and use it when presenting the mathematical formula to users, but the token used has no impact on the interpretation of the model and it does *not* enter into the SBML component identifier namespace. In other words, the SBML model cannot refer to `t` in the example above. The content of the `csymbol` element is for rendering purposes only and can be ignored by the parser.

As a further example, the following XML fragment encodes the equation $k + \text{delay}(x, 0.1)$ or alternatively $k_t + x_{t-0.1}$:

```

10 <math xmlns="http://www.w3.org/1998/Math/MathML">
11   <apply>
12     <plus/>
13     <ci> k </ci>
14     <apply>
15       <csymbol encoding="text" definitionURL="http://www.sbml.org/sbml/symbols/delay">
16         delay
17       </csymbol>
18       <ci> x </ci>
19       <cn> 0.1 </cn>
20     </apply>
21   </apply>
22 </math>

```

Note that the URI in the value of `definitionURL`, as all URIs, is intended to serve as a unique identifier and is not intended to be dereferenced as an Internet address. There is nothing actually located at the address <http://www.sbml.org/sbml/symbols/delay>.

3.4.7 Simulation time

The principal use of SBML is to represent quantitative dynamical models whose behaviors manifest themselves over time. In defining an SBML model using constructs such as reactions, time is most often implicit and does not need to be referred to in the mathematical expressions themselves. However, sometimes an explicit time dependency needs to be stated, and for this purpose, the *time* `csymbol` (described above in Section 3.4.6) may be used. This *time* symbol refers to “instantaneous current time” in a simulation, frequently given the literal name t in one’s equations.

An assumption in SBML is that “start time” or “initial time” in a simulation is zero, that is, if t_0 is the initial time in the system, $t_0 = 0$. This corresponds to the most common scenario. Initial conditions in SBML take effect at time $t = 0$. There is no mechanism in SBML for setting the initial time to a value other than 0. To refer to a different time in a model, one approach is to define a **Parameter** for a new time variable and use an **AssignmentRule** in which the assignment expression subtracts a value from the `csymbol` *time*. For example, if the desired offset is 2 time units, the MathML expression would be

```

39 <math xmlns="http://www.w3.org/1998/Math/MathML">
40   <apply>
41     <minus/>
42     <csymbol encoding="text" definitionURL="http://www.sbml.org/sbml/symbols/time"> t
43   </csymbol>
44     <cn> 2 </cn>
45   </apply>
46 </math>

```

SBML’s assignment rules (Section 4.11.3) can be used to express mathematical statements that hold true at all moments, so using an assignment rule with the expression above will result in the value being equal to $t - 2$ at every point in time. A parameter assigned this value could then be used elsewhere in the model, its value could be plotted by a simulator, etc.

3.4.8 Initial conditions and special considerations

The identifiers of **Species**, **Compartment**, **Parameter**, and **Reaction** object instances in a given SBML model refer to the main variables in a model. Depending on certain attributes of these objects (e.g., the **attribute constant** on species, compartments and parameters—this and other conditions are explained in the relevant sections elsewhere in this document), some of the variables may have constant values throughout a simulation, and others' values may change. These changes in values over time are determined by the system of equations constructed from the model's reactions, initial assignments, rules, and events.

As described in Section 3.4.7, an SBML model's simulation is assumed to begin at $t = 0$. The availability of the *delay* **csymbol** (Section 3.4.6) introduces the possibility that at $t \geq 0$, mathematical expressions in a model may draw on values of model components from time *prior* to $t = 0$. A simulator may therefore need to compute the values of variables at time points $t_i \leq 0$ to allow the calculation of values required for the evaluation of delay expressions in the model for $t \geq 0$. If there are no delays in the model, then $t_i = 0$.

The following is how the definitions of the model should be applied:

1. At time t_i :

- Every **Species**, **Compartment**, and **Parameter** whose definition includes an initial value is assigned that value. If an **element** has **constant**="false", its value may be changed by other constructs or reactions in a model according to the steps below; if **constant**="true", only an **InitialAssignment** can override the value.
- All **InitialAssignment** definitions take effect at t_i and continue to have effect up to and including $t = 0$, overriding any initial values on **Species**, **Compartment** and **Parameter**. Since **InitialAssignments** contain mathematical formulas, different values may be computed at each time step t in $t_i \leq t \leq 0$.

2. For time $t \geq t_i$:

- **AssignmentRule** and **AlgebraicRule** definitions are in effect from this point in time forward and may influence the values of **Species** quantity, **Compartment** size, and **Parameter** values. (Note there cannot be both an **AssignmentRule** and an **InitialAssignment** for the same identifier; see Section 4.11.)

3. At time $t = 0$:

- The system of equations constructed by combining **AssignmentRule** equations, **AlgebraicRule** equations, **RateRule** equations, and the equations constructed from the **Reaction** definitions in the model, are used to obtain consistent initial conditions for numerical solver algorithms. (Note that there cannot be both an **AssignmentRule** and a **RateRule** for the same identifier, or both an **AssignmentRule** and an **InitialAssignment** for the same identifier; see Section 4.11.3.)
- **Constraint** definitions begin to take effect (and a constraint violation may result; see Section 4.12).

4. For time $t > 0$:

- **RateRule** definitions can begin to take effect.
- **Event** definitions can begin to take effect. (Note that an **Event** cannot be defined to change the value of a variable that is also the subject of an **AssignmentRule**; see Section 4.14.)
- System simulation proceeds.

To reiterate: in modeling situations that do not involve the use of the *delay* **csymbol**, $t_i = 0$, but this does not alter the steps above.

3.4.9 MathML expression data types

MathML operators in SBML each return results in one of two possible types: boolean and numerical. By *numerical* type, we mean either (1) a number in MathML real, integer, rational, or "e-notation" format; or (2) the **csymbol** for delay or the **csymbol** for time described in Section 3.4.6. The following guidelines summarize the different possible cases.

1 The relational operators (**eq**, **neq**, **gt**, **lt**, **geq**, **leq**), the logical operators (**and**, **or**, **xor**, **not**), and the boolean
2 constants (**false**, **true**) always return boolean values. As noted in Section 3.4.4, the numbers 0 and 1 do
3 not count as boolean values in MathML contexts in SBML.

4 The type of an operator referring to a **FunctionDefinition** is determined by the type of the top-level operator
5 of the expression in the **math** element of the **FunctionDefinition** instance, and can be boolean or numerical.

6 All other operators, values and symbols return numerical results.

7 The roots of the expression trees used in the following contexts must yield boolean values:

- 8 • the arguments of the MathML logical operators (**and**, **or**, **xor**, **not**);
- 9 • the second argument of a MathML **piece** operator;
- 10 • the **trigger** element of an SBML **Event**; and
- 11 • the **math** element of an SBML **Constraint**.

12 The roots of the expression trees used in the following contexts can optionally yield boolean values:

- 13 • the arguments to the **eq** and **neq** operators;
- 14 • the first arguments of MathML **piece** and **otherwise** operators; and
- 15 • the top level expression of a function definition.

16 The roots of expression trees in other contexts must yield numerical values.

17 The type of expressions should be used consistently. The set of expressions that make up the first arguments
18 of the **piece** and **otherwise** operators within the same **piecewise** operator should all return values of the
19 same type. The arguments of the **eq** and **neq** operators should return the same type.

20 **3.4.10 Consistency of units in mathematical expressions and treatment of unspecified units**

21 Strictly speaking, physical validity of mathematical formulas requires not only that physical quantities added
22 to or equated with each other have the same fundamental dimensions and units of measurement; it also
23 requires that the application of operators and functions to quantities produces sensible results. Yet, in
24 real-life models today, these conditions are often and sometimes legitimately disobeyed. SBML must again
25 balance the ability of modelers to represent whatever model they desire against the need for the represented
26 models to be self-consistent and detailed enough for different software tools to be able to manipulate and
27 analyze them.

28 The approach taken in SBML is to focus requirements principally on preventing inconsistency of a model.
29 In a nutshell, the SBML requirement is that *what is expressed in a model must be consistent to the extent*
30 *that this can be determined using the information available in the model.*

31 *Unit requirements for mathematical expressions*

32 The SBML requirement for consistency of units is defined in terms of dimensional analysis applied recursively
33 to every operator and function and every argument to them. Some conditions *must* be true, and others *should*
34 be true (but are not absolutely required for model validity). The following conditions must be true in a model:

- 35 1. All arguments to the following operators must have the same units (regardless of what those units
36 happen to be): **plus**, **minus**, **eq**, **neq**, **gt**, **lt**, **geq**, **leq**.
- 37 2. The units of each argument in a call to a **FunctionDefinition** must match the units expected by the
38 **lambda** expression within the **math** expression of that **FunctionDefinition** instance.
- 39 3. All of the possible return values from **piece** and **otherwise** subelements of a **piecewise** expression
40 must have the same units, regardless of what those units are. (Otherwise, the **piecewise** expression
41 would return values having different units depending on which case evaluated to true.)
- 42 4. For the **delay** **csymbol** (Section 3.4.6) function, which has the form $delay(x, d)$, the second argument
43 d must match the model's unit of *time* (i.e., the “**time**” built-in unit).

In addition to the above, the following conditions *should* hold true (and software developers may wish to consider having their software warn users if one or more of the following conditions is not true):

1. The units of each argument to the following operators should be “dimensionless”: `exp`, `ln`, `log`, `factorial`, `sin`, `cos`, `tan`, `sec`, `csc`, `cot`, `sinh`, `cosh`, `tanh`, `sech`, `csch`, `coth`, `arcsin`, `arccos`, `arctan`, `arcsec`, `arccsc`, `arccot`, `arcsinh`, `arccosh`, `arctanh`, `arcsech`, `arccsch`, `arccoth`.
2. The two arguments to `power`, which are of the form $power(a, b)$ with the meaning a^b , should be as follows: (1) if the second argument is an integer, then the first argument can have any units; (2) if the second argument b is a rational number n/m , it must be possible to derive the m -th root of $(a\{\text{units}\})^n$, where $\{\text{units}\}$ signifies the units associated with a ; otherwise, (3) the units of the first argument must be “dimensionless”. The second argument (b) should always have units of “dimensionless”.
3. The two arguments to `root`, which are of the form $root(n, a)$ with the meaning $\sqrt[n]{a}$ and where the degree n is optional (defaulting to “2”), should be as follows: (1) if the optional degree qualifier n is an integer, then it must be possible to derive the n -th root of a ; (2) if the optional degree qualifier n is a rational n/m then it must be possible to derive the n -th root of $(a\{\text{units}\})^m$, where $\{\text{units}\}$ signifies the units associated with a ; otherwise, (3) the units of a must be “dimensionless”.
4. Since the units of literal numbers cannot be specified directly in SBML (see below), it is possible for the units of a `FunctionDefinition` object’s return value to be effectively different in different contexts where it is called. If a `FunctionDefinition`’s mathematical formula contains literal constants (i.e., numbers within MathML `cn` elements), the units of the constants should be identical in all contexts the function is called.

The units of other operators such as `abs`, `floor`, and `ceiling`, can be anything.

The final bulleted item above, regarding `FunctionDefinition`, warrants additional elaboration. An example may help illustrate the problem. Suppose the formula $x + 5$ is defined as a function, where x is an argument. The literal number 5 in SBML has unspecified units. If this function is called with an argument in moles, the only possible consistent unit for the return value is mole. If in another context in the same model, the function is called with an argument in seconds, the function return value can only be treated as being in seconds. Now suppose that a modeler decides to change all uses of seconds to milliseconds in the model. To make the function definition return the same quantity in terms of seconds, the 5 in the formula would need to be changed, but doing so would change the result of the function everywhere it is called—with the wrong consequences in the context where moles were intended. This illustrates the subtle danger of using numbers with unspecified units in function definitions. There are at least two approaches for avoiding this: (1) define separate functions for each case where the units of the constants are supposed to be different; or (2) declare the necessary constants as `Parameter` objects in the model (with declared units!) and pass those parameters as arguments to the function, avoiding the use of literal numbers in the function’s formula.

Treatment of unspecified units

There are only two ways to introduce numbers with unspecified units into mathematical formulas in SBML: using literal numbers (i.e., numbers enclosed in MathML `cn` elements), and using `Parameter` objects defined without unit declarations. All other quantities, in particular species and compartments, always have unit declarations (whether explicit or the defaults).

If an expression contains literal numbers and/or `Parameter` objects without declared units, the consistency or inconsistency of units may be impossible to determine. In the absence of a verifiable *inconsistency*, an expression in SBML is accepted as-is; the writer of the model is assumed to have written what they intended. Nevertheless, this is *not* equivalent to assuming the expression *does* have consistent units.

In some cases, it may be possible to determine that expressions containing unspecified units are inconsistent regardless of what units would be attributed to the unspecified quantities. For example, the expression

$$\frac{dX}{dt} = \frac{[Y] \cdot [Z]^n}{[Z]^m + 1} \cdot V$$

1 with X , Y and Z in units of substance, V in units of volume, and $m \neq n$, cannot ever be consistent, no
2 matter what units the literal 1 takes on. (This also illustrates the need not to stop verifying the units of an
3 expression immediately upon encountering an unspecified quantity—the rest of the expression may still be
4 profitably evaluated and checked for inconsistency.)

5 We advise modelers and software tools to declare the units of all quantities in a model, insofar as this is
6 possible, using the various mechanisms provided for this in SBML. Fully declared units can allow software
7 tools to perform dimensional analysis on the units of mathematical expressions, and such analysis can be
8 valuable in helping modelers produce correct models. In addition, it can allow model-wide operations such
9 as conversion or rescaling of units. The lack of declared units on quantities in a model does not render it
10 invalid SBML, but it reduces the types of consistency checks and useful operations (such as conversions and
11 translations) that software systems can perform.

12 **3.4.11 SBML does not define implicit unit conversions**

13 Implicit unit conversions do not exist in SBML. Consider the following example. Suppose that in some
14 model, a species S_1 has been declared as having a mass of 1 kg, and a second species S_2 has been declared
15 as having a mass of 500 g. What should be the result of evaluating an expression such as $S_1 > S_2$? If the
16 numbers alone are considered,

$$17 \quad 1 > 500$$

18 would evaluate to “false”, but if the units were implicitly converted by the software tool interpreting the
19 model,

$$20 \quad 1 \text{ kg} > 500 \text{ g}$$

21 would evaluate to “true”. This is a trivial example, but the problem for SBML is that implicit unit
22 conversions of this kind can lead to controversial situations where even humans do not agree on the answer.
23 Consequently, SBML only requires that mathematical expressions be evaluated numerically. It is up to
24 the model writer to ensure that the units on both sides of an expression match, by inserting explicit unit
25 conversion factors if necessary.

26 It is simple enough to avoid expressions with mixed units such as in the example above: a modeler or a
27 software tool can define a parameter that acts as a conversion factor, and then multiply one of the terms by
28 this parameter. Thus, a model could include a parameter “*g_per_kg*” with a value of 1000, and the expression
29 could be written as

$$30 \quad 1 * \textit{g_per_kg} > 500$$

31 which will then evaluate to “true”.

4 SBML components

In this section, we define each of the major components of SBML. We use the UML notation described in Section 1.4.3 for defining classes of objects. We also illustrate the use of SBML components by giving partial model definitions in XML. Section 7 provides many full examples of SBML in XML.

4.1 The SBML container

All well-formed XML documents must begin with an *XML declaration*, which specifies both the version of XML assumed and the document character encoding. The declaration begins with the characters `<?xml` followed by the XML **version** and **encoding** attributes. SBML Level 2 uses XML version 1.0 and requires a document encoding of UTF-8. Following this XML declaration, the outermost portion of a model expressed in SBML Level 2 Version 3 consists of an object of class **Sbml**, defined in Figure 9. This class contains two required attributes, for the SBML **level** and **version**, and one required subelement called **model** whose class is **Model**.

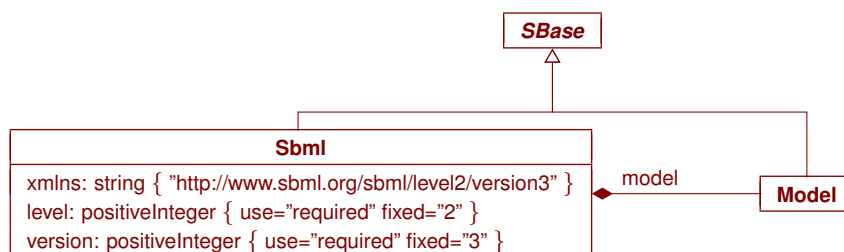


Figure 9: The definition of **Sbml** for SBML Level 2 Version 3. The class **Model** is defined in Section 4.2. Note that both **Sbml** and **Model** are subclasses of **Sbase**, and therefore inherit the attributes of that abstract class.

The following is an abbreviated example of these XML elements for an SBML Level 2 Version 3 document:

```
<?xml version="1.0" encoding="UTF-8"?>
<sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
  ...
</sbml>
```

The attribute **xmlns** declares the default XML namespace used within the **sbml** element. The URI for SBML Level 2 Version 3 is <http://www.sbml.org/sbml/level2/version3>. All elements must be placed in this namespace either by assigning the default namespace as shown above, or using a tag prefix on every element.

An SBML XML document must not contain elements or attributes in the SBML namespace that are not defined in this SBML Level 2 Version 3 Release 1 specification. Documents containing unknown elements or attributes placed in the SBML namespace do not conform to this SBML specification.

Readers may wonder why the SBML top-level XML element uses both a namespace URI identifying the SBML level and version, as well as separate XML attributes giving the level and version. Why is the information duplicated? There are several reasons. First, XML is only one possible serialization of SBML (albeit an extremely popular one at this time). Though most of this document is written with XML in mind, it is the intention behind the design of SBML that its object structure should be implementable in other languages and software systems. Programmatic access is easier if the level and version information are accessible directly as data rather than have to be extracted from a string. Second, generic high-level XML parsers may not give their calling programs access to the value of the **xmlns** attribute. Providing the information via separate attributes is a good backup measure. And finally, earlier in the history of SBML, it was expected that only the level needed to be encoded as part of the namespace URI (e.g., <http://www.sbml.org/sbml/level1>) because it was hoped that changes within levels would not require XML Schema changes. This has proven to be false, but SBML Level 1 (both versions) and the first version of SBML Level 2 still subscribe to this principle. This means that for these variants of SBML, software tools must look for a **version** attribute on the top-level element. For backwards compatibility with software that expects this, it makes more sense to keep the version and level attributes.

4.2 Model

The definition of **Model** is shown in Figure 10 on the following page. Only one instance of a **Model** object is allowed per instance of an SBML Level 2 Version 3 Release 1 document or data stream, and it must be located inside the `<sbml> ... </sbml>` element as described in Section 4.1.

The **Model** object has an optional attribute, `id`, used to give the model an identifier. The identifier must be a text string conforming to the syntax permitted by the `SId` data type described in Section 3.1.7. **Model** also has an optional `name` attribute, of type `string`. The `name` and `id` attributes must be used as described in Section 3.3.

Model serves as a container for components of classes **FunctionDefinition**, **UnitDefinition**, **CompartmentType**, **SpeciesType**, **Compartment**, **Species**, **Parameter**, **InitialAssignment**, **Rule**, **Constraint**, **Reaction** and **Event**. Instances of the classes are placed inside instances of classes **ListOfFunctionDefinitions**, **ListOfUnitDefinitions**, **ListOfCompartmentTypes**, **ListOfSpeciesTypes**, **ListOfCompartments**, **ListOfSpecies**, **ListOfParameters**, **ListOfInitialAssignments**, **ListOfRules**, **ListOfConstraints**, **ListOfReactions**, and **ListOfEvents**. The “list” classes are defined in Figure 10. All of the lists are optional, but if a given list container is present within the model, the list must not be empty; that is, it must have length one or more. The resulting XML data object for a full model containing every possible list would have the following form:

```
<?xml version="1.0" encoding="UTF-8"?>
<sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
  <model id="My_Model">
    <listOfFunctionDefinitions>
      one or more <functionDefinition> ... </functionDefinition> elements } optional
    </listOfFunctionDefinitions>
    <listOfUnitDefinitions>
      one or more <unitDefinition> ... </unitDefinition> elements } optional
    </listOfUnitDefinitions>
    <listOfCompartmentTypes>
      one or more <compartmentType> ... </compartmentType> elements } optional
    </listOfCompartmentTypes>
    <listOfSpeciesTypes>
      one or more <speciesType> ... </speciesType> elements } optional
    </listOfSpeciesTypes>
    <listOfCompartments>
      one or more <compartment> ... </compartment> elements } optional
    </listOfCompartments>
    <listOfSpecies>
      one or more <species> ... </species> elements } optional
    </listOfSpecies>
    <listOfParameters>
      one or more <parameter> ... </parameter> elements } optional
    </listOfParameters>
    <listOfInitialAssignments>
      one or more <initialAssignment> ... </initialAssignment> elements } optional
    </listOfInitialAssignments>
    <listOfRules>
      one or more elements of subclasses of Rule } optional
    </listOfRules>
    <listOfConstraints>
      one or more <constraint> ... </constraint> elements } optional
    </listOfConstraints>
    <listOfReactions>
      one or more <reaction> ... </reaction> elements } optional
    </listOfReactions>
    <listOfEvents>
      one or more <event> ... </event> elements } optional
    </listOfEvents>
  </model>
</sbml>
```

Although all the lists are optional, there are dependencies between SBML components such that defining some components requires defining others. An example is that defining a species requires defining a compartment, and defining a reaction requires defining a species. The dependencies are explained throughout the text.

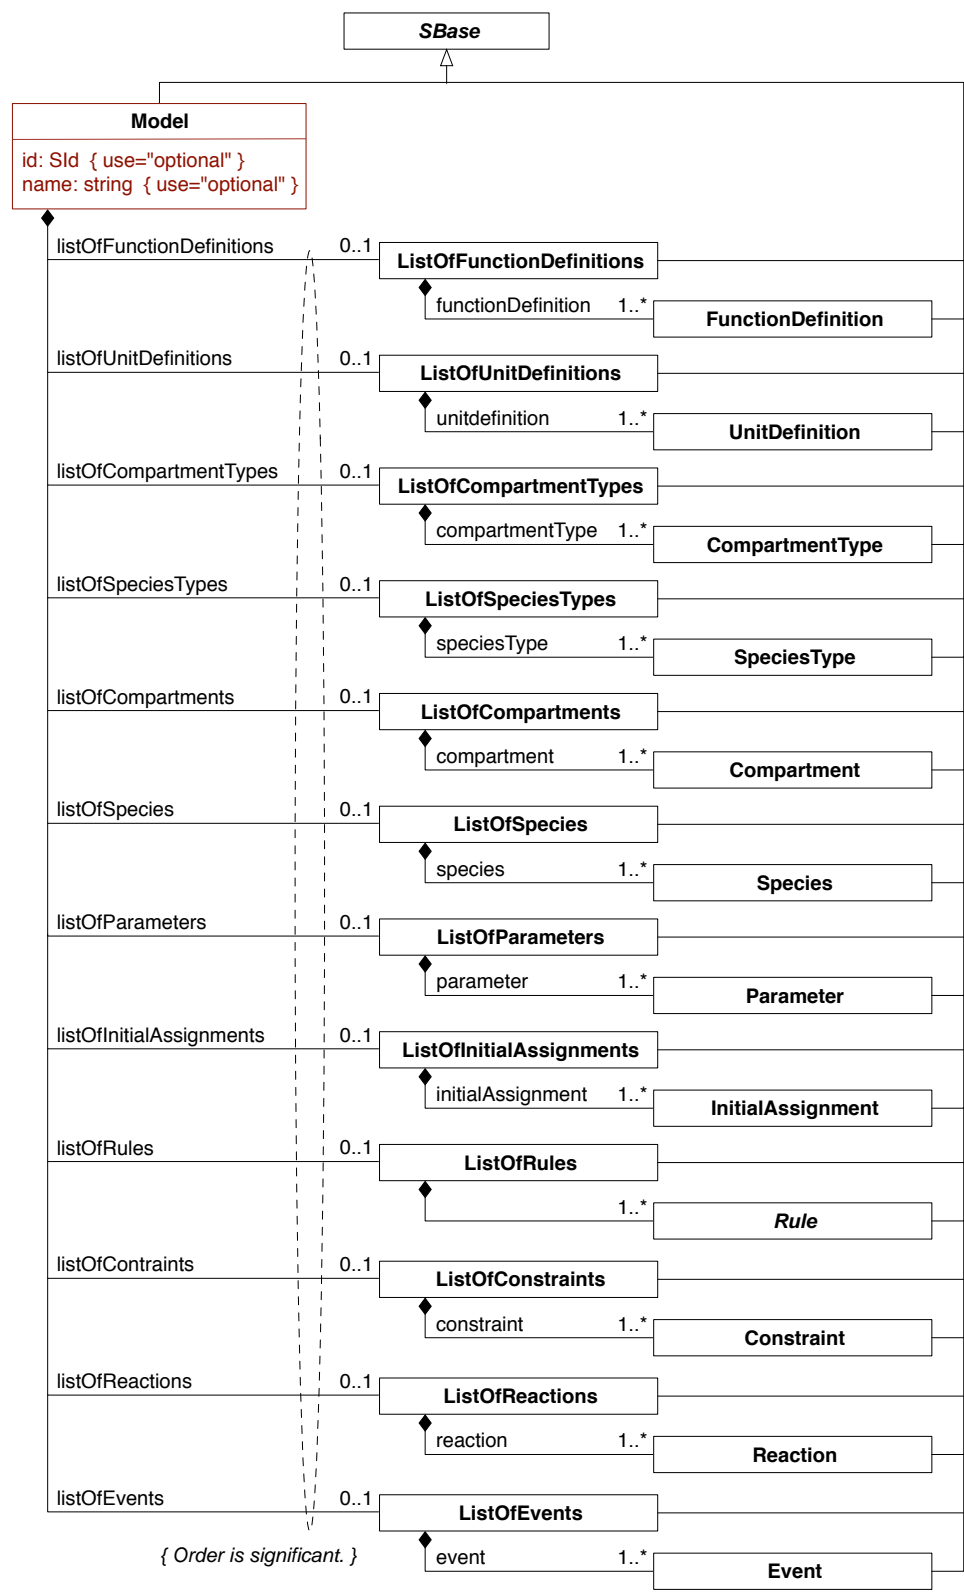


Figure 10: The definition of *Model* and the many helper classes *ListOfFunctionDefinitions*, *ListOfUnitDefinitions*, *ListOfCompartmentTypes*, *ListOfSpeciesTypes*, *ListOfCompartments*, *ListOfSpecies*, *ListOfParameters*, *ListOfInitialAssignments*, *ListOfRules*, *ListOfConstraints*, *ListOfReactions*, and *ListOfEvents*.

4.2.1 The ListOf container classes

The various **ListOf_____** classes defined in Figure 10 are merely containers used for organizing the main components of an SBML model. All are derived from the abstract class **Sbase** (Section 3.2), and inherit **Sbase**'s various attributes and subelements such as **metaid** and **annotation**, although in SBML Level 2 Version 3 Release 1 there are no defined SBO terms for the **sboTerm** attribute. The **ListOf_____** classes do not add any attributes of their own.

Readers may wonder about the motivations for using the **ListOf_____** containers. A simpler approach in XML might be to place the components all directly at the top level. We chose instead to group them within XML elements named after **listOfClassNames**, in part because we believe this helps organize the components and makes visual reading of models in XML easier. More importantly, the fact that the container classes are derived from **Sbase** means that software tools can add information about the lists themselves into each list container's **annotation**, a feature that a number of today's software tools exploit.

4.2.2 The sboTerm attribute

The **Model** object has an optional **sboTerm** attribute of type **SBOTerm** (see Sections 3.1.9 and 5). Values of this attribute must be chosen from identifiers referring to a modeling framework defined in SBO (i.e., terms derived from **SBO:0000004**, "modeling framework"). The SBO term chosen should be the most precise (narrow) term that defines the mathematical framework assumed by the *entire* model. An example might be "non-spatial discrete framework", which would mean the *entire* model was created under the assumption it will be simulated in a tool that turns substance quantities into discrete numbers and applies a stochastic simulation algorithm (e.g., Gillespie, 1977) to the model. Other frameworks are possible.

A dependency exists between the SBO term referenced by a **Model** instance and the SBO labels on individual **KineticLaw** instances within the model. In SBO, each definition of a rate expression itself includes an SBO reference to a mathematical framework. When a given **KineticLaw** element refers to a rate law term in SBO, it therefore simultaneously also identifies a framework. (It is possible for a model to mix frameworks, which is why it is important to enable individual reactions to indicate which one is assumed.) By contrast, if all the reactions in a given model use the same framework, it makes sense to characterize the model as a whole as assuming this. Thus, the label on **Model** should apply to the model *as a whole*, meaning that if **Model**'s **sboTerm** attribute has a value, *all* of that model's reaction rate expressions are assumed to be expressed using the same mathematical framework, and the application reading the model may ignore frameworks indicated by the definitions of the SBO terms referenced by the **KineticLaw** elements's **sboTerm** attributes. As discussed in Section 5, SBO labels are optional information on a model.

The absence of an **sboTerm** value on a **Model** instance means that either the model has not been labeled with SBO terms at all, or no modeling framework term in SBO at the time the model was created was deemed suitable for characterizing the model as a whole. If a **Model** instance has no value for its **sboTerm** attribute, the definitions of the SBO terms referenced by individual **KineticLaw** elements in the model may indicate the framework assumed for each reaction separately.

4.3 Function definitions

The **FunctionDefinition** object associates an identifier with a function definition. This identifier can then be used as the function called in subsequent MathML **apply** elements. **FunctionDefinition** is shown in Figure 11 on the next page.

Function definitions in SBML (also informally known as "user-defined functions") have *purposefully* limited capabilities. As is made more clear below, a function cannot reference parameters or other model quantities outside of itself; values must be passed as parameters to the function. Moreover, recursive and mutually-recursive functions are not permitted. The purpose of these limitations is to balance power against complexity of implementation. With the restrictions as they are, function definitions could be implemented as textual substitutions—they are simply macros. Software implementations therefore do not need the full function-definition machinery typically associated with programming languages.

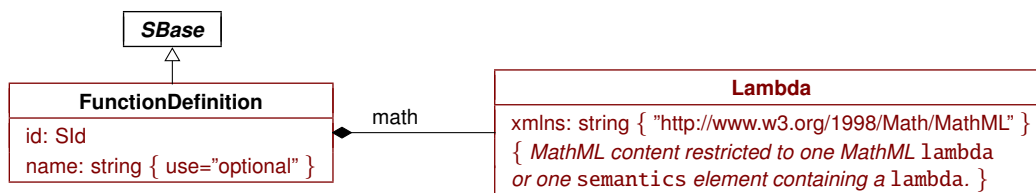


Figure 11: The definition of class **FunctionDefinition**. The contents of the **Lambda** class is a single MathML lambda expression (or a lambda surrounded by a semantics element). A function definition must contain exactly one math element defined by the **Lambda** class; note also that **Lambda** is not derived from **Sbase**, which means that the attributes defined on **Sbase** are not available on the math element. A sequence of one or more instances of **FunctionDefinition** objects can be located in an instance of **ListOfFunctionDefinitions** in **Model**, as shown in Figure 10.

4.3.1 The id and name attributes

The **id** and **name** attributes have types **SId** and **string**, respectively, and operate in the manner described in Section 3.3. MathML **ci** elements in an SBML model can refer to the function defined by a **FunctionDefinition** using the value of its **id** attribute.

4.3.2 The math element

The **math** element is a container for MathML content that defines the function. The content of this element can only be a MathML **lambda** element or a MathML **semantics** element containing a **lambda** element. The **lambda** element must begin with zero or more **bvar** elements, followed by any other of the elements in the MathML subset listed in Section 3.4.1 *except* **lambda** (i.e., a **lambda** element cannot contain another **lambda** element). This is the only place in SBML where a **lambda** element can be used. The function defined by a **FunctionDefinition** is only available for use in other MathML elements that follow the **FunctionDefinition** definition in the model. (These restrictions prevent definition of recursive and mutually-recursive functions.)

A further restriction on the content of **math** is that it cannot contain references to variables other than the variables declared to the **lambda** itself. That is, the contents of MathML **ci** elements inside the body of the **lambda** can only be the variables declared by its **bvar** elements, or the identifiers of other **FunctionDefinitions** defined earlier in the model. This restriction also applies to the **csymbol** for *time*. Functions must be written so that all variables or parameters used in the MathML content are passed to them via their function parameters.

Note that **FunctionDefinition** does not have a separate attribute for defining the units of the value returned by the function. The units associated with the function's return value, when the function is called from within MathML expressions elsewhere in SBML, are simply the overall units of the expression in **FunctionDefinition**'s **math** when applied to the arguments supplied in the call to the function. Ascertaining these units requires performing dimensional analysis on the expression. (Readers may wonder why there is no attribute. The reason is that having a separate attribute for declaring the units would not only be redundant, but also lead to the potential for having conflicting information. In the case of a conflict between the declared units and those of the value actually returned by the function, the only logical resolution rule would be to assume that the correct units are those of the expression anyway.)

4.3.3 The sboTerm attribute

FunctionDefinition inherits an optional **sboTerm** attribute of type **SBOTerm** from its parent class **Sbase** (see Sections 3.1.9 and 5). When a value is given to this attribute in a **FunctionDefinition** instance, it must be an SBO identifier referring to a mathematical expression (i.e., terms derived from **SBO:0000064**, "mathematical expression"). The relationship is of the form "the function definition is a X", where X is the SBO term. The term chosen should be the most precise (narrow) one that captures the role of the function in the model.

As discussed in Section 5, SBO labels are optional information on a model. Applications are free to ignore **sboTerm** values. A model must be interpretable without the benefit of SBO labels.

4.3.4 Examples

The following abbreviated SBML example shows a **FunctionDefinition** object instance defining `pow3` as the identifier of a function computing the mathematical expression x^3 , and after that, the invocation of that function in the mathematical formula of a rate law. Note how the invocation of the function uses its identifier.

```
<model>
  ...
  <listOfFunctionDefinitions>
    <functionDefinition id="pow3">
      <math xmlns="http://www.w3.org/1998/Math/MathML">
        <lambda>
          <bvar><ci> x </ci></bvar>
          <apply>
            <power/>
            <ci> x </ci>
            <cn> 3 </cn>
          </apply>
        </lambda>
      </math>
    </functionDefinition>
  </listOfFunctionDefinitions>
  ...
  <listOfReactions>
    <reaction id="reaction_1">
      ...
      <kineticLaw>
        <math xmlns="http://www.w3.org/1998/Math/MathML">
          <apply>
            <ci> pow3 </ci>
            <ci> S1 </ci>
          </apply>
        </math>
      </kineticLaw>
      ...
    </reaction>
  </listOfReactions>
  ...
</model>
```

4.4 Unit definitions

Units of measurement may be supplied in a number of contexts in an SBML model. The units of the following mathematical entities can be specified explicitly: the size of a **Compartment**, the initial amount of a **Species**, and the units of constant and variable **Parameter** values. The overall units of any mathematical formula appearing in SBML are those that arise naturally from the components and mathematical expressions comprising the formula, or in other words, the units obtained by doing dimensional analysis on the formula.

Rather than requiring a complete unit definition on every **object**, SBML provides a facility for defining units that can be referenced throughout a model. In addition, every kind of SBML mathematical entity has units assigned to it from a set of built-in defaults (see Section 4.4.3 below, and also Sections 4.7.5, 4.8.5 and 4.13.5). By redefining these built-in default units, it is possible to change the units used throughout a model in a simple and consistent manner. The SBML unit definition facility uses two classes of objects, **UnitDefinition** and **Unit**. Their definitions are shown in Figure 12 on the following page and explained in more detail in Sections 4.4.1 and 4.4.2 below.

The approach to defining units in SBML is compositional; for example, $meter\ second^{-2}$ is constructed by combining a **Unit** object representing *meter* with another **Unit** object representing $second^{-2}$. The combination is wrapped inside a **UnitDefinition**, which provides for assigning an identifier and optional name to the combination. The identifier can then be referenced from elsewhere in a model.

The vast majority of modeling situations requiring new SBML unit definitions involve simple multiplicative combinations of base units and factors. An example of this might be “moles per litre per second”. What distinguishes these sorts of simpler unit definitions from more complex ones is that they may be expressed

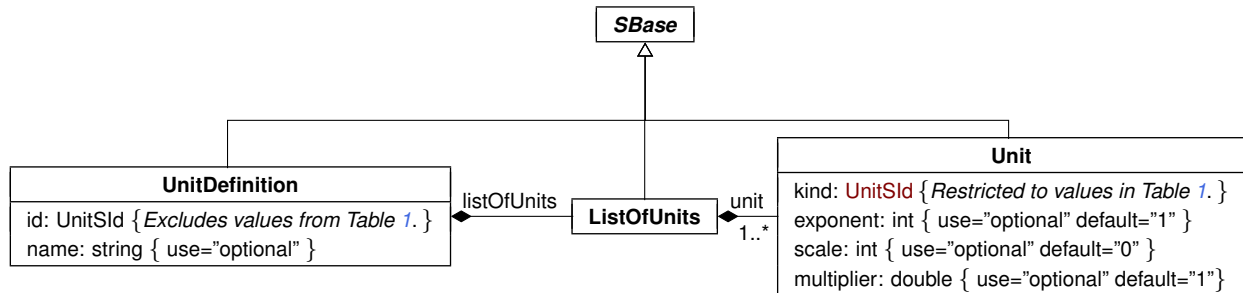


Figure 12: The definition of classes *UnitDefinition* and *Unit*. A sequence of one or more instances of *UnitDefinition* can be located in an instance of *ListOfUnitDefinitions* in *Model* (Figure 10). *ListOfUnits* has no attributes (beyond those it inherits from class *SBbase*); it merely acts as a container for one or more instances of *Unit* objects. Note that the only permitted values of *kind* on *Unit* are the reserved words in Table 1 on the following page, but these symbols are excluded from the permitted values of *UnitDefinition*'s *id* because SBML's unit system does not allow redefining the base units.

without the use of an additive offset from a zero point. The use of offsets complicates all unit definition systems, yet in the domain of SBML the real-life cases requiring offsets are few (and in fact, to the best of our knowledge, only involve temperature). Consequently, the SBML unit system has been consciously designed in a way that attempts to simplify implementation of unit support for the most common cases in systems biology, at the cost of requiring units with offsets to be handled explicitly by the modeler.

4.4.1 UnitDefinition

A unit definition in SBML consists of an instance of a *UnitDefinition* object, shown in Figure 12.

The id and name attributes

The required attribute *id* and optional attribute *name* have data types *UnitSId* and *string*, respectively. The *id* attribute is used to give the defined unit a unique identifier by which other parts of an SBML model definition can refer to it. The *name* attribute is intended to be used for giving the unit definition an optional human-readable name; see Section 3.3.2 for more guidelines about the use of names.

There are two important restrictions and guidelines about the use of unit definition *id* values:

1. The *id* of a *UnitDefinition* must *not* contain a value from Table 1, the list of reserved base unit names. This constraint simply prevents the redefinition of *base units*.
2. There is a set of predefined identifiers for the built-in default units in SBML; these identifiers are “*substance*”, “*volume*”, “*area*”, “*length*”, and “*time*”. Using one of these values for *id* in a *UnitDefinition* has the effect of redefining the model-wide default units for the corresponding quantities. We discuss this in more detail in Section 4.4.3.

The list of Units

A *UnitDefinition* object must contain one or more *Unit* objects inside a *ListOfUnits* container. Section 4.4.2 explains the meaning and use of *Unit*.

Example

The following skeleton of a unit definition illustrates an example use of *UnitDefinition*:

```

<model>
  <listOfUnitDefinitions>
    <unitDefinition id="unit1">
      <listOfUnits>
        ...
      </listOfUnits>
    </unitDefinition>
  </listOfUnitDefinitions>
</model>

```

```

1         </unitDefinition>
2         <unitDefinition id="unit2">
3             <listOfUnits>
4                 ...
5             </listOfUnits>
6         </unitDefinition>
7     </listOfUnitDefinitions>
8     ...
9 </model>

```

4.4.2 Unit

A **Unit** object represents a (possibly transformed) reference to a base unit chosen from the list in Table 1. The attribute **kind** indicates the chosen base unit, whereas the attributes **exponent**, **scale**, and **multiplier** define how the base unit is being transformed. These various attributes are described in detail below.

In SBML Level 2 Version 1, **Unit** had an additional attribute called **offset**. This attribute has been removed entirely in Level 2 Versions 2 and 3. Modelers and software authors are instead directed to use other methods of encoding units requiring offsets. The reasons for this change, and some suggestions for how to achieve equivalent effects of unit offsets, are discussed in more detail below. Another change in Version 3 is the removal of the enumeration **UnitKind** and the redefinition of **UnitSid** to include the previous **UnitKind** values as reserved symbols. This change has no net effect on permissible models or their representation.

The kind attribute

The **Unit** object class has one required attribute, **kind**, whose value must be taken from the list of reserved words given in Table 1. These reserved symbols are in the value space of **UnitSid** (Section 3.1.8).

| | | | | | |
|----------------------|-------------|----------|--------|-----------|-------|
| ampere | gram | katal | metre | second | watt |
| becquerel | gray | kelvin | mole | siemens | weber |
| candela | henry | kilogram | newton | sievert | |
| coulomb | hertz | litre | ohm | steradian | |
| <u>dimensionless</u> | <u>item</u> | lumen | pascal | tesla | |
| farad | joule | lux | radian | volt | |

Table 1: Base units defined in SBML. These symbols are predefined values of type **UnitSid** (Section 3.1.8). All are names of base or derived SI units (Bureau International des Poids et Mesures, 2000), except for “dimensionless” and “item”, which are SBML additions important for handling certain common situations. “Dimensionless” is intended for cases where a quantity is a ratio whose units cancel out, and “item” for expressing such things as “N items” (e.g., “100 molecules”). Also, note that the gram and litre are not strictly part of SI; however, they are frequently used in SBML’s areas of application and therefore are included as predefined unit identifiers. (The standard SI unit of mass is in fact the kilogram, and volume is defined in terms of cubic metres.) Comparisons of these values, like all values of type **UnitSid**, must be performed in a case-sensitive manner.

Note that the set of acceptable values for the attribute **kind** does *not* include units defined by **UnitDefinition** object. This means that the units definition system in SBML is not hierarchical: user-defined units cannot be built on top of other user-defined units, only on top of base units. SBML differs from CellML (Hedley et al., 2001) in this respect; CellML allows the construction of hierarchical unit definitions.

The exponent, scale and multiplier attributes

The optional **exponent** attribute on **Unit** represents an exponent on the unit. Its default value is “1” (one). A **Unit** object also has an optional **scale** attribute; its value must be an integer exponent for a power-of-ten multiplier used to set the scale of the unit. For example, a unit having a **kind** value of “gram” and a **scale** value of “-3” signifies $10^{-3} \times \text{gram}$, or milligrams. The default value of **scale** is “0” (zero), because $10^0 = 1$. Lastly, the optional **multiplier** attribute can be used to multiply the **kind** unit by a real-numbered factor; this enables the definition of units that are not power-of-ten multiples of SI units. For instance, a **multiplier** of 0.3048 could be used to define “foot” as a measure of length in terms of a metre. The **multiplier** attribute has a default value of “1” (one).

The unit system allows model quantities to be expressed in units other than the base units of Table 1. For analyses and computations, the consumer of the model (be it a software tool or a human) will want to convert all model quantities to base SI units for purposes such as verifying the consistency of units throughout the model. Suppose we begin with a quantity having numerical value y when expressed in units $\{u\}$. The relationship between y and a quantity y_b expressed in base units $\{u_b\}$ is

$$y_b \{u_b\} = y \{u\} \left(\frac{w \{u_b\}}{\{u\}} \right) \quad (1)$$

The term in the parentheses on the right-hand side is a factor w for converting a quantity in units $\{u\}$ to another quantity in units $\{u_b\}$. The ratio of units leads to canceling of $\{u\}$ in the equation above and leaves a quantity in units $\{u_b\}$. It remains to define this factor. In terms of the SBML unit system, it is:

$$\{u\} = (\text{multiplier} \cdot 10^{\text{scale}} \{u_b\})^{\text{exponent}} \quad (2)$$

where the dot (\cdot) represents simple scalar multiplication. The variables **multiplier**, **scale**, and **exponent** in the equation above correspond to the **attributes** with the same names in the **Unit object** defined in Figure 12. The exponent in the equation above may make it more difficult to grasp the relationship immediately; so let us suppose for the moment that **exponent**="1". Then, it is easy to see that

$$\{u\} = \text{multiplier} \cdot 10^{\text{scale}} \{u_b\}$$

Dividing both sides by $\{u\}$ produces the ratio in the parenthesized portion of Equation 1, which means that $w = \text{multiplier} \cdot 10^{\text{scale}}$. To take a concrete example, one foot expressed in terms of the metre (a base unit) requires **multiplier**="0.3048", **exponent**="1", and **scale**="0":

$$\text{foot} = 0.3048 \cdot 10^0 \cdot \text{metre}$$

leading to a conversion between quantities of

$$y_b \text{ metres} = 0.3048 \cdot y \text{ feet}$$

Given a quantity of, say, $y = 2$, the conversion results in $y_b = 0.6096$. To relate this to SBML terms more concretely, the following fragment of SBML illustrates how this is represented using the **Unit** and **UnitDefinition** constructs:

```
<listOfUnitDefinitions>
  <unitDefinition id="foot">
    <listOfUnits>
      <unit kind="metre" multiplier="0.3048"/>
    </listOfUnits>
  </unitDefinition>
</listOfUnitDefinitions>
```

The case above is the simplest possible situation, involving the transformation of quantities from a single defined unit $\{u\}$ into a quantity expressed in a single base unit $\{u_b\}$. If, instead, multiple base units $\{u_{b_1}\}, \{u_{b_2}\}, \dots, \{u_{b_n}\}$ are involved, the following equation holds (where the m_i terms are the **multiplier** values, the s_i terms are the **scale** values, and the x_i terms are the **exponent** values):

$$\begin{aligned} \{u\} &= (m_1 \cdot 10^{s_1} \{u_{b_1}\})^{x_1} \cdot (m_2 \cdot 10^{s_2} \{u_{b_2}\})^{x_2} \cdot \dots \cdot (m_n \cdot 10^{s_n} \{u_{b_n}\})^{x_n} \\ &= m_1^{x_1} \cdot m_2^{x_2} \cdot \dots \cdot m_n^{x_n} \cdot 10^{(s_1 x_1 + s_2 x_2 + \dots + s_n x_n)} \{u_{b_1}\}^{x_1} \{u_{b_2}\}^{x_2} \dots \{u_{b_n}\}^{x_n} \end{aligned} \quad (3)$$

Software developers should take care to track the exponents carefully because they can be negative integers. The overall use of Equation 3 is analogous to that of Equation 2, and leads to the following final expression. First, to simplify, let

$$\begin{aligned} m &= m_1^{x_1} \cdot m_2^{x_2} \cdot \dots \cdot m_n^{x_n} \\ p &= s_1 x_1 + s_2 x_2 + \dots + s_n x_n \end{aligned}$$

1 then,

$$2 \quad y_b \{u_{b_1}\} \{u_{b_2}\} \dots \{u_{b_n}\} = y \{u\} \left(\frac{m \cdot 10^p \{u_{b_1}\}^{x_1} \{u_{b_2}\}^{x_2} \dots \{u_{b_n}\}^{x_n}}{\{u\}} \right) \quad (4)$$

3
4 Some additional points are worth discussing about the unit scheme introduced so far. First, and most
5 importantly, the equations above are formulated with the assumption that the base units do not require
6 an additive offset as part of their definition. *When temperature values in units other than kelvin are being*
7 *considered, then a different interpretation must be made, as discussed below.*

8 A second point is that care is needed to avoid seemingly-obvious but incorrect translations of units described
9 in textbooks. The scheme above makes it easy to formulate statements such as “1 foot = 0.3048 metres”
10 in the most natural way. However, the most common expression of the relationship between temperature in
11 Fahrenheit and kelvin, “ $T_{Fahrenheit} = 1.8 \cdot (T_{kelvin} - 273.15) + 32$ ” might lead one to believe that defining
12 Fahrenheit degrees in terms of kelvin degrees involves using **multiplier**=“1.8”. *Not so*, when degree changes
13 are being considered and not temperature values. Converting *temperature values* is different from expressing
14 a relationship between degree measurements. The proper value for the multiplier in the latter case is 5/9,
15 i.e., **multiplier**=“0.555556” (where we picked an arbitrary decimal precision). If, on the other hand, the
16 actual temperature is relevant to a quantity (e.g., if a model uses a quantity that has particular values at
17 particular temperatures), then offsets are required in the unit calculations and a formula must be used as
18 discussed above.

19 *Handling units requiring the use of offsets in SBML Level 2 Version 3*

20 Unit definitions and conversions requiring offsets cannot be done using the simple approach above. The
21 most general case, involving offsets, multipliers and exponents, requires a completely different approach to
22 defining units than what has been presented up to this point.

23 In previous versions of SBML, not only was the general case incorrectly presented (i.e., in the same terms
24 described above, when in reality a different approach is required), but few if any developers even attempted
25 to support offset-based units in their software. In the development of SBML Level 2 Version 2, a consensus
26 among SBML developers emerged that a fully generalized unit scheme is so confusing and complicated that
27 it actually impedes interoperability. Level 2 Version 2 and Version 3 acknowledge this reality by reducing
28 and simplifying the unit system, specifically by removing the **offset attribute** on **Unit** and Celsius as a
29 pre-defined unit, and by describing approaches for handling Celsius and other temperature units. This is a
30 backwards-incompatible change relative to SBML Level 2 Version 1 and SBML Level 1 Version 2, but it is
31 believed to have limited real-life impact because so few tools and models appeared to have employed this
32 feature anyway. By simplifying the unit system to the point that it only involves multiplicative factors as
33 described above, we expect that more software tools will be able to support the SBML unit system from this
34 point forward, ultimately improving interoperability.

35 We first address the question of how to handle units that *do* require offsets:

- 36 • *Handling Celsius.* A model in which certain quantities are temperatures measured in degrees Celsius
37 can be converted straightforwardly to a model in which those temperatures are in kelvin. A software
38 tool could do this by performing a straightforward substitution using the following relationship:

$$39 \quad T_{kelvin} = T_{Celsius} + 273.15 \quad (5)$$

40 In every mathematical formula of the model where a quantity (call it x) in degrees Celsius appears,
41 replace x with $x_k + 273.15$ where x_k is now in kelvin. An alternative approach would be to use a
42 **FunctionDefinition** to define a function encapsulating this relationship above and then using that in the
43 rest of the model as needed. Since Celsius is a commonly-used unit, software tools could help users by
44 providing users with the ability to express temperatures in Celsius in the tools' interfaces, and making
45 substitutions automatically when writing out the SBML.

- 46 • *Handling other units requiring offsets.* The only other units requiring offsets in SBML's domain of
47 common applications are other temperature units such as Fahrenheit. Few modern scientists employ

Fahrenheit degrees; therefore, this is an unusual situation. The complication inherent in converting between degrees Fahrenheit and kelvin is that both a multiplier and an offset are required:

$$T_{kelvin} = \frac{T_F + 459.67}{1.8} \quad (6)$$

One approach to handling this is to use a **FunctionDefinition** to define a function encapsulating the relationship above, then to substitute a call to this function wherever the original temperature in Fahrenheit appears in the model's mathematical formulas. Here is a candidate definition as an example:

```
<functionDefinition id="Fahrenheit_to_kelvin">
  <math xmlns="http://www.w3.org/1998/Math/MathML">
    <lambda>
      <bvar><ci> temp_in_fahrenheit </ci></bvar>
      <apply>
        <divide/>
        <apply>
          <plus/>
          <ci> temp_in_fahrenheit </ci>
          <cn> 459.67 </cn>
        </apply>
        <cn> 1.8 </cn>
      </apply>
    </lambda>
  </math>
</functionDefinition>
```

An alternative approach not requiring the use of function definitions is to use an **AssignmentRule** for each variable in Fahrenheit units. The **AssignmentRule** could compute the conversion from Fahrenheit to (say) kelvin, assign its value to a variable (in Kelvin units), and then that variable could be used elsewhere in the model. Still another approach is to rewrite the mathematical formulas of a model to directly incorporate the conversion Equation 6 wherever the quantity appears.

All of these approaches provide general solutions to the problem of supporting any units requiring offsets in the unit system of SBML Level 2 Version 2 and Version 3. It can be used for other temperature units requiring an offset (e.g., degrees Rankine, degrees Réaumur), although the likelihood of a real-life model requiring such other temperature units seems exceedingly small.

In summary, **the removal of offset does not** impede the creation of models using alternative units. If conversions are needed, then converting between temperature in degrees Celsius and thermodynamic temperature can be handled rather easily by the simple substitution described above. For the rarer case of Fahrenheit and other units requiring combinations of multipliers and offsets, users are encouraged to employ the power of **FunctionDefinition**, **AssignmentRule**, or other constructs in SBML.

Examples

The following example illustrates the definition of an abbreviation “**mmols**” for the units $mmol\ l^{-1}\ s^{-1}$:

```
<listOfUnitDefinitions>
  <unitDefinition id="mmols">
    <listOfUnits>
      <unit kind="mole" scale="-3"/>
      <unit kind="litre" exponent="-1"/>
      <unit kind="second" exponent="-1"/>
    </listOfUnits>
  </unitDefinition>
</listOfUnitDefinitions>
```

4.4.3 Built-in units

There are five special unit identifiers in SBML, listed in Table 2 on the next page, corresponding to the five types of quantities that can play roles in SBML reactions: substance, volume, area, length and time. All

SBML mathematical entities apart from parameters have default units drawn from these predefined values. Table 2 lists the default values; all of the defaults have `multiplier="1"` and `scale="0"`.

| Identifier | Possible Scalable Units | Default Units |
|------------------------|---|---------------|
| <code>substance</code> | mole, item, gram, kilogram, dimensionless | mole |
| <code>volume</code> | litre, cubic metre, dimensionless | litre |
| <code>area</code> | square metre, dimensionless | square metre |
| <code>length</code> | metre, dimensionless | metre |
| <code>time</code> | second, dimensionless | second |

Table 2: SBML's built-in units. *The identifiers in the left-most column are values of `UnitSid` (Section 3.1.8).*

Redefinition of built-in units

Table 2 also lists alternative base units that are allowed as the basis of redefined values. For example, a redefinition of the built-in unit of time must be based on units of seconds. Within certain limits, a model may change the built-in units by reassigning the keywords `substance`, `length`, `area`, `time`, and `volume` in a `UnitDefinition`. The limitations on redefinitions of base units are the following:

1. The `UnitDefinition` of a redefined built-in unit can only contain a single `Unit` object within it.
2. The value of the `kind` attribute in a `Unit` instance must be drawn from one of the values in the second column of the appropriate row of Table 2.

Examples

The following example illustrates how to change the built-in units of volume to be millilitres. If this definition appeared in a model, the units of volume on all components that did not explicitly specify different units would be changed to millilitres.

```

<model>
  ...
  <listOfUnitDefinitions>
    <unitDefinition id="volume">
      <listOfUnits>
        <unit kind="litre" scale="-3"/>
      </listOfUnits>
    </unitDefinition>
  </listOfUnitDefinitions>
  ...
</model>

```

4.4.4 References to units

An `attribute` that defines the units of a mathematical entity (e.g., the `attribute` `units` on `Parameter`) can refer to a defined unit whose identifier is chosen from among the following:

- A new unit identifier defined by a `UnitDefinition` as described at the start of Section 4.4;
- The `base` units listed in Table 1 on page 37; and
- The `built-in` units defined in Section 4.4.3 and listed in Table 2. (These are “`substance`”, “`volume`”, “`area`”, “`length`”, and “`time`”.)

Software developers are asked to pay special attention to the units used in an SBML model. Different users and developers sometimes are accustomed to making different assumptions about units, and these assumptions may not correspond to what is actually defined in SBML. The numerical values in a model become meaningless if the corresponding units are not those being assumed. Sections 3.4.3, 4.8.5 and 4.13.5 have particularly important notes about the usage of units in SBML.

4.5 Compartment types

A *compartment type* in SBML is a grouping construct used to establish a relationship between multiple *compartments* (Section 4.7). A compartment type is represented by the **CompartmentType** object class, defined in Figure 13.

In **SBML Level 2 Version 3**, a compartment type only has an identity, and this identity can only be used to indicate that particular compartments belong to this type. This may be useful for conveying a modeling intention, such as when a model contains many similar compartments, either by their biological function or the reactions they carry; without a compartment type construct, it would be impossible in the language of SBML to indicate that all of the compartments share an underlying conceptual relationship because each SBML compartment must be given a unique and separate identity. Compartment types have no mathematical meaning in **SBML Level 2 Version 3**—they have no effect on a model’s mathematical interpretation. Simulators and other numerical analysis software may ignore **CompartmentType** objects and references to them in a model.

There is no mechanism in SBML for representing hierarchies of compartment types. One **CompartmentType** object instance cannot be the subtype of another **CompartmentType** object; SBML provides no means of defining such relationships.

4.5.1 The id and name attributes

As with other major class of objects in SBML, **CompartmentType** has a mandatory attribute, **id**, used to give the compartment type an identifier. The identifier must be a text string conforming to the syntax permitted by the **SId** data type described in Section 3.1.7. **CompartmentType** also has an optional **name** attribute, of type **string**. The **name** and **id** attributes must be used as described in Section 3.3.

4.5.2 The sboTerm attribute

CompartmentType inherits an optional **sboTerm** attribute of type **SBOTerm** from its parent class **Sbase** (see Sections 3.1.9 and 5). When a value is given to this attribute in a **CompartmentType** instance, the value must be an SBO identifier referring to a participant physical type (i.e., terms derived from **SBO:0000236**, “participant physical type”). The relationship is of the form “the compartment type *is a* X”, where X is the SBO term. The term chosen should be the most precise (narrow) one that captures the role of the compartment type in the model.

4.5.3 Examples

The following partial SBML example illustrates a compartment type used to relate together many individual compartments in a hypothetical model.

```
<model>
  ...
  <listOfCompartmentTypes>
    <compartmentType id="mitochondria"/>
  </listOfCompartmentTypes>
  <listOfCompartments>
    <compartment id="m1" size="0.013" compartmentType="mitochondria" outside="cell"/>
    <compartment id="m2" size="0.013" compartmentType="mitochondria" outside="cell"/>
    <compartment id="m3" size="0.013" compartmentType="mitochondria" outside="cell"/>
    <compartment id="cell" size="190.0"/>
  </listOfCompartments>
  ...
</model>
```

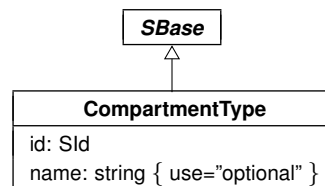


Figure 13: The definition of class **CompartmentType**. A sequence of one or more instances of **CompartmentType** objects can be located in an instance of **ListOfCompartmentTypes** in **Model**, as shown in Figure 10.

4.6 Species types

The term *species type* refers to reacting entities independent of location. These include simple ions (e.g., protons, calcium), simple molecules (e.g., glucose, ATP), large molecules (e.g., RNA, polysaccharides, and proteins), and others. The **SpeciesType** object class is intended to represent these entities. Its definition is shown in Figure 14.

SpeciesType objects are included in SBML to enable **Species** (Section 4.8) of the same type to be related together. It is a conceptual construct; the existence of **SpeciesType** objects in a model has no effect on the model's numerical interpretation. Except for the requirement for uniqueness of species/species type combinations located in compartments (described in Section 4.8.2), simulators and other numerical analysis software may ignore **SpeciesType** structures and references to them in a model.

There is no mechanism in SBML for representing hierarchies of species types. One **SpeciesType** object instance cannot be the subtype of another **SpeciesType**; SBML provides no means of defining such relationships.

4.6.1 The id and name attributes

As with other major object classes in SBML, **SpeciesType** has a mandatory attribute, **id**, used to give the species type an identifier. The identifier must be a text string conforming to the syntax permitted by the **SId** data type described in Section 3.1.7. **SpeciesType** also has an optional **name** attribute, of type **string**. The **name** and **id** attributes must be used as described in Section 3.3.

4.6.2 The sboTerm attribute

SpeciesType inherits an optional **sboTerm** attribute of type **SBOTerm** from its parent class **Sbase** (see Sections 3.1.9 and 5). When a value is given to this attribute in a **SpeciesType** instance, the value must be an SBO identifier referring to a participant physical type (i.e., terms derived from **SBO:0000236**, "participant physical type"). The relationship is of the form "the species type *is a* X", where X is the SBO term. The term chosen should be the most precise (narrow) one that captures the role of the compartment type in the model.

4.6.3 Example

The following XML fragment is an example of two **SpeciesType** objects embedded in an SBML model.

```
<model>
  ...
  <listOfSpeciesTypes>
    <speciesType id="ATP"/>
  </listOfSpeciesTypes>
  <listOfCompartments>
    <compartment id="cytosol"/>
    <compartment id="mitochon"/>
  </listOfCopartments>
  <listOfSpecies>
    <species id="ATPc" speciesType="ATP" compartment="cytosol" initialConcentration="1"/>
    <species id="ATPm" speciesType="ATP" compartment="mitochon" initialConcentration="2"/>
  </listOfSpecies>
  ...
</model>
```

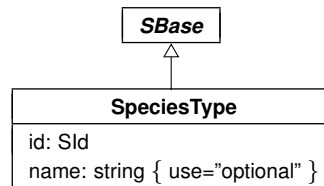


Figure 14: The definition of class **SpeciesType**. A sequence of one or more instances of **SpeciesType** objects can be located in an instance of **ListOfSpeciesTypes** in **Model**, as shown in Figure 10.

4.7 Compartments

A *compartment* in SBML represents a bounded space in which species are located. Compartments do not necessarily have to correspond to actual structures inside or outside of a biological cell, although models are often designed that way. The definition of **Compartment** is shown in Figure 15.

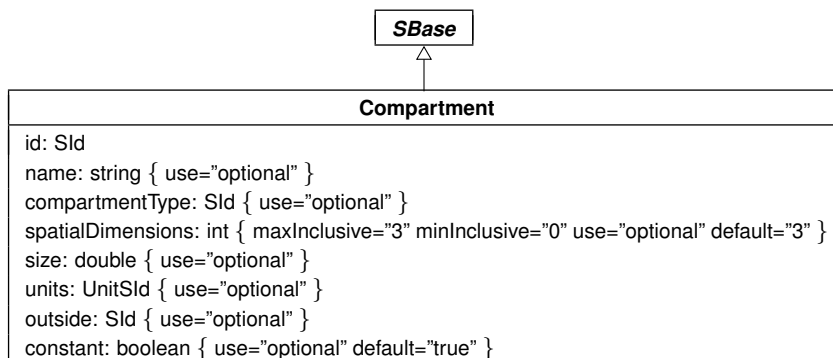


Figure 15: The definition of class **Compartment**. A sequence of one or more instances of **Compartment** objects can be located in an instance of **ListOfCompartments** in **Model**, as shown in Figure 10.

It is important to note that although compartments are optional in the overall definition of **Model** (see Section 4.2), every species in an SBML model must be located in a compartment. This in turn means that if a model defines any species, the model must also define at least one compartment. The reason is simply that species represent physical things, and therefore must exist *somewhere*. Compartments represent the *somewhere*.

4.7.1 The **id** and **name** attributes

Compartment has one required attribute, **id**, of type **SId**, to give the compartment a unique identifier by which other parts of an SBML model definition can refer to it. A compartment can also have an optional **name** attribute of type **string**. Identifiers and names must be used according to the guidelines described in Section 3.3.

4.7.2 The **compartmentType** attribute

Each compartment in a model may optionally be designated as belonging to a particular compartment type. The optional attribute **compartmentType** of type **SId** is used identify the compartment type represented by the **Compartment** object. The **compartmentType** attribute's value must be the identifier of a **CompartmentType** instance defined in the model. If the **compartmentType** attribute is not present on a particular compartment definition, a unique virtual compartment type is assumed for that compartment, and no other compartment can belong to that compartment type.

The values of **compartmentType** attributes on compartments have no effect on the numerical interpretation of a model. Simulators and other numerical analysis software may ignore **compartmentType** attributes.

4.7.3 The **spatialDimensions** attribute

A **Compartment** object has an optional attribute **spatialDimensions**, whose value must be a positive integer indicating the number of spatial dimensions possessed by the compartment. The maximum value is "3", meaning a three-dimensional structure (a volume). Other permissible values are "2" (for a two-dimensional area), "1" (for a one-dimensional curve), and "0" (for a point). The default value is "3". Note that the number of spatial dimensions possessed by a compartment affects certain aspects of the compartment's size and units-of-size; see the following two subsections.

4.7.4 The size attribute

Each compartment has an optional floating-point attribute named **size**, representing the initial total size of the compartment. The size may be a volume (if the compartment is a three-dimensional one), or it may be an area (if the compartment is two-dimensional), or a length (if the compartment is one-dimensional).

It is important to note that in SBML Level 2, a missing **size** value *does not imply that the compartment size is 1*. There is no default value of compartment size. (This is unlike the definition of compartment **volume** in SBML Level 1.) When the **spatialDimensions** attribute does not have a value of “0”, a missing value for **size** for a given compartment signifies that the value either is unknown, or to be obtained from an external source, or determined by an initial assignment (Section 4.10) or a rule (Section 4.11) elsewhere in the model. The **size** attribute must not be present if the **spatialDimensions** attribute has a value of “0”; otherwise, a logical inconsistency would exist because a zero-dimensional object cannot have a physical size.

A compartment’s size is set by its **size** attribute exactly once. If the compartment’s **constant** attribute value is “true” (the default), then the size is fixed and cannot be changed except by an **InitialAssignment** in the model (and if **spatialDimensions**=“0”, it cannot be changed by any **InitialAssignment** either). These methods of setting the size differ in that the **size** attribute can only be used to set the compartment size to a literal scalar value, whereas **InitialAssignment** allows the value to be set using an arbitrary mathematical expression. If the compartment’s **constant** attribute is “false”, the size value may be overridden by an **InitialAssignment** or changed by an **AssignmentRule** or **AlgebraicRule**, and in addition, for simulation time $t > 0$, it may also be changed by a **RateRule** or **Events**. (However, some constructs are mutually exclusive; see Sections 4.11 and 4.14.) It is not an error to define **size** on a compartment and also redefine the value using an **InitialAssignment**, but the **size** value in that case is ignored. Section 3.4.8 provides additional information about the semantics of assignments, rules and values for simulation time $t \leq 0$.

For the reasons given above, the **size** attribute on a compartment must be defined as optional; however, *it is extremely good practice to specify values for compartment sizes* when such values are available. There are two major technical reasons for this. First, models ideally should be instantiable in a variety of simulation frameworks. A commonly-used one is the discrete stochastic framework (Gillespie, 1977; Wilkinson, 2006) in which species are represented as item counts (e.g., molecule counts). If species’ initial quantities are given in terms of concentrations or densities, it is impossible to convert the values to item counts without knowing compartment sizes. Second, and more importantly, if a model contains multiple compartments whose sizes are not all identical to each other, it is impossible to quantify the reaction rate expressions without knowing the compartment volumes. The reason for the latter is that reaction rates in SBML are defined in terms of *substance/time* (see Section 4.13.5), and when species quantities are given in terms of concentrations or densities, the compartment sizes become factors in the reaction rate expressions.

4.7.5 The units attribute

The units associated with the compartment’s **size** value may be set using the optional attribute **units**. The default units, and the kinds of units allowed as values of the **units** attribute, interact with the number of spatial dimensions of the compartment. The value of the **units** attribute of a **Compartment** object must be one of the base units from Table 1, or the built-in units “volume”, “area”, “length” or “dimensionless”, or a new unit defined by a unit definition in the enclosing model, subject to the restrictions detailed in Table 3.

| Value of attribute spatialDimensions | size allowed? | units allowed? | Allowable kinds of units | Default value of attribute units |
|--|----------------------|-----------------------|-----------------------------------|--|
| “3” | yes | yes | units of volume, or dimensionless | “volume” |
| “2” | yes | yes | units of area, or dimensionless | “area” |
| “1” | yes | yes | units of length, or dimensionless | “length” |
| “0” | no | no | (no units allowed) | |

Table 3: The units permitted for compartment sizes. If **spatialDimensions**=“0”, the compartment’s **units** attribute must be left unset. Units of volume means litres, cubic metres, or units derived from them; units of area means square metres or units derived from square metres; and units of length means metres or units derived from metres. (See also Table 2 on page 41 and Table 1 on page 37.)

1 The units of the compartment size, as defined by the `units` attribute or (if `units` is not set) the default
2 value listed in Table 3 on the preceding page, are used in the following ways when the compartment has a
3 `spatialDimensions` value greater than “0”:

- 4 • The value of the `units` attribute is used as the units of the compartment identifier when the identi-
5 fier appears as a numerical quantity in a mathematical formula expressed in MathML (discussed in
6 Section 3.4.3).
- 7 • The `math` element of an `AssignmentRule` or `InitialAssignment` referring to this compartment must have
8 identical units (see Sections 4.11.3 and 4.10).
- 9 • In `RateRule` objects that set the rate of change of the compartment’s size (Section 4.11.4), the units of
10 the rule’s `math` element must be identical to the compartment’s `units` attribute divided by the default
11 *time* units. (In other words, the units for the rate of change of compartment size are *compartment*
12 *size/time* units.)
- 13 • When a `Species` is to be treated in terms of concentrations or density, the units of the spatial size
14 portion of the concentration value (i.e., the denominator in the units formula *substance/size*) are those
15 indicated by the value of the `units` attribute on the compartment in which the species is located.

16 Compartments with `spatialDimensions`=“0” require special treatment in this framework. If a compartment
17 has no size or dimensional units, how should such a compartment’s identifier be interpreted when it appears
18 in mathematical formulas? The answer is that such a compartment’s identifier should not appear in mathe-
19 matical formulas in the first place—it has no value, and its value cannot change (Section 4.7.6). Note also
20 that a zero-dimensional compartment is a point, and species located at points can only be described in terms
21 of amounts, not spatially-dependent measures such as concentration. Since SBML `KineticLaw` formulas are
22 already in terms of *substance/time* and not (say) *concentration/time*, volume or other factors in principle
23 are not needed for species located in zero-dimensional compartments.

24 4.7.6 The constant attribute

25 A `Compartment` also has an optional boolean attribute called `constant` that indicates whether the com-
26 partment’s size stays constant or can vary during a simulation. A value of “false” indicates the com-
27 partment’s size can be changed by other constructs in SBML. A value of “true” indicates the compart-
28 ment’s size cannot be changed by any other construct except `InitialAssignment`. In the special case of
29 `spatialDimensions`=“0”, the value cannot be changed by `InitialAssignment` either. The default value for the
30 `constant` attribute is “true” because in the most common modeling scenarios at the time of this writing,
31 compartment sizes remain constant. The `constant` attribute must default to or be set to “true” if the value
32 of the `spatialDimensions` attribute is “0”, because a zero-dimensional compartment cannot ever have a
33 size.

34 4.7.7 The outside attribute

35 The optional attribute `outside` of type `SId` can be used to express containment relationships between com-
36 partments. If present, the value of `outside` for a given compartment must be the `id` attribute value of
37 another compartment already defined in the model. Doing so means that the other compartment contains
38 it or is “outside” of it. This enables the representation of simple topological relationships between compart-
39 ments, for those simulation systems that can make use of the information (e.g., for drawing simple diagrams
40 of compartments).

41 There are two restrictions on the containment relationships in SBML. First, because a compartment with
42 `spatialDimensions` of “0” has no size, such a compartment cannot act as the container of any other com-
43 partment *except* compartments that *also* have `spatialDimensions` values of “0”. Second, the directed graph
44 formed by representing `Compartment` objects as vertexes and the `outside` attribute values as edges must be
45 acyclic. The latter condition is imposed to prevent a compartment from being contained inside itself.

46 Although containment relationships are partly taken into account by the compartmental localization of
47 reactants and products, it is not always possible to determine purely from the reaction equations whether

one compartment is meant to be located within another. In the absence of a value for **outside**, compartment definitions in SBML Level 2 do not have any implied spatial relationships between each other. For many modeling applications, the transfer of substances described by the reactions in a model sufficiently express the relationships between the compartments. (As discussed in Section 8.1, SBML Level 3 is expected to introduce the ability to define geometries and spatial qualities.)

4.7.8 The **sboTerm** attribute

Compartment inherits an optional **sboTerm** attribute of type **SBOTerm** from its parent class **Sbase** (see Sections 3.1.9 and 5). When a value is given to this attribute in a **Compartment** instance, the value must be an SBO identifier referring to a participant physical type (i.e., terms derived from **SBO:0000236**, “participant physical type”). The relationship is of the form “the compartment *is a* X”, where X is the SBO term. The term chosen should be the most precise (narrow) one that captures the role of the compartment type in the model.

4.7.9 Examples

The following example illustrates two compartments in an abbreviated SBML example of a model definition:

```
<model>
  ...
  <listOfCompartments>
    <compartment id="cytosol" size="2.5"/>
    <compartment id="mitochondria" size="0.3"/>
  </listOfCompartments>
  ...
</model>
```

The following is an example of using **outside** to model a cell membrane. To express that a compartment with identifier “B” has a membrane that is modeled as another compartment “M”, which in turn is located within another compartment “A”, one would write:

```
<model>
  ...
  <listOfCompartments>
    <compartment id="A"/>
    <compartment id="M" spatialDimensions="2" outside="A"/>
    <compartment id="B" outside="M"/>
  </listOfCompartments>
  ...
</model>
```

4.8 Species

A *species* refers to a pool of reacting entities of a specific *species type* that take part in reactions and are located in a specific *compartment*. The **Species** object class is intended to represent these pools. Its definition is shown in Figure 16 on the next page. Although the exact definition of **Species** given here has changed from the definition in the specification of SBML Level 2 Version 1 (i.e., through the introduction of **SpeciesType**), the concept represented by **Species** remains the same.

In previous versions of SBML Level 2, the class **Species** included an attribute called **spatialSizeUnits**, which allowed explicitly setting the units of *size* for initial concentration. SBML Level 2 Version 3 removes this attribute for two reasons. First, its semantics were confusing and introduced an implicit unit conversion. **Compartment** has its own size-setting attribute, and a separate attribute on **Species** meant that a species’ initial concentration could involve spatial size units that are different from the units of the compartment in which the species was located. Since (a) the **spatialSizeUnits** attribute determined the size units of a species’ quantity when that species appeared in reaction rate formulas, and (b) the compartment may have its own different spatial units, when both the species and the compartment appeared in reaction rate formulas, one of the quantities would have had to be converted into the same spatial units as the other. In other words, modelers and software tools would have had to insert explicit conversion factors into kinetic

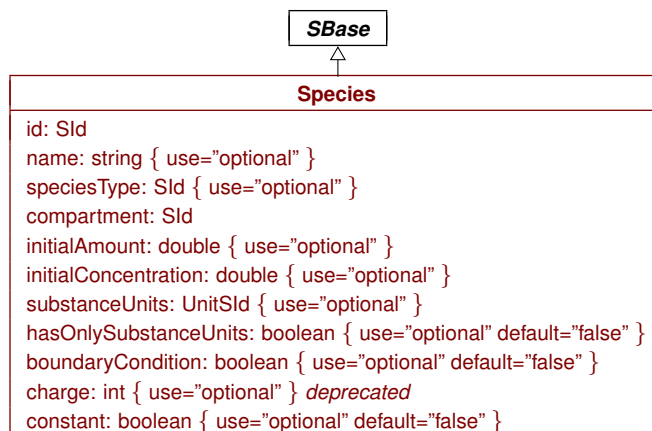


Figure 16: The definition of class **Species**. A sequence of one or more instances of **Species** objects can be located in an instance of **ListOfSpecies** in **Model**, as shown in Figure 10.

rate formulas. Such conversions can be difficult to achieve and prone to errors. Second, although the **spatialSizeUnits** attribute provided some new functionality, it could also be argued to be redundant: a compartment's definition logically *should* be the entity controlling its own spatial size information, in the same way that a species' definition controls its quantity. For these reasons, and because few software packages seemed to take account of the implicit unit conversion, it was deemed better to remove the **spatialSizeUnits** attribute from **Species**.

4.8.1 The **id** and **name** attributes

As with other major objects in SBML, **Species** has a mandatory attribute, **id**, used to give the species an identifier. The identifier must be a text string conforming to the syntax permitted by the **SId** data type described in Section 3.1.7. **Species** also has an optional **name** attribute, of type **string**. The **name** and **id** attributes must be used as described in Section 3.3.

4.8.2 The **speciesType** attribute

Each species in a model may optionally be designated as belonging to a particular species type. The optional attribute **speciesType** of type **SId** is used to identify the species type of the chemical entities that make up the pool represented by the **Species** object. The attribute's value must be the identifier of an existing **SpeciesType** object. If the **speciesType** attribute is not present on a particular species definition, it means the pool contains chemical entities of a type unique to that pool; in effect, a virtual species type is assumed for that species, and no other species can belong to that species type.

There can be only one species of a given species type in any given compartment of a model. More specifically, for all **Species** objects having a value for the **speciesType** attribute, the pair

(**speciesType** attribute value, **compartment** attribute value)

must be unique across the set of all **Species** objects in a model.

The value of **speciesType** attributes on species have no effect on the numerical interpretation of a model. Simulators and other numerical analysis software may ignore **speciesType** attributes.

4.8.3 The **compartment** attribute

The required attribute **compartment**, also of type **SId**, is used to identify the compartment in which the species is located. The attribute's value must be the identifier of an existing **Compartment** object. It is important to note that there is no default value for the **compartment** attribute on **Species**; every species in an SBML model must be assigned a compartment, and consequently, a model must define at least one compartment if that model contains any species.

4.8.4 The initialAmount and initialConcentration attributes

The optional `attributes` `initialAmount` and `initialConcentration`, both having a data type of `double`, are used to set the initial quantity of the species in the compartment where the species is located. These `attributes` are mutually exclusive; i.e., *only one* can have a value on any given instance of a `Species` object.

Missing `initialAmount` and `initialConcentration` values implies that their values either are unknown, or to be obtained from an external source, or determined by an initial assignment (Section 4.10) or rule (Section 4.11) elsewhere in the model. In the case where a species' compartment has a `spatialDimensions` value of "0", the species cannot have a value for `initialConcentration` because the concepts of "concentration" and "density" break down when a container has zero dimensions.

A species' initial quantity is set by the `initialAmount` or `initialConcentration` `attributes` exactly once. If the species' `constant` `attribute` is "true" (the default), then the size is fixed and cannot be changed except by an `InitialAssignment`. These methods differ in that the `initialAmount` and `initialConcentration` `attributes` can only be used to set the species quantity to a literal scalar value, whereas `InitialAssignment` allows the value to be set using an arbitrary mathematical expression. If the species' `constant` `attribute` is "false", the species' quantity value may be overridden by an `InitialAssignment` or changed by `AssignmentRule` or `AlgebraicRule`, and in addition, for $t > 0$, it may also be changed by a `RateRule` or `Events`. (However, some constructs are mutually exclusive; see Sections 4.11 and 4.14.) It is not an error to define `initialAmount` or `initialConcentration` on a species and also redefine the value using an `InitialAssignment`, but the `initialAmount` or `initialConcentration` setting in that case is ignored. Section 3.4.8 provides additional information about the semantics of assignments, rules and values for simulation time $t \leq 0$.

The units of the value in the `initialAmount` `attribute` are set by the `substanceUnits` `attribute` on `Species`. The units of the value in the `initialConcentration` `attribute` are *substance/size* units. The units of *substance* are those defined in the `substanceUnits`, and the *size* units are those given in the definition of the size of the `Compartment` in which the species is located.

4.8.5 The substanceUnits and hasOnlySubstanceUnits attributes

The units associated with a species' quantity, referred to as the *units of the species*, are determined via the optional `attributes` `substanceUnits` and `hasOnlySubstanceUnits`, in combination with the units of the size defined for the compartment object in which the species are located.

The `attribute` `hasOnlySubstanceUnits` takes on boolean values and defaults to "false". This `attribute`'s role is to indicate whether the units of the species, when the species identifier appears in mathematical formulas, are intended to be concentration or amount. Although it may seem as though this intention could be determined based on whether `initialConcentration` or `initialAmount` is set, the fact that these two `attributes` are optional means that a separate flag is needed. (Consider the situation where neither is set, and instead the species' quantity is established by an `InitialAssignment` or `AssignmentRule`.)

The possible values of *units of the species* are summarized in Table 4. The *units of the species* are of the form *substance/size* units (i.e., *concentration* units, using a broad definition of concentration) if the compartment's `spatialDimensions` is non-zero and `hasOnlySubstanceUnits` has the value "false". The *units of the species* are of the form *substance* if `hasOnlySubstanceUnits` has the value "true" or `spatialDimensions` is zero. (This dependence is due to the fact that a zero-dimensional compartment cannot support concentrations or densities.) The units of *substance* are those defined by the `substanceUnits` `attribute`, and the *size* units are those of the size of the compartment in which the species is located. This compartment is the one identified by the species' `compartment` `attribute`.

| value of <code>hasOnlySubstanceUnits</code> | <i>units of the species</i> when <code>spatialDimensions</code> is greater than 0 | <i>units of the species</i> when <code>spatialDimensions</code> is 0 |
|---|---|--|
| false (default) | <i>substance/size</i> | <i>substance</i> |
| true | <i>substance</i> | <i>substance</i> |

Table 4: How to interpret the value the `Species` `hasOnlySubstanceUnits` `attribute`.

As an aside, we note that treating species in terms of *substance* units (i.e., discrete quantities such as molecule counts) rather than concentrations is common when using discrete stochastic simulation frameworks (Gillespie, 1977; Wilkinson, 2006). The appropriate way of accomplishing this in SBML is to set `hasOnlySubstanceUnits="true"` in the species' definitions.

The value chosen for `substanceUnits` must be either a base unit from Table 1 on page 37, a built-in unit from Table 2 on page 41, or a new unit defined by a unit definition in the enclosing model. The chosen units for `substanceUnits` must be dimensionless, mole, item, kilogram, gram, or units derived from these. The `substanceUnits` attribute defaults to the built-in unit "substance" shown in Table 2 on page 41.

The *units of the species* are used in the following ways:

- The species identifier has these units when the identifier appears as a numerical quantity in a mathematical formula expressed in MathML (discussed in Section 3.4.3).
- The `math` element of an `AssignmentRule` or `InitialAssignment` referring to this species must have identical units (see Sections 4.11.3 and 4.10).
- In `RateRule` objects that set the rate of change of the species' quantity (Section 4.11.4), the units of the rule's `math` element must be identical to the *units of the species* divided by the model's *time* units.

4.8.6 The constant and boundaryCondition attributes

The `Species` object has two optional boolean attributes named `constant` and `boundaryCondition`, used to indicate whether and how the quantity of that species can vary during a simulation. Table 5 shows how to interpret the combined values of the `boundaryCondition` and `constant` attributes.

| constant value | boundaryCondition value | can have assignment or rate rule | can be reactant or product | species' quantity can be changed by |
|----------------|-------------------------|----------------------------------|----------------------------|--|
| true | true | no | yes | (never changes) |
| false | true | yes | yes | rules and events |
| true | false | no | no | (never changes) |
| false | false | yes | yes | reactions <i>or</i> rules (but not both), and events |

Table 5: How to interpret the values of the `constant` and `boundaryCondition` attributes on `Species`. Note that column four is specifically about reactants and products and not also about species acting as modifiers; the latter are by definition unchanged by reactions.

By default, when a species is a product or reactant of one or more reactions, its quantity is determined by those reactions. In SBML, it is possible to indicate that a given species' quantity is *not* determined by the set of reactions even when that species occurs as a product or reactant; i.e., the species is on the *boundary* of the reaction system, and its quantity is not determined by the reactions. The boolean attribute `boundaryCondition` can be used to indicate this. The value of the attribute defaults to "false", indicating the species *is* part of the reaction system.

The `constant` attribute indicates whether the species' quantity can be changed at all, regardless of whether by reactions, rules, or constructs other than `InitialAssignment`. The default value is "false", indicating that the species' quantity can be changed, since the purpose of most simulations is precisely to calculate changes in species quantities. Note that the initial quantity of a species can be set by an `InitialAssignment` irrespective of the value of the `constant` attribute.

In practice, a `boundaryCondition` value of "true" means a differential equation derived from the reaction definitions should not be generated for the species. However, the species' quantity may still be changed by `AssignmentRule`, `RateRule`, `AlgebraicRule`, `Event`, and `InitialAssignment` constructs if its `constant` attribute is "false". Conversely, if the species' `constant` attribute is "true", then its value cannot be changed by anything except `InitialAssignment`.

1 A species having `boundaryCondition="false"` and `constant="false"` can appear as a product and/or
 2 reactant of one or more reactions in the model. If the species is a reactant or product of a reaction, it must
 3 not also appear as the target of any [AssignmentRule](#) or [RateRule](#) object in the model. If instead the species
 4 has `boundaryCondition="false"` and `constant="true"`, then it cannot appear as a reactant or product,
 5 or as the target of any [AssignmentRule](#), [RateRule](#) or [EventAssignment](#) object in the model.

6 The example model in section 7.6 contains all four possible combinations of the `boundaryCondition` and
 7 `constant` attributes on `species` elements. Section 7.7 gives an example of how one can translate into ODEs
 8 a model that uses `boundaryCondition` and `constant` attributes.

9 4.8.7 The charge attribute

10 The optional attribute `charge` takes an integer indicating the charge on the species (in terms of electrons,
 11 not the SI unit coulombs). This may be useful when the species is a charged ion such as calcium (Ca^{2+}).
 12 The charge attribute in SBML has been deprecated since Level 2 Version 2.

13 4.8.8 The sboTerm attribute

14 [Species](#) inherits an optional `sboTerm` attribute of type [SBOTerm](#) from its parent class [Sbase](#) (see Sections 3.1.9
 15 and 5). When a value is given to this attribute in a [Species](#) instance, the value must be an SBO identifier
 16 referring to a participant physical type (i.e., terms derived from [SBO:0000236](#), “participant physical type”).
 17 The relationship is of the form “the species *is a* X”, where X is the SBO term. The term chosen should be
 18 the most precise (narrow) one that captures the role of the compartment type in the model.

19 4.8.9 Example

20 The following example shows two species definitions within an abbreviated SBML model definition. The
 21 example shows that species are listed under the heading `listOfSpecies` in the model:

```
22 <model>
23   ...
24   <listOfSpecies>
25     <species id="Glucose" compartment="cell" initialConcentration="4"/>
26     <species id="Glucose_6_P" compartment="cell" initialConcentration="0.75"/>
27   </listOfSpecies>
28   ...
29 </model>
```

30 4.9 Parameters

31 A [Parameter](#) is used in SBML to define a symbol associated with a value; this symbol can then be used
 32 in mathematical formulas in a model. By default, parameters have constant value for the duration of a
 33 simulation, and for this reason are called “parameters” instead of “variables” in SBML, although in truth,
 34 SBML parameters can be either. The definition of [Parameter](#) is shown in Figure 17.

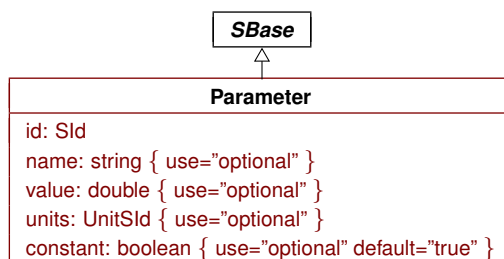


Figure 17: The definition of class [Parameter](#). A sequence of one or more instances of [Parameter](#) objects can be located in an instance of [ListOfParameters](#) in [Model](#), as shown in Figure 10.

Parameters can be defined in two places in SBML: in lists of parameters defined at the top level in a **Model instance**, and within individual reaction definitions (as described in Section 4.13). Parameters defined at the top level are *global* to the whole model; parameters that are defined within a reaction are local to the particular reaction and (within that reaction) *override* any global parameters having the same identifiers (See Section 3.3.1 for further details).

The use of the term *parameter* in SBML sometimes leads to confusion among readers who have a particular notion of what something called “parameter” should be. It has been the source of heated debate, but despite this, no one has yet found an adequate replacement term that does not have different connotations to different people and hence leads to confusion among *some* subset of users. Perhaps it would have been better to have two constructs, one called “constants” and the other called “variables”. The current approach in SBML is simply more parsimonious, using a single **Parameter** construct with the boolean flag **constant** indicating which flavor it is. In any case, readers are implored to look past their particular definition of a “parameter” and simply view SBML’s **Parameter** as a single mechanism for defining both constants and (additional) variables in a model. (We write *additional* because the species in a model are usually considered to be the central variables.) After all, software tools are not required to expose to users the actual names of particular SBML constructs, and thus tools can present to their users whatever terms their designers feel best matches their target audience.

4.9.1 The id and name attributes

Parameter has one required **attribute**, **id**, of type **SId**, to give the parameter a unique identifier by which other parts of an SBML model definition can refer to it. A parameter can also have an optional **name attribute** of type **string**. Identifiers and names must be used according to the guidelines described in Section 3.3.

4.9.2 The value attribute

The optional **attribute value** determines the value (of type **double**) assigned to the identifier. A missing **value** implies that the **value** either is unknown, or to be obtained from an external source, or determined by an initial assignment (Section 4.10) or a rule (Section 4.11) elsewhere in the model.

A parameter’s value is set by its **value attribute** exactly once. If the parameter’s **constant attribute** has the value “**true**” (the default), then the value is fixed and cannot be changed except by an **InitialAssignment**. These two methods of setting the parameter’s value differ in that the **value attribute** can only be used to set it to a literal scalar value, whereas **InitialAssignment** allows the value to be set using an arbitrary mathematical expression. If the parameter’s **constant attribute** has the value “**false**”, the parameter’s value may be overridden by an **InitialAssignment** or changed by **AssignmentRule** or **AlgebraicRule**, and in addition, for simulation time $t > 0$, it may also be changed by a **RateRule** or **Events**. (However, some of these constructs are mutually exclusive; see Sections 4.11 and 4.14.) It is not an error to define **value** on a parameter and also redefine the value using an **InitialAssignment**, but the **value** in that case is ignored. Section 3.4.8 provides additional information about the semantics of assignments, rules and values for simulation time $t \leq 0$.

4.9.3 The units attribute

The units associated with the value of the parameter are specified by the attribute **units**. The value assigned to the parameter’s **units attribute** must be chosen from one of the following possibilities: one of the base unit identifiers from Table 1 on page 37; one of the built-in unit identifiers appearing in first column of Table 2 on page 41; or the identifier of a new unit defined in the list of unit definitions in the enclosing **Model** definition. There are no constraints on the units that can be chosen from these sets. There are no default units for parameters.

The units of the parameter are used in the following ways:

- When the parameter identifier appears in mathematical formulas expressed in MathML in a model, the units associated with the value are those declared by the parameter’s **units attribute**.
- The units of the **math** element of an **AssignmentRule**, **InitialAssignment** or **EventAssignment** setting the value of the parameter must have units identical to those declared by the parameter’s **units attribute**.

- The units of the `math` element of a `RateRule` that references the parameter be identical to *parameter units/time*, where *parameter units* are the units declared for the parameter using the `units` attribute and *time* is the model-wide `time` units.

The fact that parameter units are optional, and that no defaults are defined, means that models can define parameters with undeclared units. If such parameters appear in mathematical expressions elsewhere in a model, it may not be possible for a software tool to verify the consistency of units used in the expressions. Modelers and software tools should therefore assign units to parameters whenever possible.

4.9.4 The constant attribute

The `Parameter` object has an optional boolean attribute named `constant` which indicates whether the parameter's value can vary during a simulation. The attribute's default value is "true". A value of "false" indicates the parameter's value can be changed by rules (see Section 4.11) and that the value is actually intended to be the initial value of the parameter.

Parameters local to a reaction (i.e., those defined within a `Reaction`'s `KineticLaw` object, as described in Section 4.13.5) cannot be changed by rules and therefore are implicitly always constant; thus, parameter definitions within `Reaction` objects should *not* have their `constant` attribute set to "false".

What if a global parameter has its `constant` attribute set to "false", but the model does not contain any rules, events or other constructs that ever change its value over time? Although the model may be suspect, this situation is not strictly an error. A value of "false" for `constant` only indicates that a parameter *can* change value, not that it *must*.

4.9.5 The sboTerm attribute

The `Parameter` object inherits from `Sbase` the optional `sboTerm` attribute of type `SBOTerm` (see Sections 3.1.9 and 5). When a value is given to this attribute in a parameter definition, the value must be an SBO identifier referring to a quantitative parameter defined in SBO (i.e., terms derived from SBO:0000002, "quantitative parameter"). The relationship is of the form "the SBML parameter *is a* X", where X is the SBO term. The term chosen should be the most precise (narrow) one that captures the role of the parameter in the model.

As discussed in Section 5, SBO labels are optional information on a model. Applications are free to ignore `sboTerm` values. A model must be interpretable without the benefit of SBO labels.

4.9.6 Example

The following is an example of parameters defined at the `Model` level:

```

<model>
  ...
  <listOfParameters>
    <parameter id="tau2" value="3e-2" units="second"/>
    <parameter id="Km1" value="10.7" units="moleperlitre"/>
  </listOfParameters>
  ...
</model>

```

4.10 Initial assignments

SBML Level 2 Version 3 provides two ways of assigning initial values to entities in a model. The simplest and most basic is to set the values of the appropriate attributes in the relevant components; for example, the initial value of a model parameter (whether it is a constant or a variable) can be assigned by setting its `value` attribute directly in the model definition (Section 4.9). However, this approach is not suitable when the value must be calculated, because the initial value attributes on different components such as species, compartments, and parameters are single values and not mathematical expressions. This is the reason for the introduction of `InitialAssignment`: to permit the calculation of the value of a constant or the initial value of a variable from the values of other quantities in a model. The definition of `InitialAssignment` is shown in Figure 18 on the following page.

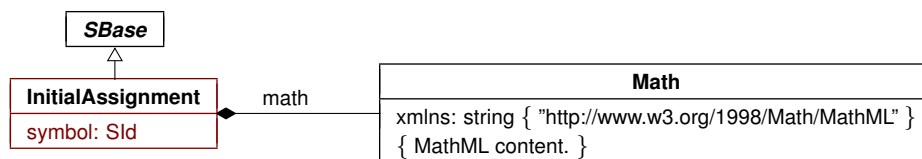


Figure 18: The definition of class *InitialAssignment*. The contents of the *Math* class can be any MathML permitted in SBML; see Section 3.4.1. A sequence of one or more instances of *InitialAssignment* objects can be located in an instance of *ListOfInitialAssignments* in *Model*, as shown in Figure 10.

As explained below, the provision of *InitialAssignment* does not mean that models necessarily must use this construct when defining initial values of quantities. If a value can be set using the relevant *attribute* of a component in a model, then that approach may be more efficient and more portable to other software tools. *InitialAssignment* should be used when the other mechanism is insufficient for the needs of a particular model.

Initial assignments have some similarities to assignment rules (Section 4.11.3). The main differences are (a) an *InitialAssignment* can set the value of a constant whereas an *AssignmentRule* cannot, and (b) unlike *AssignmentRule*, an *InitialAssignment* definition only applies up to and including the beginning of simulation time, i.e., $t \leq 0$, while an *AssignmentRule* applies at all times.

4.10.1 The symbol *attribute*

InitialAssignment contains the *attribute* *symbol*, of type *SId*. The value of this *attribute* in an *InitialAssignment* object can be the identifier (i.e., the value of the *id attribute*) of a *Compartment*, *Species* or *global Parameter* elsewhere in the model. The purpose of the *InitialAssignment* is to define the initial value of the constant or variable referred to by the *symbol attribute*. (The *attribute*'s name is *symbol* rather than *variable* because it may assign values to constants as well as variables in a model; see Section 4.10.4 below.)

An initial assignment cannot be made to reaction identifiers, that is, the *symbol attribute* value of an *InitialAssignment* cannot be an identifier that is the *id attribute* value of a *Reaction* object in the model. This is identical to a restriction placed on rules (see Section 4.11.5).

4.10.2 The math *element*

The *math element* contains a MathML expression that is used to calculate the value of the constant or the initial value of the variable. The units of the value computed by the formula in the *math element must be* identical to be the units associated with the identifier given in the *symbol attribute*. (That is, the units are the units of the species, compartment, or parameter, as appropriate for the kind of object identified by the value of *symbol*.)

4.10.3 The *sboTerm attribute*

InitialAssignment inherits from *Sbase* an optional *sboTerm attribute* of type *SBOTerm* (see Sections 3.1.9 and 5). When a value is given to this *attribute* in an initial assignment definition, the value must be a valid SBO identifier referring to a mathematical expression (i.e., terms derived from *SBO:0000064*, “mathematical expression”). The *InitialAssignment* object should have a “is a” relationship with the SBO term, and the term should be the most precise (narrow) term that captures the role of the *InitialAssignment* in the model.

As discussed in Section 5, SBO labels are optional information on a model. Applications are free to ignore *sboTerm* values. A model must be interpretable without the benefit of SBO labels.

4.10.4 Semantics of initial assignments

The value calculated by an *InitialAssignment* object overrides the value assigned to the given symbol by the object defining that symbol. For example, if a *Compartment*'s *size* is set in its definition, and the model also contains an *InitialAssignment* having that compartment's *id* as its *symbol* value, then the interpretation is that the *size* assigned in the *Compartment* definition should be ignored and the value assigned based on the

1 computation defined in the **InitialAssignment**. Initial assignments can take place for **Compartment**, **Species**
2 and **global Parameter** objects regardless of the value of their **constant attribute**.

3 This does not mean that a definition of a symbol can be omitted if there is an **InitialAssignment** object for
4 that symbol; the symbols must always be defined even if they are assigned a value separately. For example,
5 there must be a **Parameter** definition for a given parameter if there is an **InitialAssignment** for that parameter.

6 The actions of all **InitialAssignment** objects are in general terms the same, but differ in the precise details
7 depending on the type of variable being set:

- 8 • *In the case of a species*, an **InitialAssignment** sets the referenced species' initial quantity (*concentration*
9 or *amount of substance*) to the value determined by the formula in **math**. (See Section 4.8.5 for an
10 explanation of how the units of the species' quantity are determined.)
- 11 • *In the case of a compartment*, an **InitialAssignment** sets the referenced compartment's initial size to the
12 size determined by the formula in **math**. The overall units of the formula **must be the same as** the units
13 specified for the size of the compartment. (See Section 4.7.5 for an explanation of how the units of the
14 compartment's size are determined.)
- 15 • *In the case of a parameter*, an **InitialAssignment** sets the referenced parameter's initial value to that
16 determined by the formula in **math**. The overall units of the formula **must be the same as** the units
17 defined for the parameter. (See Section 4.9.3 for an explanation of how the units of the parameter are
18 determined.)

19 In the context of a simulation, initial assignments establish values that are in effect prior to and including
20 the start of simulation time, i.e., $t \leq 0$. Section 3.4.8 provides information about the interpretation of
21 assignments, rules, and entity values for simulation time up to and including the start time $t = 0$; this is
22 important for establishing the initial conditions of a simulation if the model involves expressions containing
23 the *delay csymbol* (Section 3.4.6).

24 There cannot be two initial assignments for the same symbol in a model; that is, a model must not contain
25 two or more **InitialAssignment** objects that both have the same identifier as their **symbol attribute** value.
26 A model must also not define initial assignments *and* assignment rules for the same entity. That is, there
27 cannot be *both* an **InitialAssignment** and an **AssignmentRule** for the same symbol in a model, because both
28 kinds of constructs apply prior to and at the start of simulated time—allowing both to exist for a given
29 symbol would result in indeterminism). (See also Section 4.11.5.)

30 The ordering of **InitialAssignment** objects is not significant. The combined set of **InitialAssignment**, **Assign-**
31 **mentRule** and **KineticLaw** objects form a set of assignment statements that must be considered as a whole. The
32 combined set of assignment statements should not contain algebraic loops: a chain of dependency between
33 these statements should terminate. (More formally, consider the directed graph of assignment statements
34 where nodes are a model's assignment statements and directed arcs exist for each occurrence of a symbol in
35 an assignment statement **math attribute**. The directed arcs in this graph start from the statement assigning
36 the symbol and end at the statement that contains the symbol in their **math elements**. Such a graph must
37 be acyclic.) Examples of valid and invalid set of assignment statements are given in Section 4.11.5.

38 Finally, it is worth being explicit about the expected behavior in the following situation. Suppose (1) a given
39 symbol has a value x assigned to it in its definition, and (2) there is an initial assignment having the identifier
40 as its **symbol** value and reassigning the value to y , *and* (3) the identifier is also used in the mathematical
41 formula of a second initial assignment. What value should the second initial assignment use? It is y , the
42 value assigned to the symbol by the first initial assignment, not whatever value was given in the symbol's
43 definition. This follows directly from the behavior at the defined at the beginning of this section and in
44 Section 3.4.8: if an **InitialAssignment** object exists for a given symbol, then the symbol's value is overridden
45 by that initial assignment.

46 4.10.5 Example

47 The following example shows how the species “**x**” can assigned the initial value $2 \times y$, where “**y**” is an
48 identifier defined elsewhere in the model:

```

1      <model>
2          ...
3          <listOfSpecies>
4              <species id="x" initialConcentration="5"/>
5          </listOfSpecies>
6          ...
7          <listOfInitialAssignments>
8              <initialAssignment symbol="x">
9                  <math xmlns="http://www.w3.org/1998/Math/MathML">
10                     <apply>
11                         <times/>
12                         <ci> y </ci>
13                         <cn> 2 </cn>
14                     </apply>
15                 </math>
16             </initialAssignment>
17         </listOfInitialAssignments>
18         ...
19     </model>

```

20 The next example illustrates the more complex behavior discussed above, when a symbol has a value assigned
21 in its definition but there also exists an **InitialAssignment** for it *and* another **InitialAssignment** uses its value
22 in its mathematical formula.

```

23     <model>
24         ...
25         <listOfSpecies>
26             <species id="x" initialConcentration="5"/>
27         </listOfSpecies>
28         ...
29         <listOfInitialAssignments>
30             <initialAssignment symbol="x">
31                 <math xmlns="http://www.w3.org/1998/Math/MathML">
32                     <cn> 2 </cn>
33                 </math>
34             </initialAssignment>
35             <initialAssignment symbol="othersymbol">
36                 <math xmlns="http://www.w3.org/1998/Math/MathML">
37                     <apply>
38                         <times/>
39                         <ci> x </ci>
40                         <cn> 2 </cn>
41                     </apply>
42                 </math>
43             </initialAssignment>
44         </listOfInitialAssignments>
45         ...
46     </model>

```

47 The value of “othersymbol” in the SBML excerpt above will be “4”. The case illustrates the rule of thumb
48 that if there is an initial assignment for a symbol, the value assigned to the symbol in its definition must be
49 ignored and the value created by the initial assignment used instead.

50 **4.11 Rules**

51 In SBML, *Rules* provide additional ways to define the values of variables in a model, their relationships,
52 and the dynamical behaviors of those variables. *Rules* enable the encoding of relationships that cannot be
53 expressed using reactions alone (Section 4.13) nor by the assignment of an initial value to a variable in a
54 model (Section 4.10).

55 SBML separates rules into three subclasses for the benefit of model analysis software. The three subclasses
56 are based on the following three different possible functional forms (where x is a variable, f is some arbitrary
57 function returning a numerical result, \mathbf{V} is a vector of variables that does not include x , and \mathbf{W} is a vector
58 of variables that may include x):

| | | |
|-------------------|-------------------------------------|-------------------------|
| <i>Algebraic</i> | left-hand side is zero: | $0 = f(\mathbf{W})$ |
| <i>Assignment</i> | left-hand side is a scalar: | $x = f(\mathbf{V})$ |
| <i>Rate</i> | left-hand side is a rate-of-change: | $dx/dt = f(\mathbf{W})$ |

In their general form given above, there is little to distinguish between *assignment* and *algebraic* rules. They are treated as separate cases for the following reasons:

- *Assignment* rules can simply be evaluated to calculate intermediate values for use in numerical methods;
- SBML needs to place restrictions on assignment rules, for example the restriction that assignment rules cannot contain algebraic loops (discussed further in Section 4.11.5);
- Some simulators do not contain numerical solvers capable of solving unconstrained algebraic equations, and providing more direct forms such as assignment rules may enable those simulators to process models they could not process if the same assignments were put in the form of general algebraic equations;
- Those simulators that *can* solve these algebraic equations make a distinction between the different categories listed above; and
- Some specialized numerical analyses of models may only be applicable to models that do not contain *algebraic* rules.

The approach taken to covering these cases in SBML is to define an abstract *Rule* object class containing an *element*, *math*, to hold the right-hand side expression, then to derive subtypes of *Rule* that add *attributes* to distinguish the cases of algebraic, assignment and rate rules. Figure 19 gives the definitions of *Rule* and the subtypes derived from it. The figure shows there are three subtypes, *AlgebraicRule*, *AssignmentRule* and *RateRule* derived directly from *Rule*. These correspond to the cases *Algebraic*, *Assignment*, and *Rate* described above respectively.

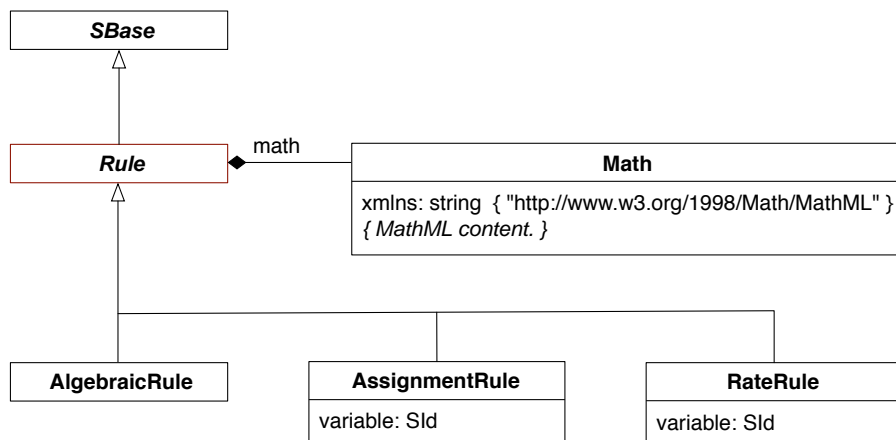


Figure 19: The definition of *Rule* and derived types *AlgebraicRule*, *AssignmentRule* and *RateRule*.

4.11.1 Common attributes in *Rule*

The classes derived from *Rule* inherit *math* and the attributes and elements from *Sbase*, including *sboTerm*.

The *math* element

A *Rule* object has a required *element* called *math*, containing a MathML expression defining the mathematical formula of the rule. This MathML formula must return a numerical value. The formula can be an arbitrary expression referencing the variables and other entities in an SBML model. The interpretation of *math* and the units of the formula are described in more detail in Sections 4.11.2, 4.11.3 and 4.11.4 below.

1 *The `sboTerm` attribute*

2 The **Rule** object class inherits from **Sbase** the optional `sboTerm` attribute of type **SBOTerm** (see Sections 3.1.9
3 and 5). When a value is given to this attribute, it must be a valid SBO identifier referring to a mathematical
4 expression defined in SBO (i.e., terms derived from **SBO:0000064**, “mathematical expression”). The **AlgebraicRule**,
5 **AssignmentRule**, or **RateRule** object should have a “is a” relationship with the SBO term, and the
6 term should be the most precise (narrow) term that captures the role of that rule in the model.

7 As discussed in Section 5, SBO labels are optional information on a model. Applications are free to ignore
8 `sboTerm` values. A model must be interpretable without the benefit of SBO labels.

9 **4.11.2 AlgebraicRule**

10 The rule type **AlgebraicRule** is used to express equations that are neither assignments of model variables nor
11 rates of change. **AlgebraicRule** does not add any attributes to the basic **Rule**; its role is simply to distinguish
12 this case from the other cases. An example of the use of **AlgebraicRule** is given in Section 7.5.

13 In the context of a simulation, algebraic rules are in effect at all times, $t \geq 0$. For purposes of evaluating
14 expressions that involve the *delay* `csymbol` (Section 3.4.6), algebraic rules are considered to apply also at
15 $t \leq 0$. Section 3.4.8 provides additional information about the semantics of assignments, rules, and entity
16 values for simulation time $t \leq 0$.

17 **AlgebraicRule** is unlike **AssignmentRule** and **RateRule** in that there are no specific definitions of the meaning
18 of the mathematical expression in the **AlgebraicRule**. There is only one requirement on the units, which is
19 that the units of the overall expression must be consistent under dimensional analysis.

20 The ability to define arbitrary algebraic expressions in an SBML model introduces the possibility that a
21 model is mathematically overdetermined by the overall system of equations constructed from its rules and
22 reactions. An SBML model must not be overdetermined; this is discussed in Section 4.11.5 below.

23 **4.11.3 AssignmentRule**

24 The rule type **AssignmentRule** is used to express equations that set the values of variables. The left-hand
25 side (the `variable` attribute) of an assignment rule can refer to the identifier of a **Species**, **Compartment**,
26 or **Parameter** object in the model (but not a reaction). The entity identified must not have its `constant`
27 attribute set to “true”. The effects of an **AssignmentRule** are in general terms the same, but differ in the
28 precise details depending on the type of variable being set:

- 29 • *In the case of a species*, an **AssignmentRule** sets the referenced species’ quantity (*concentration* or
30 *amount of substance*) to the value determined by the formula in `math`. The units of the formula in
31 `math` must be the same as the *units of the species* (Section 4.8.5) for the species identified by the
32 `variable` attribute of the **AssignmentRule**.

33 *Restrictions:* There must not be both an **AssignmentRule** `variable` attribute and a **SpeciesReference**
34 `species` attribute having the same value, unless that species has its `boundaryCondition` attribute
35 set to “true”. In other words, an assignment rule cannot be defined for a species that is created or
36 destroyed in a reaction unless that species is defined as a boundary condition in the model.

- 37 • *In the case of a compartment*, an **AssignmentRule** sets the referenced compartment’s size to the value
38 determined by the formula in `math`. The overall units of the formula in `math` must be the same as the
39 units of the size of the compartment (Section 4.7.5).
- 40 • *In the case of a parameter*, an **AssignmentRule** sets the referenced parameter’s value to that determined
41 by the formula in `math`. The overall units of the formula in `math` must be the same as the units defined
42 for the parameter (Section 4.9.3).

43 In the context of a simulation, assignment rules are in effect at all times, $t \geq 0$. For purposes of evaluating
44 expressions that involve the *delay* `csymbol` (Section 3.4.6), assignment rules are considered to apply also at
45 $t \leq 0$. Section 3.4.8 provides additional information about the semantics of assignments, rules, and entity
46 values for simulation time $t \leq 0$.

1 A model must not contain more than one **AssignmentRule** or **RateRule** object having the same value of
2 **variable**; in other words, in the set of all assignment rules and rate rules in an SBML model, each variable
3 appearing in the left-hand sides can only appear once. This simply follows from the fact that an indeterminate
4 system would result if a model contained more than one assignment rule for the same variable or both an
5 assignment rule and a rate rule for the same variable.

6 Similarly, a model must also not contain *both* an **AssignmentRule** and an **InitialAssignment** for the same
7 variable, because both kinds of constructs apply prior to and at the start of simulation time, i.e., $t \leq 0$. If a
8 model contained both an initial assignment and an assignment rule for the same variable, an indeterminate
9 system would result. (See also Section 4.10.4.)

10 The value calculated by an **AssignmentRule** object overrides the value assigned to the given symbol by the
11 object defining that symbol. For example, if a **Compartment**'s **size** is set in its definition, and the model also
12 contains an **AssignmentRule** having that compartment's **id** as its **variable** value, then the **size** assigned
13 in the **Compartment** definition is ignored and the value assigned based on the computation defined in the
14 **AssignmentRule**. This does *not* mean that a definition for a given symbol can be omitted if there is an
15 **AssignmentRule** object for it. For example, there must be a **Parameter** definition for a given parameter if
16 there is an **AssignmentRule** for that parameter.

17 4.11.4 RateRule

18 The rule type **RateRule** is used to express equations that determine the rates of change of variables. The
19 left-hand side (the **variable attribute**) can refer to the identifier of a species, compartment, or parameter
20 (but not a reaction). The entity identified must have its **constant attribute** set to “false”. The effects of a
21 **RateRule** are in general terms the same, but differ in the precise details depending on which variable is being
22 set:

- 23 • *In the case of a species*, a **RateRule** sets the rate of change of the species' quantity (*concentration* or
24 *amount of substance*) to the value determined by the formula in **math**. The overall units of the formula
25 in **math** must be *species quantity/time*, where the *time* units are the built-in units of time described in
26 Section 4.4 and the *species quantity* units are the *units of the species* as defined in Section 4.8.5.

27 *Restrictions:* There must not be both a **RateRule variable attribute** and a **SpeciesReference species**
28 **attribute** having the same value, unless that species has its **boundaryCondition attribute** is set to
29 “true”. This means a rate rule cannot be defined for a species that is created or destroyed in a
30 reaction, unless that species is defined as a boundary condition in the model.

- 31 • *In the case of a compartment*, a **RateRule** sets the rate of change of the compartment's size to the value
32 determined by the formula in **math**. The overall units of the formula must be *size/time*, where the *time*
33 units are the built-in units of time described in Section 4.4 and the *size* units are the units of size on
34 the compartment (Section 4.7.5).
- 35 • *In the case of a parameter*, a **RateRule** sets the rate of change of the parameter's value to that determined
36 by the formula in **math**. The overall units of the formula must be $x/time$, where x are the units of the
37 parameter (Section 4.9.3).

38 In the context of a simulation, rate rules are in effect for simulation time $t > 0$. Other types of rules and
39 initial assignments are in effect at different times; Section 3.4.8 describes these conditions.

40 As mentioned in Section 4.11.3 for **AssignmentRule**, a model must not contain more than one **RateRule** or
41 **AssignmentRule** object having the same value of **variable**; in other words, in the set of all assignment rules
42 and rate rules in an SBML model, each variable appearing in the left-hand sides can only appear once. This
43 simply follows from the fact that an indeterminate system would result if a model contained more than one
44 assignment rule for the same variable or both an assignment rule and a rate rule for the same variable.

45 4.11.5 Additional restrictions on rules

46 An important design goal of SBML rule semantics is to ensure that a model's simulation and analysis results
47 will not be dependent on when or how often rules are evaluated. To achieve this, SBML needs to place

two additional restrictions on rule use in addition to the conditions described above regarding the use of **AlgebraicRule**, **AssignmentRule** and **RateRule**. The first concerns algebraic loops in the system of assignments in a model, and the second concerns overdetermined systems.

The model must not contain algebraic loops

The combined set of **InitialAssignment**, **AssignmentRule** and **KineticLaw** objects constitute a set of assignment statements that should be considered as a whole. (A **KineticLaw** object is counted as an assignment because it assigns a value to the symbol contained in the **id** attribute of the **Reaction** object in which it is defined.) This combined set of assignment statements must not contain algebraic loops—dependency chains between these statements must terminate. To put this more formally, consider a directed graph in which nodes are assignment statements and directed arcs exist for each occurrence of an SBML species, compartment or parameter symbol in an assignment statement’s **math element**. Let the directed arcs point from the statement assigning the symbol to the statements that contain the symbol in their **math element** expressions. This graph must be acyclic.

SBML does not specify when or how often rules should be evaluated. Eliminating algebraic loops ensures that assignment statements can be evaluated any number of times without the result of those evaluations changing. As an example, consider the following equations:

$$x = x + 1, \quad y = z + 200, \quad z = y + 100$$

If this set of equations were interpreted as a set of assignment statements, it would be invalid because the rule for x refers to x (exhibiting one type of loop), and the rule for y refers to z while the rule for z refers back to y (exhibiting another type of loop).

Conversely, the following set of equations would constitute a valid set of assignment statements:

$$x = 10, \quad y = z + 200, \quad z = x + 100$$

The model must not be overdetermined

An SBML model must not be overdetermined; that is, a model must not define more equations than there are unknowns in a model. An SBML model that does not contain **AlgebraicRule** objects cannot be overdetermined.

Assessing whether a given continuous, deterministic, mathematical model is overdetermined does not require dynamic analysis; it can be done by analyzing the system of equations created from the model. One approach is to construct a bipartite graph in which one set of vertices represents the variables and the other the set of vertices represents the equations. Place edges between vertices such that variables in the system are linked to the equations that determine them. For algebraic equations, there will be edges between the equation and each variable occurring in the equation. For ordinary differential equations (such as those defined by rate rules or implied by the reaction rate definitions), there will be a single edge between the equation and the variable determined by that differential equation. A mathematical model is overdetermined if the maximal matchings (Chartrand, 1977) of the bipartite graph contain disconnected vertexes representing equations. (If one maximal matching has this property, then all the maximal matchings will have this property; i.e., it is only necessary to find one maximal matching.) Appendix D describes a method of applying this procedure to specific SBML data objects.

4.11.6 Example of rule use

This section contains an example set of rules. Consider the following set of equations:

$$k = \frac{k_3}{k_2}, \quad s_2 = \frac{k \cdot x}{1 + k_2}, \quad A = 0.10 \cdot x$$

This can be encoded by the following scalar rule set (where the definitions of \mathbf{x} , \mathbf{s} , \mathbf{k} , \mathbf{k}_2 , \mathbf{k}_3 and \mathbf{A} are assumed to be located elsewhere in the model and not shown in this abbreviated example):

```

1      <listOfRules>
2          <assignmentRule variable="k">
3              <math xmlns="http://www.w3.org/1998/Math/MathML">
4                  <apply> <divide/> <ci> k3 </ci> <ci> k2 </ci> </apply>
5              </math>
6          </assignmentRule>
7          <assignmentRule variable="s2">
8              <math xmlns="http://www.w3.org/1998/Math/MathML">
9                  <apply>
10                     <divide/>
11                     <apply> <times/> <ci> k </ci> <ci> x </ci> </apply>
12                     <apply> <plus/> <cn> 1 </cn> <ci> k2 </ci> </apply>
13                 </apply>
14             </math>
15         </assignmentRule>
16         <assignmentRule variable="A">
17             <math xmlns="http://www.w3.org/1998/Math/MathML">
18                 <apply> <times/> <cn> 0.10 </cn> <ci> x </ci> </apply>
19             </math>
20         </assignmentRule>
21     </listOfRules>

```

4.12 Constraints

The **Constraint** object is a mechanism for stating the assumptions under which a model is designed to operate. The *constraints* are statements about permissible values of different quantities in a model. Figure 20 shows the definition of the **Constraint** object class.

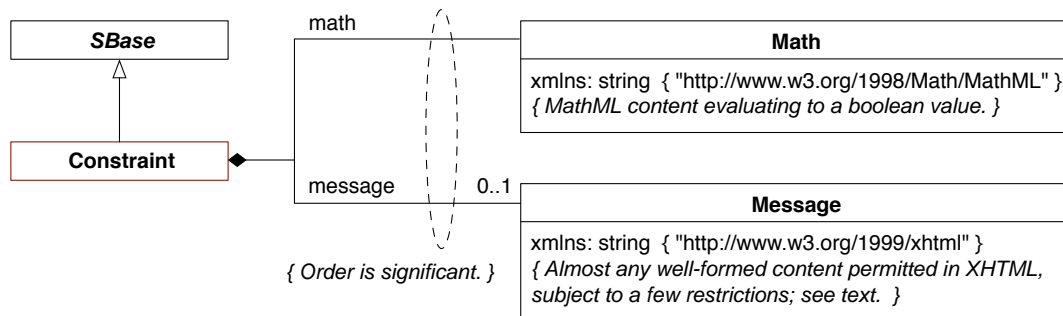


Figure 20: The definition of class **Constraint**. The contents of the **Math** class can be any MathML permitted in SBML, but it must return a boolean value. As shown above, an instance of **Constraint** can also contain zero or one instances of **Message**; this element is simply a wrapper (in the XML form, within `<message> . . . </message>` tags) for XHTML content. The same guidelines for XHTML content as explained in Section 3.2.3 for notes on **SBbase** also apply to the XHTML within messages in a **Constraint**. A sequence of one or more instances of **Constraint** objects can be located in an instance of **ListOfConstraints** in **Model**, as shown in Figure 10.

The essential meaning of a constraint is this: if a dynamical analysis of a model (such as a simulation) reaches a state in which a constraint is no longer satisfied, the results of the analysis are deemed invalid beginning with that point in time. The exact behavior of a software tool, upon encountering a constraint violation, is left up to the software; *however*, a software tool must somehow indicate to the user when a model's constraints are no longer satisfied. (Otherwise, a user may not realize that the analysis has reached an invalid state and is potentially producing nonsense results.) If a software tool does not have support for constraints, it should indicate this to the user when encountering a model containing constraints.

4.12.1 The math element

Constraint has one required subelement, **math**, containing a MathML formula defining the condition of the constraint. This formula must return a boolean value of "true" when the model is in a *valid* state. The formula can be an arbitrary expression referencing the variables and other entities in an SBML model. The evaluation of **math** and behavior of constraints are described in more detail in Section 4.12.4 below.

4.12.2 The message element

A **Constraint** object has an optional **element** called **message**. This can contain a message in XHTML format that may be displayed to the user when the condition of the constraint in **math** evaluates to a value of “false”. Software tools are not required to display the message, but it is recommended that they do so as a matter of best practice.

The XHTML content within a **message** element must follow the same restrictions as for the **notes** element on **Sbase** described in Section 3.2.3. For example, **message** must not contain an XML declaration or a DOCTYPE declaration, and the permitted content can only take one of the following general forms: (1) a complete XHTML document beginning with the element `<html>` and ending with `</html>`; (2) the “body” portion of a document beginning with the element `<body>` and ending with `</body>`; or (3) XHTML content that is permitted within a `<body> ... </body>` elements. Appendix F describes one approach to reading the **message** content.

4.12.3 The sboTerm attribute

The **Constraint** object inherits from **Sbase** the optional **sboTerm** attribute of type **SBOTerm** (see Sections 3.1.9 and 5). When a value is given to this attribute in a constraint definition, the value must be a valid SBO identifier referring to a mathematical expression (i.e., terms derived from **SBO:0000064**, “mathematical expression”). The **Constraint** should have an “is a” relationship with the SBO term, and the term should be the most precise (narrow) term that captures the role of the **Constraint** in the model.

As discussed in Section 5, SBO labels are optional information on a model. Applications are free to ignore **sboTerm** values. A model must be interpretable without the benefit of SBO labels.

4.12.4 Semantics of constraints

In the context of a simulation, a **Constraint** has effect at all times $t \geq 0$. Each **Constraint**’s **math** element is first evaluated after any **InitialAssignment** definitions in a model at $t = 0$ and can conceivably trigger at that point. (In other words, a simulation could fail a constraint immediately.)

Constraint definitions *cannot and should not* be used to compute the dynamical behavior of a model as part of, for example, simulation. Constraints may be used as input to non-dynamical analysis, for instance by expressing flux constraints for flux balance analysis.

The results of a simulation of a model containing a constraint are invalid from any simulation time at and after a point when the function given by the **math** returns a value of “false”. Invalid simulation results do not make a prediction of the behavior of the biochemical reaction network represented by the model. The precise behavior of simulation tools is left undefined with respect to constraints. If invalid results are detected with respect to a given constraint, the **message** element (Section 4.12.2) may optionally be displayed to the user. The simulation tool may also halt the simulation or clearly delimit in output data the simulation time point at which the simulation results become invalid.

SBML does not impose restrictions on duplicate **Constraint** definitions or the order of evaluation of **Constraint** objects in a model. It is possible for a model to define multiple constraints all with the same **math** element. Since the failure of any constraint indicates that the model simulation has entered an invalid state, a system is not required to attempt to detect whether other constraints in the model have failed once any one constraint has failed.

4.12.5 Example

As an example, the following SBML fragment demonstrates the constraint that species S_1 should only have values between 1 and 100:

```
<model>
...
<listOfConstraints>
  <constraint>
    <math xmlns="http://www.w3.org/1998/Math/MathML">
```

```

1         <apply>
2             <and/>
3                 <apply> <lt/> <cn> 1 </cn> <ci> S1 </ci> </apply>
4                 <apply> <lt/> <ci> S1 </ci> <cn> 100 </cn> </apply>
5             </apply>
6         </math>
7         <message>
8             <p xmlns="http://www.w3.org/1999/xhtml"> Species S1 is out of range. </p>
9         </message>
10        </constraint>
11    </listOfConstraints>
12    ...
13 </model>

```

4.13 Reactions

A *reaction* represents any transformation, transport or binding process, typically a chemical reaction, that can change the quantity of one or more species. In SBML, a reaction is defined primarily in terms of the participating reactants and products (and their corresponding stoichiometries), along with optional modifier species, an optional rate at which the reaction takes place, and optional parameters. These various parts of a reaction are recorded in the SBML **Reaction** object class and other supporting data classes, defined in Figure 21 on the next page.

4.13.1 Reaction

Each reaction in an SBML model is defined using an instance of a **Reaction** object. As shown in Figure 21 on the following page, it contains several scalar attributes and several lists of objects.

The id and name attributes

As with most other main kinds of objects in SBML, the **Reaction** object class includes a mandatory attribute called **id**, of type **SId**, and an optional attribute **name**, of type **string**. The **id** attribute is used to give the reaction a unique identifier in the model. This identifier can be used in mathematical formulas elsewhere in an SBML model to represent the rate of that reaction; this usage is explained in detail in Section 4.13.7 below. The **name** attribute can be used to give the reaction a more free-form, descriptive name. The **name** and **id** attributes must be used as described in Section 3.3.

The lists of reactants, products and modifiers

The species participating as reactants, products, and/or modifiers in a reaction are declared using lists of **SpeciesReference** and/or **ModifierSpeciesReference** instances stored in **listOfReactants**, **listOfProducts** and **listOfModifiers**. **SpeciesReference** and **ModifierSpeciesReference** are described in more detail in Sections 4.13.3 and 4.13.4 below.

Certain restrictions are placed on the appearance of species in reaction definitions:

- The ability of a species to appear as a reactant or product of any reaction in a model is governed by certain flags in that species' definition; see Section 4.8.6 for more information.
- Any species appearing in the mathematical formula of the **kineticLaw** of a **Reaction** instance must be declared in at least one of that **Reaction**'s lists of reactants, products, and/or modifiers. Put another way, it is an error for a reaction's kinetic law formula to refer to species that have not been declared for that reaction.
- A reaction definition can contain an empty list of reactants *or* an empty list of products, but it must have at least one reactant or product; in other words, a reaction without any reactant or product species is not permitted. (This restriction does not apply to modifier species, which remain optional in all cases.)

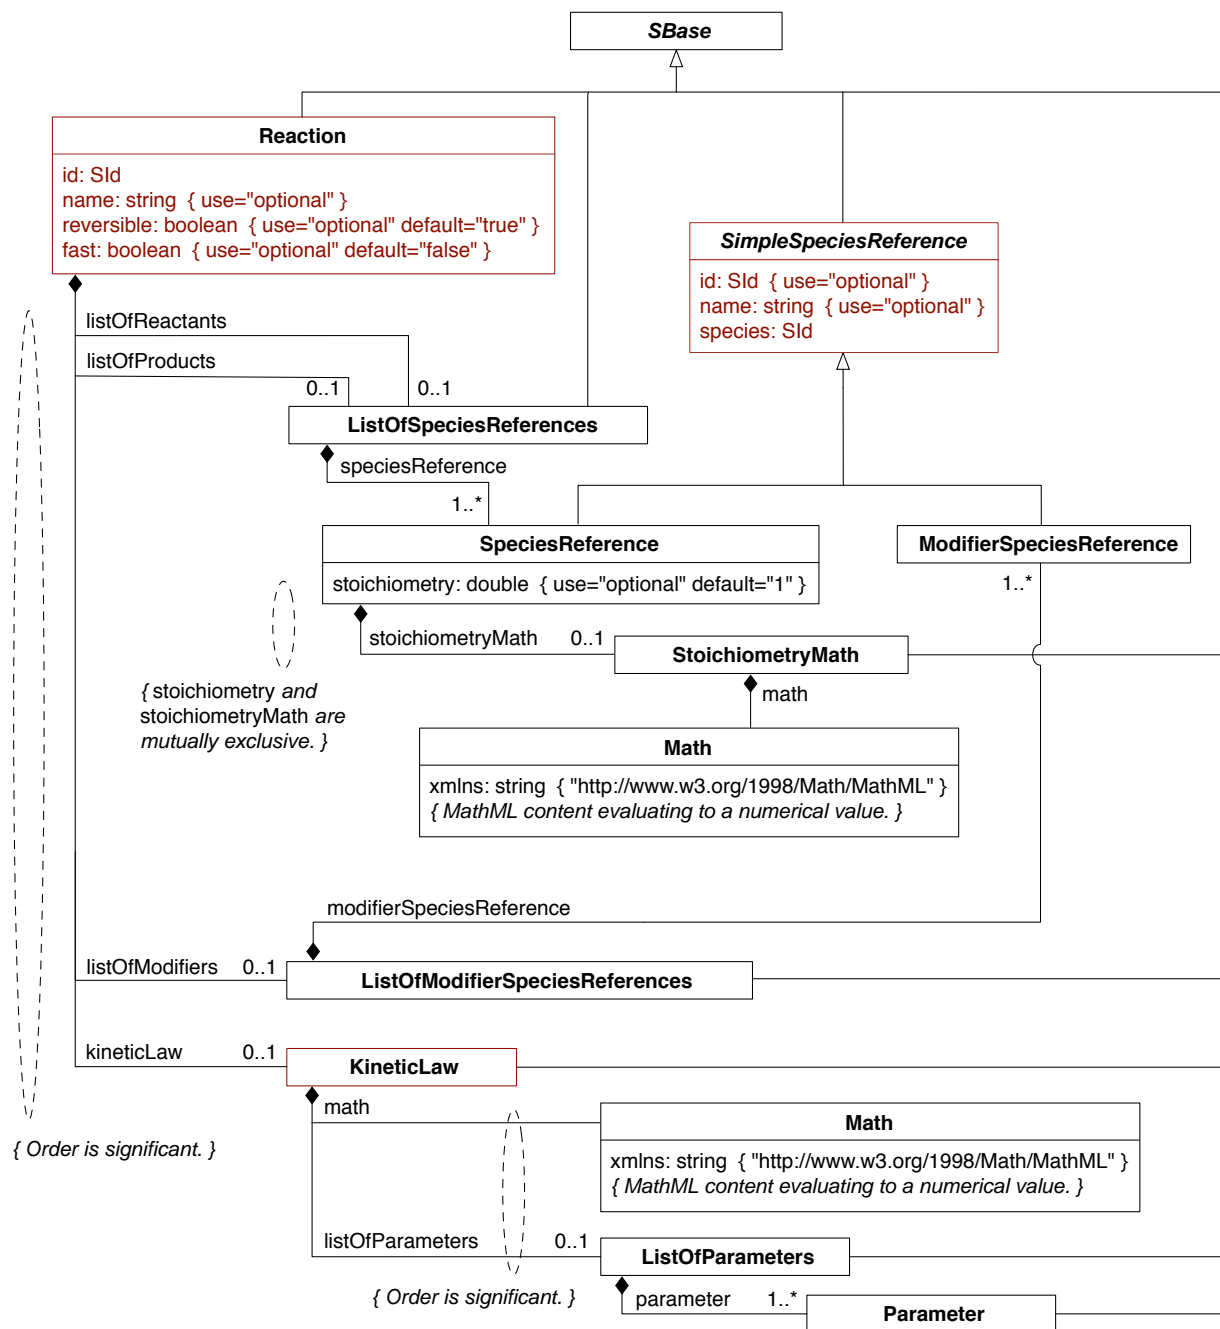


Figure 21: The definitions of classes *Reaction*, *KineticLaw*, *SpeciesReference*, *ModifierSpeciesReference*, as well as the container classes *ListOfSpeciesReferences*, *ListOfModifierSpeciesReferences*, and *ListOfParameters*. Note that *SimpleSpeciesReference* is an abstract class used only to provide some common attributes to its derived classes. The class *Parameter* is defined in Section 4.9.

1 *The kineticLaw element*

2 A reaction can contain up to one *KineticLaw* object in the *kineticLaw* element of the *Reaction*. This “kinetic
 3 law” defines the speed at which the process defined by the reaction takes place. A detailed description of
 4 *KineticLaw* is left to Section 4.13.5 below.

Note that the inclusion of a **KineticLaw** object in an instance of a **Reaction** component is optional; however, in general there is no useful default that can be substituted in place of a missing rate expression in a reaction. Moreover, a reaction's rate cannot be defined in any other way in SBML—**InitialAssignment**, **AssignmentRule**, **RateRule**, **AlgebraicRule**, **Event**, and other constructs in SBML cannot be used to set the reaction rate separately. Nevertheless, for some modeling applications, reactions without any defined rate can be perfectly acceptable.

The reversible attribute

The optional boolean attribute **reversible** indicates whether the reaction is reversible. The default is “true”.

To say that a reaction is *reversible* is to say it can proceed in either the forward or the reverse direction. Although the reversibility of a reaction can sometimes be deduced by inspecting its rate expression, this is not always the case, especially for complicated expressions. Having a separate attribute supports the ability to perform some kinds of model analyses in the absence of performing a time-course simulation of the model. Moreover, the need in SBML to allow rate expressions (i.e., **KineticLaw**) to be optional leads to the need for a separate flag indicating reversibility. Information about reversibility in the absence of a **KineticLaw** in a **Reaction** is useful in certain kinds of structural analyses such as elementary mode analysis.

Mathematically, the **reversible** attribute on **Reaction** has no impact on the construction of the equations giving the overall rates of change of each species quantity in a model. A concrete explanation may help illustrate this. Suppose a model consists of multiple reactions, of which two particular irreversible reactions R_f and R_r are actually the forward and reverse processes of the same underlying reaction. The product species of R_f necessarily will be the reactants of R_r , and the reactants of R_f will be the products of R_r . Let $f_f(\mathbf{X})$ and $f_r(\mathbf{X})$ be the SBML kinetic rate formulas of R_f and R_r , respectively, with \mathbf{X} representing the species, parameters and compartments in the model. For the sake of this example, suppose we are using a continuous differential equation framework to simulate the system of reactions. Then for each species, we need to construct an expression representing the overall rate of change of that species' amount in the model. This overall expression will be a sum of the contributions of all the relevant rate formulas,

$$\frac{dS}{dt} = \dots - n \cdot f_f(\mathbf{X}) + n \cdot f_r(\mathbf{X}) + \dots$$

where S is a reactant species of R_f and a product of R_r , n is the effective stoichiometry of S in R_f (which by implication must be the same as its stoichiometry in R_r), and “...” indicates other rate formulas in the model involving the particular species S . Now, contrast this to the case of an identical second SBML model, except that instead of having separate **Reaction** definitions for the forward and reverse reactions, this model has a single **Reaction** R_c labeled as reversible and whose reactants and products are the same as those of R_f in the first model. The rate of this reaction will be a formula $f_c = f_f(\mathbf{X}) - f_r(\mathbf{X})$. In constructing an expression representing the overall rate of change for the species S involved in that reaction, we will have

$$\begin{aligned} \frac{dS}{dt} &= \dots - n \cdot f_c(\mathbf{X}) + \dots \\ &= \dots - n \cdot f_f(\mathbf{X}) + n \cdot f_r(\mathbf{X}) + \dots \end{aligned}$$

In other words, the result is the same final expression for the rate of change of a species. Although in this simple example we used an expression for f_c that had clearly separated terms, in the general case the expression may have a more complicated form.

Note that labeling a reaction as irreversible is an assertion that the reaction always proceeds in the given forward direction. (Why else would it be flagged as irreversible?) This implies the rate expression in the **KineticLaw** always has a non-negative value during simulations. Software tools could provide a means of optionally testing that this condition holds. The presence of reversibility information in two places (i.e., the rate expression and the **reversible** flag) leaves open the possibility that a model could contain contradictory information, but the creation of such a model would be an error on the part of the software generating it.

1 *The fast attribute*

2 The optional boolean attribute **fast** is another optional boolean attribute of **Reaction**. The attribute's
3 default value is "false".

4 Previous definitions of SBML indicated that software tools could ignore this attribute if they did not im-
5 plement support for the corresponding concept; however, further research has revealed that this is incorrect
6 and **fast cannot be ignored** if it is set to "true". **SBML Level 2 Version 3** therefore stipulates that if a
7 model has any reactions with **fast** set to "true", a software tool must be able to respect the attribute or
8 else indicate to the user that it does not have the capacity to do so. Analysis software cannot ignore the
9 value of the **fast** attribute because doing so may lead to different results as compared to a software system
10 that *does* make use of **fast**.

11 When a model contains true values for **fast** on any of its reactions, it is an indication that the creator of
12 the model is distinguishing different time scales of reactions in the system. The model's reaction definitions
13 are divided into two sets by the values of the **fast** attributes. The set of reactions having **fast**="true"
14 (known as *fast reactions*) should be assumed to be operating on a time scale significantly faster than the
15 other reactions (the *slow reactions*). Fast reactions are considered to be instantaneous relative to the slow
16 reactions.

17 Software tools must use a pseudo steady-state approximation for the set of fast reactions when constructing
18 the system of equations for the model. More specifically, the set of reactions that have the **fast** attribute
19 set to "true" forms a subsystem that should be described by a pseudo steady-state approximation in rela-
20 tionship to all other reactions in the model. Under this description, relaxation from any initial condition or
21 perturbation from any intermediate state of this subsystem would be infinitely fast. Appendix E provides a
22 technical explanation of an approach to solving systems with fast reactions.

23 The correctness of the approximation requires a significant separation of time scales between the fast reactions
24 and other processes. This is not trivial to estimate a priori, and may even change over the course of a
25 simulation, but can reasonably be assessed a posteriori in most cases.

26 *The sboTerm attribute on Reaction*

27 The **Reaction** object class inherits from **Sbase** the optional **sboTerm** attribute of type **SBOterm** (see Sec-
28 tions 3.1.9 and 5). When a value is given to this attribute in a reaction definition, the value must be a
29 valid SBO identifier referring to an event (i.e., terms derived from **SBO:0000231**, "event"). The **Reaction**
30 object should have an "is a" relationship with the SBO term. The SBO term chosen should be the most
31 precise (narrow) term that defines the event represented by the reaction as a whole. For example, a given
32 reaction in a model might represent *binding*, which has a term in SBO (identifier **SBO:0000177**, definition
33 "Interaction between several biochemical entities that results in the formation of a non-covalent complex").
34 The corresponding **Reaction** instance in SBML would have **sboTerm**="SBO:0000177".

35 As mentioned elsewhere, SBO labels are optional information on a model. Applications are free to ignore
36 **sboTerm** values, and a model must be interpretable without the benefit of SBO labels. Section 5 gives more
37 information about this principle and the use of SBO.

38 **4.13.2 The SimpleSpeciesReference abstract type**

39 As mentioned above, every species that enters into a given reaction must appear in that reaction's lists of
40 reactants, products and/or modifiers. In an SBML model, all species that may participate in any reaction
41 are listed in the **listOfSpecies** element of the top-level **Model** instance (see Section 4.2). Lists of products,
42 reactants and modifiers in **Reaction** objects do not introduce new species, but rather, they refer back to
43 those listed in the model's top-level **listOfSpecies**. For reactants and products, the connection is made
44 using a **SpeciesReference** object; for modifiers, it is made using a **ModifierSpeciesReference** object. **Simple-**
45 **SpeciesReference**, defined in Figure 21 on page 64, is an abstract type that serves as the parent class of both
46 **SpeciesReference** and **ModifierSpeciesReference**. It is used simply to hold the attributes and elements that
47 are common to the latter two objects.

The `id` and `name` attributes

The optional identifier stored in the `id` attribute allows `SpeciesReference` and `ModifierSpeciesReference` instances to be referenced from other `object`. No SBML `object classes` currently do this; however, such `classes` are anticipated in future SBML Levels. The value of `id` must be a text string conforming to the syntax permitted by the `SId` data type described in Section 3.1.7. The `id` value (whether it is in a `SpeciesReference` or `ModifierSpeciesReference` object) exists in the global namespace of the model, as described in Section 3.3. The `id` and `name` attributes must be used as described in Section 3.3.

The `species` attribute

The `SimpleSpeciesReference` object class has a required attribute, `species`, of type `SId`. As with the other attributes, it is inherited by `SpeciesReference` and `ModifierSpeciesReference`. The value of `species` must be the identifier of a species defined in the enclosing `Model`. The species is thereby declared as participating in the reaction being defined. The precise role of that species as a reactant, product, or modifier in the reaction is determined by the subtype of `SimpleSpeciesReference` (i.e., either `SpeciesReference` or `ModifierSpeciesReference`) in which the identifier appears.

The `sboTerm` attribute

The class `SimpleSpeciesReference` inherits from `Sbase` the optional `sboTerm` attribute of type `SBOTerm` (see Sections 3.1.9 and 5). This means that the object classes derived from `SimpleSpeciesReference`, namely `SpeciesReference` and `ModifierSpeciesReference`, all have `sboTerm` attributes. When a value is given to this attribute, it must be a valid SBO identifier referring to a participant role. The appropriate SBO term depends on whether the object is a reactant, product or modifier. If a reactant, then it should be a term in the `SBO:0000010`, “reactant” hierarchy; if a product, then it should be a term in the `SBO:0000011`, “product” hierarchy; and if a modifier, then it should be a term in the `SBO:0000019`, “modifier” hierarchy. The `SpeciesReference` and `ModifierSpeciesReference` instances should have an “is a” relationship to the term identified by the SBO identifier.

The SBO terms chosen should be the most precise (narrow) one that defines the role of the species in the reaction. An SBO reactant term can only be assigned to the `sboTerm` attribute of a `SpeciesReference` instance when that instance is contained in the list of reactants in the containing `Reaction` instance. Similarly, an SBO product term can only be assigned to the `sboTerm` attribute of a `SpeciesReference` instance when that instance is contained in the list of products in the containing `Reaction` instance.

4.13.3 `SpeciesReference`

The `Reaction` object class provides a way to express which species act as reactants and which species act as products in a reaction. In a given reaction, references to those species acting as reactants and/or products are made using instances of `SpeciesReference` objects in `Reaction`'s lists of reactants and products. The `SpeciesReference` structure inherits the mandatory attribute `species` and optional attributes `id`, `name`, and `sboTerm`, from the parent type `SimpleSpeciesReference`; see Section 4.13.2 for their definitions. It also defines attribute `stoichiometry` and element `stoichiometryMath`, described below.

The `species` attribute value must be the identifier of an existing species defined in the enclosing `Model`; the species is thereby designated as a reactant or product in the reaction. Which one it is (i.e., reactant or product) is indicated by whether the `SpeciesReference` appears in the `Reaction`'s reactant or product lists.

The `stoichiometry` attribute and `stoichiometryMath` element

Product and reactant stoichiometries can be specified using either `stoichiometry` or `stoichiometryMath` in a `SpeciesReference` object. The `stoichiometry` attribute is of type `double` and should contain values greater than zero (0). The `stoichiometryMath` element is implemented as an element containing a `MathML` expression. These two are mutually exclusive; only one of `stoichiometry` or `stoichiometryMath` should be defined in a given `SpeciesReference` instance. When neither the attribute nor the element is present, the value of `stoichiometry` in the `SpeciesReference` instance defaults to “1”.

For maximum interoperability, **SpeciesReference**'s **stoichiometry** attribute should be used in preference to **stoichiometryMath** when a species' stoichiometry is a simple scalar number (integer or decimal). When the stoichiometry is a rational number, or when it is a more complicated formula, **stoichiometryMath** must be used. The MathML expression in **stoichiometryMath** may also refer to identifiers of entities in a model (except reaction identifiers), as discussed in Section 3.4.3. However, the only species identifiers that can be used in **stoichiometryMath** are those referenced in the **Reaction** list of reactants, products and modifiers.

The **stoichiometry** attribute and the **stoichiometryMath** element, when either is used, is each interpreted as a factor applied to the reaction rate to give the rate of change of the species identified by the **species** attribute. This is the normal interpretation of a stoichiometry, but in SBML, one additional consideration has to be taken into account. The reaction rate, which is the result of the **KineticLaw** **math** element, is always in the model's *substance per time* units. However, the rate of change of the species will involve the species' *substance* units (i.e., the units identified by the **Species** object's **substanceUnits** attribute), and these units may be different from the model's default *substance* units. If the units are different, the stoichiometry must incorporate a conversion factor for converting the model's *substance* units to the species' *substance* units. The conversion factor is assumed to be included in the scalar value of the **stoichiometry** attribute if **stoichiometry** is used. If instead **stoichiometryMath** is used, then the product of the model's *substance* units times the **stoichiometryMath** units must match the *substance* units of the species. Note that in either case, if the species' units and the model's default *substance* units are the same, the stoichiometry ends up being a dimensionless number and equivalent to the standard chemical stoichiometry found in textbooks. Examples and more explanations of this are given in Section 4.13.6.

The following is a simple example of a species reference for species "X0", with stoichiometry "2", in a list of reactants within a reaction having the identifier "J1":

```

23 <model>
24   ...
25   <listOfReactions>
26     <reaction id="J1">
27       <listOfReactants>
28         <speciesReference species="X0" stoichiometry="2">
29       </listOfReactants>
30       ...
31     </reaction>
32     ...
33   </listOfReactions>
34   ...
35 </model>

```

The following is a more complex example of a species reference for species "X0", with a stoichiometry formula consisting of a rational number:

```

38 <model>
39   ...
40   <listOfReactions>
41     <reaction id="J1">
42       <listOfReactants>
43         <speciesReference species="X0">
44           <stoichiometryMath>
45             <math xmlns="http://www.w3.org/1998/Math/MathML">
46               <cn type="rational"> 3 <sep/> 2 </cn>
47             </math>
48           </stoichiometryMath>
49         </speciesReference>
50       </listOfReactants>
51       ...
52     </reaction>
53     ...
54   </listOfReactions>
55   ...
56 </model>

```

A species can occur more than once in the lists of reactants and products of a given **Reaction** instance. The effective stoichiometry for a species in a reaction is the sum of the stoichiometry values given in the

SpeciesReference objects in the list of products minus the sum of stoichiometry values given in the **SpeciesReference** objects in the list of reactants. A positive value indicates the species is effectively a product and a negative value indicates the species is effectively a reactant. SBML places no restrictions on the effective stoichiometry of a species in a reaction; for example, it can be zero. In the following SBML fragment, the two reactions have the same effective stoichiometry for all their species:

```
<reaction id="x">
  <listOfReactants>
    <speciesReference species="a"/>
    <speciesReference species="a"/>
    <speciesReference species="b"/>
  </listOfReactants>
  <listOfProducts>
    <speciesReference species="c"/>
    <speciesReference species="b"/>
  </listProducts>
</reaction>
<reaction id="y">
  <listOfReactants>
    <speciesReference species="a" stoichiometry="2"/>
  </listOfReactants>
  <listOfProducts>
    <speciesReference species="c"/>
  </listProducts>
</reaction>
```

4.13.4 *ModifierSpeciesReference*

Sometimes a species appears in the kinetic rate formula of a reaction but is itself neither created nor destroyed in that reaction (for example, because it acts as a catalyst or inhibitor). In SBML, all such species are simply called *modifiers* without regard to the detailed role of those species in the model. The **Reaction** object class provides a way to express which species act as modifiers in a given reaction. This is the purpose of the list of modifiers available in **Reaction**. The list contains instances of **ModifierSpeciesReference** object.

As shown in Figure 21 on page 64, the **ModifierSpeciesReference** class inherits the mandatory attribute **species** and optional attributes **id** and **name** from the parent class **SimpleSpeciesReference**; see Section 4.13.2 for their precise definitions.

The value of the **species** attribute must be the identifier of a species defined in the enclosing **Model**; this species is designated as a modifier for the current reaction. A reaction may have any number of modifiers. It is permissible for a modifier species to appear simultaneously in the list of reactants and products of the same reaction where it is designated as a modifier, as well as to appear in the list of reactants, products and modifiers of other reactions in the model.

4.13.5 *KineticLaw*

The **KineticLaw** object class is used to describe the rate at which the process defined by the **Reaction** takes place. As shown in Figure 21 on page 64, **KineticLaw** has elements called **math** and **listOfParameters**, in addition to the attributes and elements it inherits from **Sbase**.

Previous definitions of SBML included two additional attributes called **substanceUnits** and **timeUnits**, which allowed the *substance/time* units of the reaction rate expression to be defined on a per-reaction basis. **SBML Level 2 Version 2 and Version 3 removed** these attributes for several reasons. First, the introduction in SBML Level 2 Version 2 of mass and dimensionless units as possible units of *substance*, coupled with the previous facility for defining the units of each reaction separately and the ability to use non-integer stoichiometries, lead to the possibility of creating a valid model whose reactions nevertheless could not be integrated into a system of equations without outside knowledge for converting the quantities used. (As a simple example, consider if one reaction is defined to be in grams per second and another in moles per second, and species are given in moles: converting from mass to moles would require knowing the molecular mass of the species.) Second, the ability to change the units of a reaction provided the potential of creating unintuitive and difficult-to-reconcile systems of equations, yet the feature added little functionality to SBML.

1 The *absence* of `substanceUnits` does not prevent the definition of any reactions; it only results in requiring
2 the generator of the model to be explicit about any necessary conversion factors. Third, few if any software
3 tools have ever correctly implemented support for `substanceUnits`, which made the use of this `attribute` in
4 a model an impediment to interoperability. Fourth, examination of real-life models revealed that a frequent
5 reason for using `substanceUnits` was to set the units of all reactions to the same set of substance units,
6 which is better achieved by setting the model-wide values of “`substance`”.

7 *The math element*

8 As shown in Figure 21 on page 64, `KineticLaw` has a `element` called `math` for holding a MathML formula
9 defining the rate of the reaction. The expression in `math` may refer to species identifiers, as discussed in
10 Section 3.4.3. The only `Species` identifiers that can be used in `math` are those declared in the lists of
11 reactants, products and modifiers in the `Reaction` object (see Sections 4.13.2, 4.13.3 and 4.13.4). `Parameter`
12 identifiers may be taken from the `KineticLaw`'s list of local parameters (see below) or the parameters defined
13 globally on the `Model` instance.

14 Section 4.13.6 provides important discussions about the meaning and interpretation of SBML “kinetic laws”.

15 *The list of parameters*

16 An instance of `KineticLaw` can contain a list of `one` or more `Parameter` objects (Section 4.9) which define
17 new parameters whose identifiers can be used in the `math` formula. As discussed in Section 3.3.1, reactions
18 introduce local namespaces for parameter identifiers, and within a `KineticLaw` object, a local parameter whose
19 identifier is identical to a global parameter defined in the model takes precedence over that global parameter.

20 The type of `object` used to define a parameter inside `KineticLaw` is the same `Parameter` object class used to
21 define global parameters (Section 4.9). This simplifies the SBML language and reduces the number of unique
22 types of data objects. However, there is a difference between local and global parameters: in the case of
23 parameters defined locally to a `KineticLaw`, there is no means by which the parameter values can be changed.
24 Consequently, such parameters' values are always constant, and the `constant` attribute in their definitions
25 must always have a value of “`true`” (either explicitly or left to its default value).

26 *The sboTerm attribute*

27 `KineticLaw` inherits from `Sbase` the optional attribute called `sboTerm` of type `SBOTerm` (see Section 5). When a
28 value is given to this attribute, the value must be an SBO identifier referring to a term from the `SBO:0000001`,
29 “rate law” vocabulary defined in SBO. The relationship is of the form “the kinetic law *is a* X”, where X is
30 the SBO term. The SBO term chosen should be the most precise (narrow) term that defines the type of
31 reaction rate expression encoded by the `KineticLaw` instance.

32 As explained in Section 4.2.2, there is an interaction between the SBO terms chosen for a model's `KineticLaw`
33 instances and the `sboTerm` value given to the overall `Model` instance. The interaction is as follows. In SBO,
34 each definition of a rate expression itself includes its own SBO reference to a mathematical framework. When
35 a given `KineticLaw` element refers to a rate law defined in SBO, it therefore simultaneously also identifies a
36 framework. If a `Model` instance has a value for `sboTerm`, it implies that all reaction kinetics in the model
37 assume the same mathematical framework. In that case, the frameworks referenced in the SBO definitions
38 of individual reaction rate expressions may be ignored. Consequently, if the mathematical frameworks
39 underlying the definitions of different reactions in a model *should* be observed individually (e.g., because the
40 model mixes frameworks), then there should not be a value assigned to `sboTerm` on the enclosing `Model`.

41 *Example*

42 The following is an example of a `Reaction` object that defines a reaction with identifier J_1 , in which $X_0 \rightarrow S_1$
43 at a rate given by $k \cdot [X_0] \cdot [S_2]$, where S_2 is a catalyst and k is a parameter, and the square brackets
44 symbolizes that the species quantities have units of concentration. The example demonstrates the use of
45 species references and `KineticLaw` objects. The units on the species here are the defaults of *substance/volume*
46 (see Section 4.8), and so the rate expression $k \cdot [X_0] \cdot [S_2]$ needs to be multiplied by the compartment volume
47 (represented by its identifier, “`c1`”) to produce the final units of *substance/time* for the rate expression.

```

1      <model>
2          ...
3          <listOfUnitDefinitions>
4              <unitDefinition id="per_concent_per_time">
5                  <listOfUnits>
6                      <unit kind="litre"/>
7                      <unit kind="mole" exponent="-1"/>
8                      <unit kind="second" exponent="-1"/>
9                  </listOfUnits>
10             </unitDefinition>
11         </listOfUnitDefinitions>
12         ...
13         <listOfSpecies>
14             <species id="S1" compartment="c1" initialConcentration="2.0"/>
15             <species id="S2" compartment="c1" initialConcentration="0.5"/>
16             <species id="X0" compartment="c1" initialConcentration="1.0"/>
17         </listOfSpecies>
18         ...
19         <listOfReactions>
20             <reaction id="J1">
21                 <listOfReactants>
22                     <speciesReference species="X0"/>
23                 </listOfReactants>
24                 <listOfProducts>
25                     <speciesReference species="S1"/>
26                 </listOfProducts>
27                 <listOfModifiers>
28                     <modifierSpeciesReference species="S2"/>
29                 </listOfModifiers>
30                 <kineticLaw>
31                     <math xmlns="http://www.w3.org/1998/Math/MathML">
32                         <apply>
33                             <times/> <ci> k </ci> <ci> S2 </ci> <ci> X0 </ci> <ci> c1 </ci>
34                         </apply>
35                     </math>
36                     <listOfParameters>
37                         <parameter id="k" value="0.1" units="per_concent_per_time"/>
38                     </listOfParameters>
39                 </kineticLaw>
40             </reaction>
41         </listOfReactions>
42         ...
43     </model>

```

4.13.6 Traditional rate laws versus SBML “kinetic laws”

It is important to make clear that a “kinetic law” in SBML is *not* identical to a traditional rate law. The reason is that SBML must support multicompartment models, and the **units normally** used in traditional rate laws as well as some conventional single-compartment modeling packages are problematic when used for defining reactions between multiple compartments.

When modeling species as continuous amounts (e.g., concentrations), the rate laws used are traditionally expressed in terms of *amount of substance concentration per time*, embodying a tacit assumption that reactants and products are all located in a single, constant volume. Attempting to describe reactions between multiple volumes using *concentration/time* (which is to say, *substance/volume/time*) quickly leads to difficulties. Here is an illustration of this. Suppose we have two species pools S_1 and S_2 , with S_1 located in a compartment having volume V_1 , and S_2 located in a compartment having volume V_2 . Let the volume $V_2 = 3V_1$. Now consider a transport reaction $S_1 \rightarrow S_2$ in which the species S_1 is moved from the first compartment to the second. Assume the simplest type of chemical kinetics, in which the rate of the transport reaction is controlled by the activity of S_1 and this rate is equal to some constant k times the activity of S_1 . For the sake of simplicity, assume S_1 is in a diluted solution and thus that the activity of S_1 can be taken to be equal to its concentration $[S_1]$. The rate expression will therefore be $k \cdot [S_1]$, with the units of k being *1/time*. Then:

$$\frac{d[S_2]}{dt} = -\frac{d[S_1]}{dt} = k \cdot [S_1]$$

So far, this looks normal—until we consider the number of molecules of S_1 that disappear from the compartment of volume V_1 and appear in the compartment of volume V_2 . The number of molecules of S_1 (call this n_{S_1}) is given by $[S_1] \cdot V_1$ and the number of molecules of S_2 (call this n_{S_2}) is given by $[S_2] \cdot V_2$. Since our volumes have the relationship $V_2/V_1 = 3$, the relationship above implies that $n_{S_1} = k \cdot [S_1] \cdot V_1$ molecules disappear from the first compartment **per unit of time** and $n_{S_2} = 3 \cdot k \cdot [S_1] \cdot V_1$ molecules appear in the second compartment. In other words, we have created matter out of nothing!

The problem lies in the use of concentrations as the measure of what is transferred by the reaction, because concentrations depend on volumes and the scenario involves multiple unequal volumes. The problem is not limited to using concentrations or volumes; the same problem also exists when using density, i.e., *mass/volume*, and dependency on other spatial distributions (i.e., areas or lengths). What must be done instead is to consider the number of “items” being acted upon by a reaction process irrespective of their distribution in space (volume, area or length). An “item” in this context may be a molecule, particle, mass, or other “thing”, as long as the substance measurement is independent of the size of the space in which the items are located.

For the current example, the expressions in terms of n_{S_1} and n_{S_2} are straightforward:

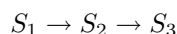
$$\frac{dn_{S_2}}{dt} = -\frac{dn_{S_1}}{dt} = k \cdot [S_1] \cdot V_1$$

Given numbers of items, it is then easy to recover concentrations by dividing the item counts of each species by the volume of the compartment in which the species is located: $[S_1] = n_{S_1}/V_1$ and $[S_2] = n_{S_2}/V_2$.

The need to support multicompartment models requires that the reaction rates in SBML to be expressed in terms of *substance/time*, rather than the more typical *substance/size/time*. As a result, modelers and software tools in general cannot insert textbook rate laws unmodified as the rate expression in the **math element** of a **KineticLaw**. The unusual term “kinetic law” was chosen to alert users to this difference. We explain the general principles of converting rate laws in the following paragraphs.

Basic cases

Let us expand the simple example above by adding a second reaction, to create the system



with the left-hand reaction’s rate (call this r_1) being given as $k_1 \cdot [S_1]$ and the rate of the right-hand reaction (call it r_2) as $k_2 \cdot [S_2]$. Also assume each species is located in a different compartment:

| | | |
|-------|------------------------------|-------------------|
| S_1 | located in compartment C_1 | with volume V_1 |
| S_2 | located in compartment C_2 | with volume V_2 |
| S_3 | located in compartment C_3 | with volume V_3 |

As before, converting the rate of the first reaction ($S_1 \rightarrow S_2$) to units of *substance/time* in this case is a simple matter of multiplying by the volume of the compartment where the reactants are located, leading to the following SBML rate formula:

$$R_1 = r_1 \cdot V_1 = -k_1 \cdot [S_1] \cdot V_1$$

The second rate expression becomes

$$R_2 = r_2 \cdot V_2 = -k_2 \cdot [S_2] \cdot V_2$$

The units of k_1 and k_2 are *1/time* (often *1/sec*, but not necessarily), as is typical for reactions that are first-order in one reactant. The expressions R_1 and R_2 are what would be written in **KineticLaw math** definitions for the two reactions in this system. The formulas give the speed of each reaction in terms of the substance change over time. The reader of the SBML model needs to combine the individual contributions of each

1 reaction to construct equations for the *overall* rates of change of each species in the model using these
 2 expressions. In terms of differential equations, these are:

$$\begin{aligned} \frac{dn_{S_1}}{dt} &= -R_1 &= -k_1 \cdot [S_1] \cdot V_1 \\ \frac{dn_{S_2}}{dt} &= +R_1 - R_2 &= +k_1 \cdot [S_1] \cdot V_1 - k_2 \cdot [S_2] \cdot V_2 \\ \frac{dn_{S_3}}{dt} &= +R_2 &= +k_2 \cdot [S_2] \cdot V_2 \end{aligned}$$

4 To recover the concentration values, we add the following to the system of equations:

$$\begin{aligned} [S_1] &= n_{S_1}/V_1 \\ [S_2] &= n_{S_2}/V_2 \\ [S_3] &= n_{S_3}/V_3 \end{aligned}$$

6 Note that this formulation works properly even if the compartment sizes V_1 , V_2 and V_3 vary during simulation.

7 Extrapolating from this example, we can now provide a general approach to translating a system of reactions
 8 involving species located in multiple compartments, for the restricted case where all reactants of any given
 9 reaction are in the same compartment (but where the compartments involved may be different for each
 10 reaction). For a species S_i involved in m reactions whose rates are given (in “textbook” form, without
 11 volume adjustments) as r_1, r_2, \dots, r_m , where the reactants of r_j are located in the compartment of size V_j ,

$$\begin{aligned} \frac{dn_{S_i}}{dt} &= \text{sign}_1 \cdot \text{stoich}_1 \cdot r_1 \cdot V_1 \\ &+ \text{sign}_2 \cdot \text{stoich}_2 \cdot r_2 \cdot V_2 \\ &+ \dots \\ &+ \text{sign}_m \cdot \text{stoich}_m \cdot r_m \cdot V_m \end{aligned} \tag{7}$$

$$[S_i] = n_{S_i}/V_i$$

13 In Equation (7), each term sign_j is “−” if S_i is a reactant in r_j and “+” if it is a product, and each term
 14 stoich_j is the stoichiometry of S_i in reaction r_j . Letting $R_j = r_j \cdot V_j$ represent the form of the rate expressions
 15 as they would be written in the **KineticLaw** math elements, then we can equivalently write

$$\begin{aligned} \frac{dn_{S_i}}{dt} &= \text{sign}_1 \cdot \text{stoich}_1 \cdot R_1 \\ &+ \text{sign}_2 \cdot \text{stoich}_2 \cdot R_2 \\ &+ \dots \\ &+ \text{sign}_m \cdot \text{stoich}_m \cdot R_m \end{aligned} \tag{8}$$

$$[S_i] = n_{S_i}/V_i$$

17 This approach preserves the use of concentration terms within the reaction rate expressions so that the core
 18 of those rate expressions can be ordinary rate laws. This is important when modeling species as continuous
 19 quantities, because most textbook rate expressions are measured in terms of concentrations, and most
 20 rate constants have units involving concentration rather than item counts. For example, the second-order
 21 rate constant in a mass-action rate law has units of $1/(M \cdot s)$, which is to say, *volume/(substance·time)*;
 22 this constant is then multiplied by two concentration terms. Reaction definitions in SBML models can be
 23 constructed by taking such expressions and multiplying them by the volume of the compartment in which
 24 the reactants are located. By contrast, if we were to simply replace concentrations of species by item counts
 25 in such rate laws, it would in most cases be incorrect. At the very least, the constants in the equations would
 26 need to be converted in some way to make such expressions valid.

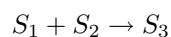
1 The preceding discussion of problems involving rate laws concerns modeling approaches that use continuous
2 quantities. There is an alternative approach to modeling that instead treats species as discrete popula-
3 tions (Wilkinson, 2006). In those cases, the rate expressions already use substance or item counts rather
4 than concentrations and there is no need to convert them.

5 A full SBML example of translating a complete multicompartmental model into ODEs is given in Section 7.7.
6 An example of translating a discrete model is given in Section 7.3.

7 *Advanced cases*

8 The explanation above applies to reactions where all of the reactants are in the same compartment. What
9 about cases where two or more reactant species are in separate compartments?

10 This is a more difficult situation, and the guidelines described above for Equation (7) cannot always be
11 applied because there will be more than one compartment size term by which the core rate expression needs
12 to be multiplied. Unfortunately, there is often no straightforward way to mechanically convert such models
13 without requiring a more significant change to the reaction rate expression. An example will help illustrate
14 the difficulty. Suppose we have a simple reaction system consisting of only



16 where S_1 , S_2 and S_3 are each located in separate compartments with volumes V_1 , V_2 and V_3 , and the rate
17 expression is given as $k \cdot [S_1] \cdot [S_2]$. (In reality, one would not use such a rate law in this case, but for the
18 sake of this example, let us ignore the fact that a mass-action rate law would actually involve an assumption
19 that all reactants are in a well-mixed solution.) A straightforward examination of the possibilities eventually
20 leads to the conclusion that in order to take account of the multiple volumes, the rate expressions in terms
21 of *substance/time* have to be written as

$$22 \quad \frac{dn_{S_i}}{dt} = -k'(V_1, V_2) \cdot ([S_1] \cdot V_1) \cdot ([S_2] \cdot V_2)$$

23 The crux of the problem is that the new factor $k'(V_1, V_2)$ is not the original k ; to make the overall units of
24 the expression work out, $k'(V_1, V_2)$ must be a function of the volumes, and its value must change if V_1 or V_2
25 changes. It is no longer a standard rate constant. In an SBML model, it is easy to define an **AssignmentRule**
26 to compute the value of k' based on k , V_1 , V_2 , and possibly other variables in the system as needed, but only
27 the modeler can determine the proper formula for their particular modeling situation. (For example, the
28 modeler may know that in their hypothesized physical system, the reaction actually takes place completely
29 in one or the other compartment and therefore the factor should be designed accordingly, or perhaps the
30 reaction takes place on a membrane between compartments and a scaling factor based on the area should
31 be used.)

32 Thus, although these models can be represented in SBML, constructing the correct rate expression in terms
33 of *substance/time* units depends on problem-specific knowledge, and we cannot provide a general recipe.

34 *Mixed species units*

35 The discussion so far has assumed that all of the species appearing in a given reaction's rate expression
36 had the same units, whether they be concentration or amounts or other. However, **Species** objects can each
37 declare units separately. What happens then?

38 It is important to realize that implicit conversions in this situation are not defined in SBML. A species
39 identifier appearing in a mathematical expression has the units attributed to that species (see Section 4.8.5
40 for a definition of the *the units of the species*). If a reaction contains species each having different units, the
41 creator of the SBML model must explicitly incorporate the necessary conversions to make the units of the
42 rate expression consistent. The most appropriate way is to include the conversion factor as part of the value
43 of **stoichiometry** or **stoichiometryMath** in the **SpeciesReference** for that species.

1 An example may help illustrate this. Suppose we have a system of two biochemical reactions with mass-action
 2 kinetics, written in the traditional form



3
 4 with the units of k_1 being $\text{liter}^2\text{s}^2\text{mol}^{-2}$ and k_2 's units being s^{-1} . Assume the reactions take place in a
 5 single compartment of volume V , but now let us throw a wrench into the problem: suppose that the species
 6 in the model are defined with mixed units as follows:

7 A is in millimoles per litre
 8 B is in grams per litre
 9 C is in items per litre

10 When biochemical reaction equations of the form (9) above are written, the units of species' quantities
 11 usually are assumed to be the same, and therefore the stoichiometries in the reaction equations (9) represent
 12 simple ratios between the quantities of the species in those units. (“One mole of this and two moles of that
 13 react to produce one mole of that other.”) If instead the quantities of the species are given in mixed units,
 14 as in the present example, the quantities not only need to be in proper stoichiometric relationships, the
 15 units also have to be made consistent. In SBML, this is done by appropriately setting the **stoichiometry**
 16 **attribute** value in the species references of the lists of reactants, products and modifiers in a **Reaction**. This
 17 then permits a properly balanced system of equations to be constructed for each species' rate of change of
 18 quantity.

19 In the present example, the SBML “kinetic law” formulas for the reaction rates will be written as **conventional**
 20 **mass-action reaction rates adjusted for volume as described previously**,

$$\begin{aligned}
 R_1 &= k_1' \cdot [A] \cdot [B]^2 \cdot V \\
 R_2 &= k_2' \cdot [C] \cdot V
 \end{aligned}$$

21
 22 where $k_2' = k_2$ for this particular example but k_1' will depend on other units as described in the paragraphs
 23 below. When these formulas are combined into overall expressions for the rates of change of A , B , and C ,
 24 the result is

$$\begin{aligned}
 dA/dt &= -a_1 \cdot R_1 + a_2 \cdot R_2 \\
 dB/dt &= -b_1 \cdot R_1 \\
 dC/dt &= +c_1 \cdot R_1 - c_2 \cdot R_2
 \end{aligned}$$

25
 26 where

27 a_1 is the SBML stoichiometry of A in reaction 1
 28 a_2 is the SBML stoichiometry of A in reaction 2
 29 b_1 is the SBML stoichiometry of B in reaction 1
 30 c_1 is the SBML stoichiometry of C in reaction 1
 31 c_2 is the SBML stoichiometry of C in reaction 2

32 We use the term *SBML stoichiometries* to highlight the fact that in this example involving mixed-units
 33 species, the values may not be identical to the biochemical stoichiometries in the reaction equations (9).
 34 And just what are the *SBML stoichiometries*? In the kind of mixed-units situation faced in this example,
 35 they must encompass both the biochemical stoichiometries and any necessary unit conversions. Thus, letting
 36 m_B stand for the molecular mass of B :

37 $a_1 = 1000$ (in each reaction event, 1 mole of A is consumed, with A expressed in millimoles)
 38 $a_2 = 3000$ (in each reaction event, 3 moles of A are produced, with A expressed in millimoles)
 39 $b_1 = 2 \cdot m_B$ (in each reaction event, 2 moles of B are consumed, with B expressed in grams)
 40 $c_1 = 6.02 \cdot 10^{23}$ (in each reaction event, 1 mole of C is produced, expressed as item counts)
 41 $c_2 = 6.02 \cdot 10^{23}$ (in each reaction event, 1 mole of C is consumed, expressed as item counts)

42 and $k_1' = k_1 \cdot 10^{-3} \cdot 1/m_B^2$.

1 What happens if the definition of the SBML built-in unit “**substance**” is changed in the model to be
2 millimoles? Then the stoichiometries must be changed to the following:

$$\begin{aligned} 3 \quad a_1 &= 1 && \text{(in each reaction event, 1 millimole of } A \text{ is consumed, expressed in millimoles)} \\ 4 \quad a_2 &= 3 && \text{(in each reaction event, 3 millimoles of } A \text{ are produced, expressed in millimoles)} \\ 5 \quad b_1 &= 2 \cdot 10^{-3} \cdot m_B && \text{(in each reaction event, 2 millimoles of } B \text{ are consumed, expressed in grams)} \\ 6 \quad c_1 &= 6.02 \cdot 10^{20} && \text{(in each reaction event, 1 millimole of } C \text{ is produced, expressed as item counts)} \\ 7 \quad c_2 &= 6.02 \cdot 10^{20} && \text{(in each reaction event, 1 millimole of } C \text{ is consumed, expressed as item counts)} \end{aligned}$$

$$8 \quad \text{and } k_1' = k_1 \cdot (10^3/m_B)^2.$$

9 What happens if instead the definition of the SBML built-in unit “**substance**” is changed in the model to
10 be “**item**”? Then the stoichiometries must be changed to the following:

$$\begin{aligned} 11 \quad a_1 &= 1/(6.02 \cdot 10^{20}) && \text{(in each reaction event, 1 item of } A \text{ is consumed, expressed in millimoles)} \\ 12 \quad a_2 &= 3/(6.02 \cdot 10^{20}) && \text{(in each reaction event, 3 items of } A \text{ are produced, expressed in millimoles)} \\ 13 \quad b_1 &= 2 \cdot 1/(6.02 \cdot 10^{23}) \cdot m_B && \text{(in each reaction event, 2 items of } B \text{ are consumed, expressed in grams)} \\ 14 \quad c_1 &= 1 && \text{(in each reaction event, 1 item of } C \text{ is produced, expressed as item counts)} \\ 15 \quad c_2 &= 1 && \text{(in each reaction event, 1 item of } C \text{ is consumed, expressed as item counts)} \end{aligned}$$

$$16 \quad \text{and } k_1' = k_1 \cdot 6.02 \cdot 10^{20} \cdot \left(\frac{6.02 \cdot 10^{23}}{m_B} \right)^2.$$

17 And finally, what happens if the definition of the SBML built-in unit “**substance**” is changed in the model
18 to be “**gram**”? Then the stoichiometries must be changed again, to the following:

$$\begin{aligned} 19 \quad a_1 &= 1000 \cdot m_A && \text{(in each reaction event, 1 gram of } A \text{ is consumed, expressed in millimoles)} \\ 20 \quad a_2 &= 3000 \cdot m_A && \text{(in each reaction event, 3 grams of } A \text{ are produced, expressed in millimoles)} \\ 21 \quad b_1 &= 2 && \text{(in each reaction event, 2 grams of } B \text{ are consumed, expressed in grams)} \\ 22 \quad c_1 &= 6.02 \cdot 10^{23} \cdot m_C && \text{(in each reaction event, 1 gram of } C \text{ is produced, expressed as item counts)} \\ 23 \quad c_2 &= 6.02 \cdot 10^{23} \cdot m_C && \text{(in each reaction event, 1 gram of } C \text{ is consumed, expressed as item counts)} \end{aligned}$$

$$24 \quad \text{and } k_1' = k_1 \cdot 10^{-3} \cdot m_A, \text{ where } m_A, m_B, \text{ and } m_C \text{ are the molecular masses of species } A, B, \text{ and } C, \text{ respectively.}$$

25 In all of these cases, one straightforward approach to computing the value of k_1' is to use an **InitialAssignment**
26 construct for setting the value of k_1' to the result of a formula such as one of those given in the examples
27 above. And finally, note that if the species units *were* the same throughout (and in most models they are),
28 the unit conversion aspects of the SBML stoichiometries would become unity, leaving only the expected
29 biochemical stoichiometry values. Isn't that nice?

30 **4.13.7 Use of reaction identifiers in mathematical expressions**

31 The value of the **id** attribute of a **Reaction** can be used as the content of a **ci** element in MathML formulas
32 elsewhere in the model. Such a **ci** element or symbol represents the rate of the given reaction as given by
33 the reaction's **KineticLaw** object. The symbol has the units of *substance/time*.

34 A **KineticLaw** object in effect forms an assignment statement assigning the evaluated value of the **math**
35 element to the symbol value contained in the **Reaction id** attribute. No other object can assign a value
36 to such a reaction symbol; i.e., the **variable** attributes of **InitialAssignment**, **RateRule**, **AssignmentRule** and
37 **EventAssignment** objects cannot contain the value of a **Reaction id** attribute.

38 The combined set of **InitialAssignment**, **AssignmentRule** and **KineticLaw** objects form a set of assignment
39 statements that should be considered as a whole. The combined set of assignment rules should not contain
40 algebraic loops: a chain of dependency between these statements should terminate. (More formally, consider
41 the directed graph of assignment statements where nodes are statements and directed arcs exist for each
42 occurrence of a symbol in a assignment statement **math** element. The directed arcs start from the statement
43 defining the symbol to the statements that contain the symbol in their **math** elements. Such a graph must
44 be acyclic.) Examples of valid and invalid set of assignment statements are given in Section 4.11.5.

4.14 Events

Model has an optional list of **Event** objects that describe the time and form of explicit instantaneous discontinuous state changes in the model. For example, an event may describe that one species quantity is halved when another species quantity exceeds a given threshold value.

An **Event** object defines when the event can occur, the variables that are affected by the event, and how the variables are affected. The effect of the event can optionally be delayed after the occurrence of the condition which invokes it. The operation of an **event** is divided into two phases (even when the event is not delayed): one when the event is *fired* and the other when the event is *executed*. The **Event** type is defined in Figure 22. The object classes **Event**, **Trigger**, **Delay** and **EventAssignment** are derived from **Sbase** (see Section 3.2). An example of a model which uses events is given below.

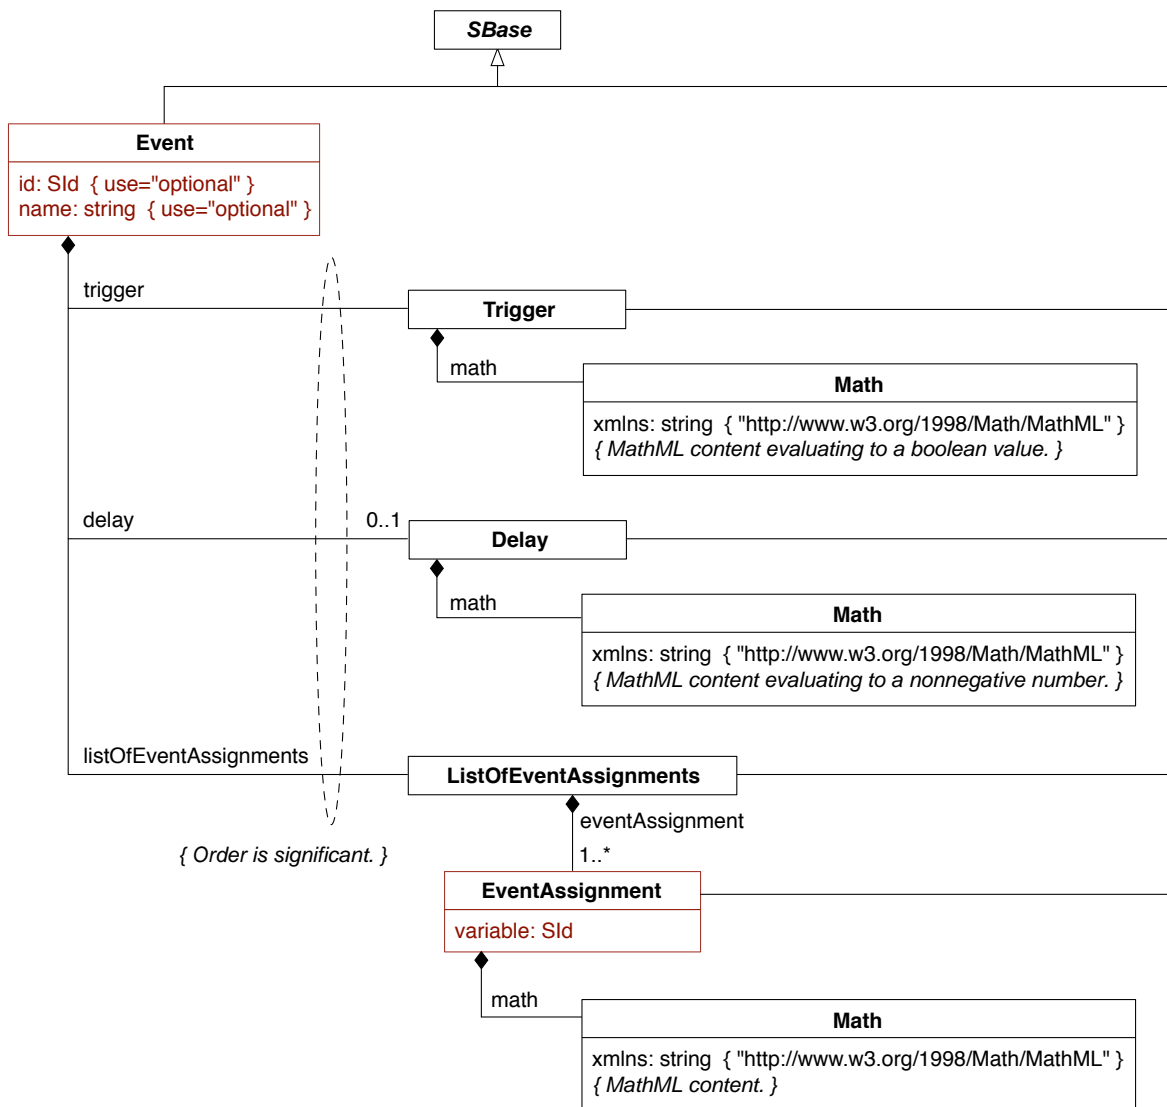


Figure 22: The definitions of **Event**, **Trigger**, **Delay** and **EventAssignment**, and the container class **ListOfEventAssignment**.

4.14.1 Event

An **Event** definition has three required parts: an identifier, a trigger condition and at least one **EventAssignment**. In addition, an event can include an optional delay. These features of **Event** are described below.

Previous definitions of **Event** in SBML Level 2 included an additional attribute called **timeUnits**, which allowed the time units of the **Delay** to be set explicitly. SBML Level 2 Version 3 removes this attribute for several reasons. First, the ability to change the time units of the delay time of an **Event** to be different from the units of time for the whole model meant that computing an **Event**'s time of triggering and its delay might have to be done using two different sets of units—a potential source of overlooked errors. Second, the ability to redefine the units of time for the delay of an **Event** became inconsistent with the lack of such an attribute on other SBML Level 2 Version 3 components involving an element of time, such as **RateRule** and **KineticLaw**. On balance, the **timeUnits** feature was judged to add needless complexity and inconsistency for little gain in functionality.

The id and name attributes

As with most components in SBML, an **Event** has **id** and **name** attributes, but in the case of **Event**, both are optional. These attributes operate in the manner described in Section 3.3.

The optional sboTerm attribute on Event

As with all SBML components derived from **Sbase**, an **Event** has an optional attribute **sboTerm** of type **SBOTerm** (see Sections 3.1.9 and 5). When a value is given to this attribute, it must be a valid term derived from SBO:0000231, “event” in SBO. The **Event** should have an “is a” relationship with the SBO term, and the term should be the most precise (narrow) term that captures the meaning of the event in the model.

As discussed in Section 5, SBO labels are optional information on a model. Applications are free to ignore **sboTerm** values. A model must be interpretable without the benefit of SBO labels.

4.14.2 Trigger

As shown in Figure 22, the **trigger** element of an **Event** must contain exactly one object of class **Trigger**. This object contains one **math** element containing a MathML expression. The expression must evaluate to a value of type **boolean**. The exact moment at which the expression evaluates to “**true**” is the time point when the **Event** is *fired*.

An event only fires when its **Trigger** expression makes the transition in value from “**false**” to “**true**”. The event will also fire at any future time points when the **trigger** make this transition; in other words, an event can fire multiple times during a simulation if its trigger condition makes the transition from “**false**” to “**true**” more than once.

An important question is whether an event can fire prior to, or at, initial simulation time, i.e., $t \leq 0$. The answer is no: an event can only be triggered immediately after initial simulation time i.e., $t > 0$.

The optional sboTerm attribute on Trigger

As with all SBML components derived from **Sbase**, **Trigger** inherits the optional attribute **sboTerm** of type **SBOTerm** (see Sections 3.1.9 and 5). When a value is given to this attribute, it must be a valid term representing a mathematical expression (i.e., a term chosen from the SBO:0000064, “mathematical expression” hierarchy). The formula in the **Trigger**'s **math** expression should have an “is a” relationship with the SBO term, and the term should capture most precise (narrow) meaning of the mathematical formula of the trigger.

4.14.3 Delay

As shown in Figure 22, an **Event** object can contain an optional **delay** element of class **Delay**. The **Delay** is derived from **Sbase** and contains a mathematical formula stored in **math**. The formula is used to compute the length of time between when the event has *fired* and when the event's assignments (see below) are actually *executed*. If no delay is present on a given **Event**, a time delay of zero is assumed.

1 The expression in **math** must be evaluated at the time the rule is *fired*. The expression must always evaluate
2 to a nonnegative number (otherwise, a nonsensical situation could arise where an event is defined to fire
3 before it is triggered!).

4 *Units of delay expressions*

5 The units of the numerical value computed by a **Delay** instance's **math** expression must match the model's
6 units of *time* (meaning the definition of the “**time**” units in the model; see Section 4.4.3). Note that, as
7 in other cases of MathML expressions in SBML, units are *not* predefined or assumed. As discussed in
8 Section 3.4.10, literal numbers (i.e., numbers enclosed in MathML **cn** elements) or expressions containing
9 only literal numbers and/or **Parameter** objects without declared units, are considered to have unspecified
10 units. In such cases, the correspondence between the needed units and the (unknown) units of the **Delay**
11 **math** expression cannot be proven, and while such expressions are not considered inconsistent, all that can
12 be assumed by model interpreters (whether software or human) is that the units *may* be consistent.

13 The following **Event** example fragment helps illustrate this:

```
14 <model>  
15   ...  
16   <listOfEvents>  
17     <event>  
18       ...  
19       <delay>  
20         <math xmlns="http://www.w3.org/1998/Math/MathML">  
21           <cn> 10 </cn>  
22         </math>  
23       </delay>  
24       ...  
25     </event>  
26   </listOfEvents>  
27   ...  
28 </model>
```

29 Note the “<cn> 10 </cn>” within the mathematical formula has no specified units. The model is not invalid
30 because of this, but a recipient of the model may justifiably be concerned about what “10” really means. (Ten
31 seconds? What if the global units of time on the model were changed from seconds to milliseconds? Would
32 the modeler remember to change “10” to “10 000”?) As discussed elsewhere, leaving units unspecified may
33 prevent software tools from performing complete validation and other useful operations such as global unit
34 conversions. A better approach is to avoid literal numbers and instead use an approach such as defining a
35 parameter with declared units, as in the following modified version of the example fragment:

```
36 <model>  
37   ...  
38   <listOfParameters>  
39     <parameter id="transcriptionDelay" value="10" units="time"/>  
40   </listOfParameters>  
41   ...  
42   <listOfEvents>  
43     <event>  
44       ...  
45       <delay>  
46         <math xmlns="http://www.w3.org/1998/Math/MathML">  
47           <ci> transcriptionDelay </ci>  
48         </math>  
49       </delay>  
50       ...  
51     </event>  
52   </listOfEvents>  
53   ...  
54 </model>
```

1 The optional `sboTerm` attribute on **Delay**

2 As with all SBML components derived from **Sbase**, **Delay** inherits an optional `sboTerm` attribute of type
3 **SBOTerm** (see Sections 3.1.9 and 5). When a value is given to this attribute, it must be a valid term derived
4 from the **SBO:0000064**, “mathematical expression” hierarchy in SBO. The **Delay** formula should have an “is
5 a” relationship with the chosen SBO term, and the term should be the most precise (narrow) term that
6 captures the meaning of the delay expression.

7 4.14.4 **EventAssignment**

8 **Event** contains a mandatory element called `listOfEventAssignments`, of class **ListOfEventAssignment**. In
9 every instance of an event definition in a model, the object’s `listOfEventAssignments` element must have
10 a non-empty list of one or more **eventAssignment** elements of class **EventAssignment**. The object class
11 **EventAssignment** has one required attribute, `variable`, and a required element, `math`. Being derived from
12 **Sbase**, it also has all the usual attributes and elements of its parent class.

13 An “event assignment” has effect when the event is *executed*; that is, at the end of any given delay period
14 (if given) following the moment that the **Event** is triggered. See Section 4.14.6 below for more information
15 about events and event assignments in SBML.

16 The variable *attribute*

17 The `variable attribute` is of type **SId** and can contain the identifier of a **Compartment**, **Species** or **Parameter**
18 instance defined in the model. When the event fires, the value of the model component identified by `variable`
19 is changed by the **EventAssignment** to the value computed by the `math element`; that is, a species’ quantity,
20 compartment’s size or parameter’s value are reset to the value computed by `math`.

21 Certain restrictions are placed on what can appear in `variable`:

- 22 • The object identified by the value of the `variable attribute` must not have its `constant attribute` set
23 to or default to “true”. (Constants cannot be affected by events.)
- 24 • The `variable attribute` must not contain the identifier of a reaction; only species, compartment and
25 parameter values may be set by an **Event**.
- 26 • The value of every `variable attribute` must be unique among the set of **EventAssignment** objects
27 within a given **Event** instance. In other words, a single event cannot have multiple **EventAssignments**
28 assigning the same variable. (All of them would be performed at the same time, when that particular
29 **Event** triggers, resulting in indeterminacy.) Separate **Event** instances can refer to the same variable.
- 30 • A variable cannot be assigned a value in an **EventAssignment** object instance and also be assigned a
31 value by an **AssignmentRule**, i.e., the value of the `variable attribute` in an **EventAssignment** instance
32 cannot be the same as the value of a `variable attribute` in a **AssignmentRule** instance. (Assignment
33 rules hold at all times, therefore it would be inconsistent to also define an event that reassigns the
34 value of the same variable.)

35 The optional `sboTerm` attribute on **EventAssignment**

36 **EventAssignment** has an optional `sboTerm attribute` of type **SBOTerm** (see Sections 3.1.9 and 5). When a
37 value is given to this `attribute`, it must be a valid SBO term identifier referring to a mathematical expression
38 (i.e., terms derived from **SBO:0000064**, “mathematical expression”). The **EventAssignment** should have an
39 “is a” relationship with the SBO term, and the term should be the most precise (narrow) term that captures
40 the form of the assignment formula in the model.

41 **EventAssignment**’s `math`

42 The `math element` contains a MathML expression that defines the new value of the variable. This expression
43 is evaluated when the **Event** is *fired* but the variable only acquires the result or new value when the **Event**
44 is *executed*. The order of the **EventAssignment** objects is not significant; the effect of one assignment cannot

affect the result of another assignment. The identifiers occurring in the MathML `ci` attributes of the **EventAssignment** object represent the value of the identifier at the point when the **Event** is *fired*.

Units of the math formula in **EventAssignment**

In all cases, as would be expected, the units of the formula contained in the `math` element of **EventAssignment** must be consistent with the units of the object identified by the `variable` attribute. More precisely:

- *In the case of a species*, an **EventAssignment** sets the referenced species' quantity (*concentration or amount of substance*) to the value determined by the formula in `math`. The units of the `math` formula must be identical to the *units of the species* as defined in Section 4.8.5.
- *In the case of a compartment*, an **EventAssignment** sets the referenced compartment's size to the size determined by the formula in `math`. The overall units of the formula must be identical to the units specified for the size of the compartment identified by the value of the **EventAssignment**'s `variable` attribute. (See Section 4.7.5 for an explanation of how the units of the compartment's size are determined.)
- *In the case of a parameter*, an **EventAssignment** sets the referenced parameter's value to that determined by the formula in `math`. The overall units of the formula must be identical to the units defined for the parameter identified by the value of the **EventAssignment**'s `variable` attribute. (See Section 4.9.3 for an explanation of how the units of the parameter are determined.)

Note that the formula placed in the `math` element has no assumed units. The consistency of the units of the formula, and the units of the entity which the assignment affects, must be explicitly established just as in the case of the value of `delay`. An approach similar to the one discussed in Section 4.14.3 may be used.

4.14.5 Example Event definitions

A example of an **Event** object follows. This structure makes the assignment $k_2 = 0$ at the point when $P_1 \leq P_2$:

```

<event>
  ...
  <listOfUnitDefinitions>
    <unitDefinition id="per_second">
      <listOfUnits>
        <unit kind="second" exponent="-1" multiplier="1" offset="0"/>
      </listOfUnits>
    </unitDefinition>
  </listOfUnitDefinitions>
  ...
  <listOfParameters>
    ...
    <parameter id="k2" value="0.05" units="per_second"/>
    <parameter id="k2reset" value="0.0" units="per_second"/>
    ...
  </listOfParameters>
  ...
  <listOfEvents>
    <event>
      <trigger>
        <math xmlns="http://www.w3.org/1998/Math/MathML">
          <apply>
            <leq/>
            <ci> P_1 </ci>
            <ci> P_2 </ci>
          </apply>
        </math>
      </trigger>
      <listOfEventAssignments>
        <eventAssignment variable="k2">
          <math xmlns="http://www.w3.org/1998/Math/MathML">

```

```

1           <ci> k2reset </ci>
2         </math>
3       </eventAssignment>
4     </listOfEventAssignments>
5   </event>
6 </listOfEvents>
7   ...
8 </model>

```

A complete example of a model using events is given in Section 7.10.

4.14.6 Detailed semantics of events

The description of events above describes the action of events in isolation from each other. This section describes how events interact.

Events whose **trigger** expression is true at the start of a simulation do not *fire* at the start of the simulation ($t = 0$). Events *fire* only when the trigger *becomes* true, i.e., the trigger expression transitions from false to true, which cannot happen at $t = 0$ but can happen at $t > 0$.

Any transition of a **trigger** expression from “false” to “true” will cause an **event** to *fire*. Consider an **event** E with delay d where the **trigger** expression makes a transition from false to true at times t_1 and t_2 . The **EventAssignment** object will have effect at $t_1 + d$ and $t_2 + d$ irrespective of the relative times of t_1 and t_2 . For example events can “overlap” so that $t_1 < t_2 < t_1 + d$ still causes an event assignments to occur at $t_1 + d$ and $t_2 + d$.

It is possible for events to *fire* other events, i.e., an event assignment can cause an event to *fire*, therefore it is possible for a model to be entirely encoded in **Event** objects.

It is entirely possible for two events to be *executed* simultaneously in simulated time. It is assumed that, although the precise time at which these events are *executed* is not resolved beyond the given point in simulated time, the order in which the events occur is resolved. This order can be significant in determining the overall outcome of a given simulation. SBML Level 2 does not define the algorithm for determining this order (the tie-breaking algorithm). As a result, the **outcomes** of simulations involving events may vary when simultaneous events occur during simulation. It is anticipated that future versions or levels of SBML will define a specific set of tie-breaking algorithms and a mechanism for models to indicate which algorithm should be applied during simulation.

Despite the absence of a specific tie-breaking algorithm, SBML event simulation is constrained as follows. When an event X *fires* another event Y and event Y has zero delay then event Y is added to the existing set of simultaneous events that are pending *execution*. Events such as Y do not have a special priority or ordering within the tie-breaking algorithm. Events X and Y form a cascade of events at the same point in simulation time. All events in a model are open to being in a cascade. The position of an event in the event list does not affect whether it can be in the cascade: Y can be triggered whether it is before or after X in the list of events. A cascade of events can be infinite (never terminate). When this occurs a simulator should indicate this has occurred; i.e., it is incorrect for the simulator to arbitrarily break the cascade and continue the simulation without at least indicating the infinite cascade occurred. A variable can change more than once when processing simultaneous events at simulation time t . The model behavior (output) for such a variable is the value of the variable at the end of processing all the simultaneous events at time t .

5 The Systems Biology Ontology and the `sboTerm` attribute

The values of `id` attributes on SBML components allow the components to be cross-referenced within a model. The values of `name` attributes on SBML components provide the opportunity to assign them meaningful labels suitable for display to humans (Section 3.3). The specific identifiers and labels used in a model necessarily must be unrestricted by SBML, so that software and users are free to pick whatever they need. However, this freedom makes it more difficult for software tools to determine, without additional human intervention, the semantics of models more precisely than the semantics provided by the SBML object classes defined in other sections of this document. For example, there is nothing inherent in a parameter with identifier “k” that would indicate to a software tool it is a first-order rate constant (if that’s what “k” happened to be in some given model). However, one may need to convert a model between different representations (e.g., Henri-Michaelis-Menten vs. elementary steps), or to use it with different modelling approaches (discrete or continuous). One may also need to relate the model components with other description formats, using a deeper semantics. Although an advanced software tool *might* be able to deduce the semantics of some model components through detailed analysis of the kinetic rate expressions and other parts of the model, this quickly becomes infeasible for any but the simplest of models.

An approach to solving this problem is to associate model components with terms from carefully curated controlled vocabularies (CVs). This is the purpose of the optional `sboTerm` attribute provided on the SBML class `Sbase`. The `sboTerm` attribute always refers to terms belonging to the Systems Biology Ontology (SBO, <http://biomodels.net/SBO/>). In this section, we discuss the `sboTerm` attribute, SBO, the motivations and theory behind their introduction, and guidelines for their use.

SBO is not part of SBML; it is being developed separately, to allow the modeling community to evolve the ontology independently of SBML. However, the terms in the ontology are being designed with SBML components in mind, and are classified into subsets that can be directly related with SBML components such as reaction rate expressions, parameters, and a few others. The use of `sboTerm` attributes is optional, and the presence of `sboTerm` on an element does not change the way the model is interpreted. Annotating SBML elements with SBO terms adds additional semantic information that may be used to convert the model into another model, or another format. Although SBO support provides an important source of information to understand the meaning of a model, software does not need to support `sboTerm` to be considered SBML-compliant.

5.1 Principles

Labeling model components with terms from shared controlled vocabularies allows a software tool to identify each component using identifiers that are not tool-specific. An example of where this is useful is the desire by many software developers to provide users with meaningful names for reaction rate equations. Software tools with editing interfaces frequently provide these names in menus or lists of choices for users. However, without a standardized set of names or identifiers shared between developers, a given software package cannot reliably interpret the names or identifiers of reactions used in models written by other tools.

The first solution that might come to mind is to stipulate that certain common reactions always have the same name (e.g., “Michaelis-Menten”), but this is simply impossible to do: not only do humans often disagree on the names themselves, but it would not allow for correction of errors or updates to the list of predefined names except by issuing new releases of the SBML specification—to say nothing of many other limitations with this approach. Moreover, the parameters and variables that appear in rate expressions also need to be identified in a way that software tools can interpret mechanically, implying that the names of these entities would also need to be regulated.

The Systems Biology Ontology (SBO) provides vocabularies for identifying common reaction rate expressions, common participant types and roles in reactions, common parameter types and their roles in rate expressions, common modeling frameworks (e.g., “continuous”, “discrete”, etc.), and common types of species and reactions. The relationship implied by an `sboTerm` value on an SBML model component is “is a”: the thing defined by that SBML component “is an” instance of the thing defined in SBO by the term having the identifier given by the `sboTerm` attribute value. By adding SBO term references on the components of a

1 model, a software tool can provide additional details using independent, shared vocabularies that can enable
2 *other* software tools to recognize precisely what the component is meant to be. Those tools can then act on
3 that information. For example, if the SBO identifier **SBO:0000049** is assigned to the concept of “first-order
4 irreversible mass-action kinetics, continuous framework”, and a given **KineticLaw** object in a model has an
5 **sboTerm** attribute with this value, then regardless of the identifier and name given to the reaction itself, a
6 software tool could use this to inform users that the reaction is a first-order irreversible mass-action reaction.
7 This kind of reverse engineering of the meaning of reactions in a model would be difficult to do otherwise,
8 especially for more complex reaction types.

9 The presence of an SBO label on a compartment, species, or reaction, can help map SBML elements to equiv-
10 alents in other standards, such as (but not limited to) BioPAX (<http://www.biopax.org/>) or the Systems
11 Biology Graphical Notation (SBGN, <http://www.sbgn.org/>). Such mappings can be used in conversion
12 procedures, or to build interfaces, with SBO becoming a kind of “glue” between standards of representation.

13 The presence of the label on a kinetic expression can also allow software tools to make more intelligent
14 decisions about reaction rate expressions. For example, an application could recognize **certain types of**
15 **reaction formulas** as being ones it knows how to solve with optimized procedures. The application could
16 then use internal, optimized code implementing the rate formula indexed by identifiers such as **SBO:0000049**
17 appearing in SBML models.

18 Although the use of SBO can be beneficial, it is critical to keep in mind that the presence of an **sboTerm**
19 value on an object *must not change the fundamental mathematical meaning* of the model. An SBML model
20 must be defined such that it stands on its own and does not depend on additional information added by
21 SBO terms for a correct mathematical interpretation. SBO term definitions will not imply any alternative
22 mathematical semantics for any SBML **object** labeled with that term. Two important reasons motivate this
23 principle. First, it would be too limiting to require all software tools to be able to understand the SBO
24 vocabularies in addition to understanding **SBML Level 2 Version 3**. Supporting SBO is not only additional
25 work for the software developer; for some kinds of applications, it may not make sense. If SBO terms on
26 a model are optional, it follows that the SBML model *must* remain unambiguous and fully interpretable
27 without them, because an application reading the model may ignore the terms. Second, we believe allowing
28 the use of **sboTerm** to alter the mathematical meaning of a model would **allow too much** leeway to shoehorn
29 inconsistent concepts into SBML **objects**, ultimately reducing the interoperability of the models.

30 5.2 Using SBO and sboTerm

31 The **sboTerm** attribute data type is always **SBOTerm**, defined in Section 3.1.9. When present in a given model
32 object instance, the **attribute**'s value must be an identifier taken from the Systems Biology Ontology (SBO;
33 <http://biomodels.net/SBO/>). This identifier must refer to a single SBO term that best defines the entity
34 encoded by the SBML object in question. An example of the type of relationship intended is: *the KineticLaw*
35 *in reaction R1 is a first-order irreversible mass action rate law.*

36 Note the careful use of the words “defines” and “entity encoded by the SBML object” in the paragraph
37 above. As mentioned, the relationship between the SBML object and the URI is:

38 The “thing” encoded by this SBML object *is an* instance of the “thing” represented by the
39 referenced SBO term.

40 5.2.1 The structure of the Systems Biology Ontology

41 The goal of SBO labeling for SBML is to clarify to the fullest extent possible the nature of **each element in**
42 **a model**. The approach taken in SBO begins with a hierarchically-structured set of controlled vocabularies
43 with **five** main divisions: (1) **modeling framework**, (2) **participant type**, (3) **quantitative parameter**, (4)
44 **mathematical expression**, and (5) **event**. Figure 23 on the following page illustrates the highest level of SBO.

45 Each of the **five** branches of Figure 23 have a hierarchy of terms underneath them. At this time, we
46 can only begin to list some initial concepts and terms in SBO; what follows is not meant to be com-
47 plete, comprehensive or even necessarily consistent with future versions of SBO. The web site for SBO

(<http://biomodels.net/SBO/>) should be consulted for the current version of the ontology. Section 5.4.1 describes how the impact of SBO changes on software applications is minimized.

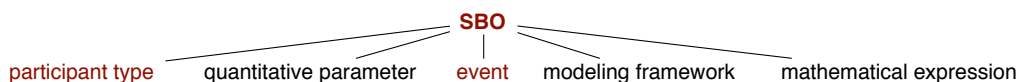


Figure 23: The five controlled vocabularies (CVs) that make up the main branches of SBO. (Other CVs in SBO may evolve in time, but are not discussed here.)

Figure 24 shows the structure for the *participant type* CV which reflects the hierarchical conceptual groupings of the concepts. For example, in reaction rate expressions, there are a variety of possible modifiers. Some classes of modifiers can be further subdivided and grouped. All of this is easy to capture in the CV. As more agreement is reached in the modeling community about how to define and name modifiers for different cases, the CV can grow to accommodate it.

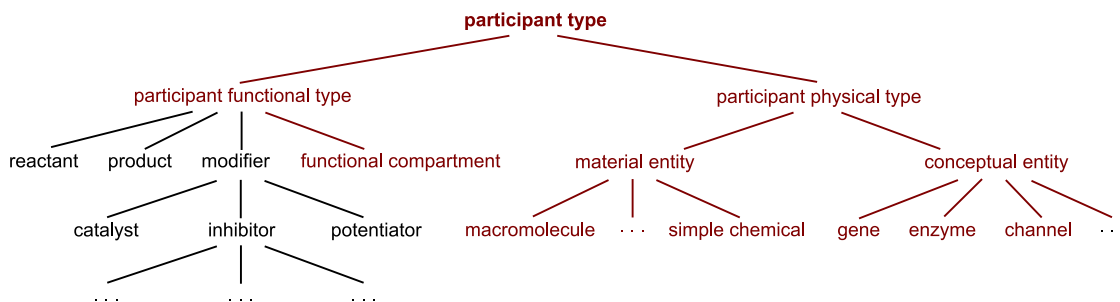


Figure 24: Partial expansion of some of the terms in the participant type CV of SBO.

Similar to the above, the controlled vocabulary for quantitative parameters also has a hierarchical structure, as illustrated in Figure 25. Note the separation of *kinetic constant* into separate terms for unimolecular, bimolecular, etc. reactions, as well as for forward and reverse reactions. The need to have separate terms for forward and reverse rate constants arises in reversible mass-action reactions. This distinction is not always necessary for all quantitative parameters; for example, there is no comparable concept for the Michaelis constant. Another distinction for some quantitative parameters is a decomposition into different versions based on the modeling framework being assumed. For example, different terms for continuous and discrete formulations of kinetic constants represent specializations of the constants for particular simulation frameworks. Not all quantitative parameters will need to be distinguished along this dimension.

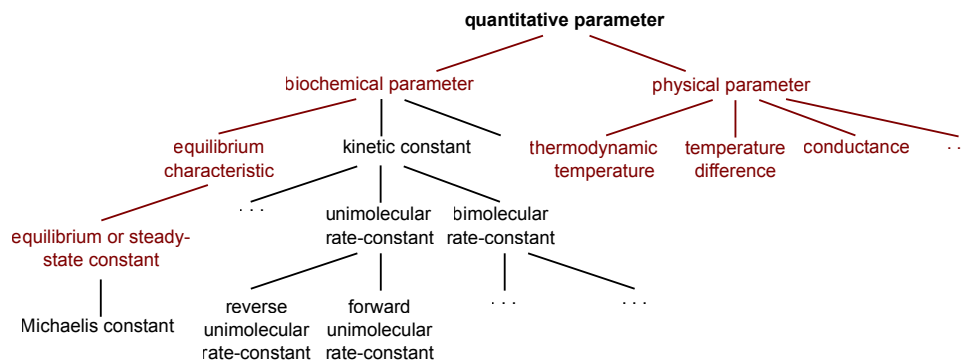


Figure 25: Partial expansion of some of the terms in the quantitative parameter CV.

1 The terms of the SBO quantitative parameter vocabulary contain mathematical formulas encoded using
 2 MathML 2.0 expressing the parameter using other SBO parameters. The main use of that approach is to
 3 avoid listing all the variants of a mathematical expression, escaping a combinatorial explosion.

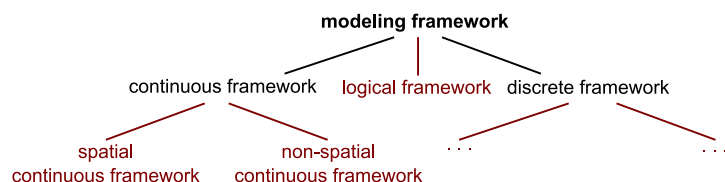


Figure 26: Partial expansion of some of the terms in the modeling framework CV.

4 The *modeling framework* controlled vocabulary is needed to support labeling models and reactions with the
 5 framework for which they are designed. Figure 26 illustrates the structure of this CV, which is at this point
 6 extremely simple, but we expect that more terms will evolve in the future.

7 The *mathematical expression vocabulary* encompasses the various mathematical expressions that constitute
 8 a model. Figure 27 illustrates a portion of the hierarchy. Rate law formulas are part of the mathematical
 9 expression hierarchy, and subdivided by successively more refined distinctions until the leaf terms represent
 10 precise statements of common reaction types. Other types of mathematical expressions are likely to be
 11 included in the future in order to be able to characterize other mathematical components of a model,
 12 namely initial assignments, assignment rules, rate rules, algebraic rules, constraints, and event triggers and
 13 assignments.

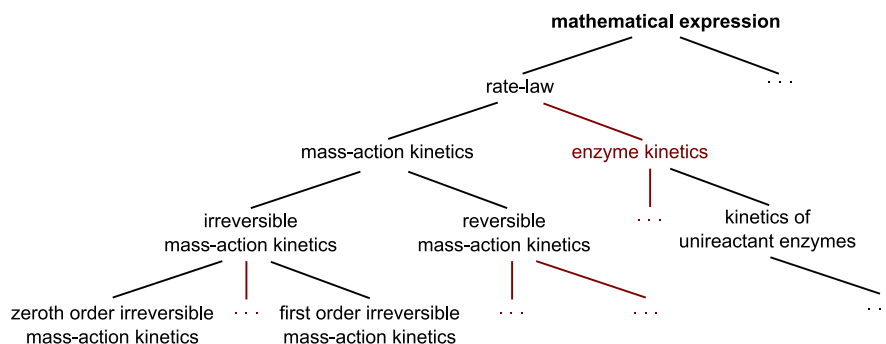


Figure 27: Partial expansion of some of the terms in the mathematical expression CV.

14 The leaf terms of the rate law portion of the SBO mathematical expression framework contain the mathe-
 15 matical formulas for the rate laws encoded using MathML 2.0. There are many potential uses for this. One is
 16 to allow a software application to obtain the formula and insert it into a model. In effect, the formulas given
 17 in the CV act as templates for what to put into an SBML **KineticLaw** definition. The MathML definition
 18 also acts as a precise statement about the rate law in question. Some of the non-leaf terms also contain rate
 19 laws encoded using MathML 2.0. However, those are generic rate-laws, containing MathML constructs not
 20 yet supported by SBML and need to be expanded into the MathML subset allowed in SBML before they
 21 can be used in conjunction with SBML models.

22 To make this discussion concrete, here is an example definition of an entry in the SBO rate law hierarchy
 23 at the time of this writing. This term represents second-order, irreversible, mass-action kinetics with one
 24 reactant, formulated for use in a continuous modeling framework:

25 *ID:* SBO:0000052

26 *Name:* second-order irreversible mass action kinetics, one reactant, continuous scheme.

Definition: Reaction scheme where the products are created from the reactants and the change of a product quantity is proportional to the product of reactant activities. The reaction scheme does not include any reverse process that creates the reactants from the products. The change of a product quantity is proportional to the square of one reactant quantity. It is to be used in a reaction modelled using a continuous framework.

Parent(s):

- SBO:0000050 second order irreversible mass action kinetics, one reactant (is a).
- SBO:0000163 irreversible mass action kinetics, continuous scheme (is a).

MathML:

```

<math xmlns="http://www.w3.org/1998/Math/MathML">
  <semantics definitionURL="http://biomodels.net/SBO/#SBO:0000062">
    <lambda>
      <bvar><ci definitionURL="http://biomodels.net/SBO/#SBO:0000036">k</ci></bvar>
      <bvar><ci definitionURL="http://biomodels.net/SBO/#SBO:0000010">R</ci></bvar>
      <apply>
        <times/>
        <ci>k</ci>
        <ci>R</ci>
        <ci>R</ci>
      </apply>
    </lambda>
  </semantics>
</math>

```

In the MathML definition of the term shown above, the bound variables in the `lambda` expression are tagged with references to terms in the SBO quantitative parameter (for `k`) and SBO participant role (for `R`) vocabularies. This makes it possible for software applications to interpret the intended meanings of the parameters in the expression.

The event branch of SBO defines types of biological processes. It lists the types of biochemical reactions, such as binding, conformational transition, or cleavage, and also the different controls that modify a biochemical reaction, such as inhibition, catalysis, etc.

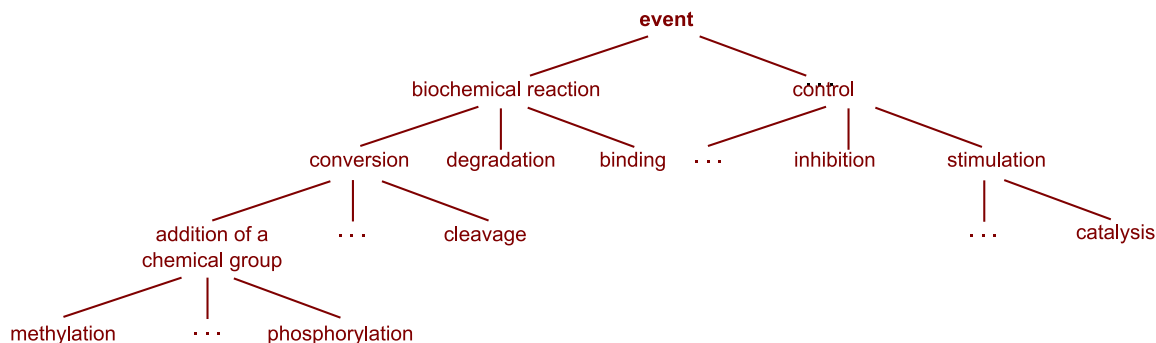


Figure 28: Partial expansion of some of the terms in the event CV.

One of the goals of SBO is to permit a tool to traverse up and down the hierarchy in order to find equivalent terms in different frameworks. The hope is that when a software tool encounters a given rate formula in a model, the formula will be a specific form (say, “mass-action kinetics, second order, one reactant, for discrete simulation”), but by virtue of the consistent organization of the reaction rate CV into framework-specific definitions, the tool should in principle be able to determine the definitions for other frameworks (say, “mass-action kinetics, second order, one reactant for *continuous* simulation”). If the software tool is designed for continuous simulation and it encounters an SBML model with rate laws formulated for discrete simulation, it could in principle look up the rate laws’ identifiers in the CV and search for alternative definitions intended for discrete simulation. And of course, the converse is true, for when a tool designed for discrete simulation encounters a model with rate laws formulated for continuous simulation.

5.2.2 Relationships between individual SBML components and SBO terms

The `sboTerm` attribute is defined on the abstract class `Sbase` and can be used in all derived elements. However, not all SBO terms can be used to annotate all SBML elements. Table 6 summarizes the relationship between SBML components and the branches within SBO that apply to that component. (There are currently no specific SBO term that correspond to the `Sbml`, `UnitDefinition`, `Unit`, and various `ListOf_____` list classes).

| SBML Component | SBO Vocabulary | Parent SBO Identifier |
|---------------------------------------|-----------------------------|-----------------------|
| <code>Model</code> | modeling framework | SBO:0000004 |
| <code>FunctionDefinition</code> | mathematical expression | SBO:0000064 |
| <code>CompartmentType</code> | participant physical type | SBO:0000236 |
| <code>SpeciesType</code> | participant physical type | SBO:0000236 |
| <code>Compartment</code> | participant physical type | SBO:0000236 |
| <code>Species</code> | participant physical type | SBO:0000236 |
| <code>Reaction</code> | event | SBO:0000231 |
| <code>Parameter</code> | quantitative parameter | SBO:0000002 |
| <code>SpeciesReference</code> | participant functional type | SBO:0000003 |
| <code>ModifierSpeciesReference</code> | participant functional type | SBO:0000003 |
| <code>KineticLaw</code> | rate law | SBO:0000001 |
| <code>InitialAssignment</code> | mathematical expression | SBO:0000064 |
| <code>AlgebraicRule</code> | mathematical expression | SBO:0000064 |
| <code>AssignmentRule</code> | mathematical expression | SBO:0000064 |
| <code>RateRule</code> | mathematical expression | SBO:0000064 |
| <code>Constraint</code> | mathematical expression | SBO:0000064 |
| <code>Event</code> | event | SBO:0000231 |
| <code>Trigger</code> | mathematical expression | SBO:0000064 |
| <code>Delay</code> | mathematical expression | SBO:0000064 |
| <code>EventAssignment</code> | mathematical expression | SBO:0000064 |

Table 6: SBML components and the main types of SBO terms that may be assigned to them. The parent identifiers are provided for guidance, but actual annotations should use more specific child terms. See text for explanation.

The parent identifiers shown in Table 6 are provided for reference. They are the highest-level terms in their respective vocabularies; however, these are *not* the terms that would be used to annotate an element in SBML, because there are more specific terms underneath the parents shown here. A software tool should use the most specific SBO term available for a given concept rather than using the top-level identifier acting as the root of that particular vocabulary.

5.2.3 Tradeoffs in using SBO terms

The SBO-based approach to annotating SBML components with controlled terms has the following strengths:

1. The syntax is minimally intrusive and maximally simple, requiring only one string-valued attribute.
2. It supports a significant fraction of what SBML users have wanted to do with controlled vocabularies.
3. It does not interfere with any other scheme. The more general annotation-based approach described in Section 6 can still be used simultaneously in the same model.

The scheme has the following weaknesses:

1. An object can only have one `sboTerm` attribute, therefore, it can only be related to a single term in SBO. (This also impacts the design of SBO: it must be structured such that a class of SBML elements can logically only be associated with one class of terms in the ontology.)
2. The only relationship that can be expressed by `sboTerm` is “is a”. It is not possible to represent different relationships (known as *verbs* in ontology-speak). This limits what can be expressed using SBO.

1 The weaknesses are not shared by the annotation scheme described in Section 6. If application's needs
2 cannot be met using SBO terms, software developers should examine the approach described in Section 6.

3 **5.3 Relationships to the SBML annotation element**

4 Another means of providing this kind of information would be to place SBO terms inside the **annotation**
5 element defined on *Sbase* (Sections 3.2 and 6). However, in the interest of making the use of SBO in SBML
6 maximally interoperable between software tools, the best-practice recommendation is to place SBO references
7 in the **sboTerm** attribute rather than inside the **annotation** element of an object.

8 Although **sboTerm** is just another kind of optional annotation in SBML, the SBO references are separated out
9 into their own attribute on SBML components, both to simplify their use for software tools and because they
10 assert a slightly stronger and more focused connection in a more regimented fashion. SBO term references
11 are intended to allow a modeler to make a statement of the form "this object is identical in meaning and
12 intention to the object defined in the term X of SBO", and do so in a way that a *software tool can interpret*
13 *unambiguously*.

14 Some software applications may have their own vocabulary of terms similar in purpose to SBO. For maximal
15 software interoperability, the best-practice recommendation in SBML is nonetheless to use SBO terms in
16 preference to using application-specific annotation schemes. Software applications should therefore attempt
17 to translate their private terms to and from SBO terms when writing and reading SBML, respectively.

18 **5.4 Discussion**

19 Here we discuss some additional points about the SBO-based approach.

20 **5.4.1 Frequency of change in the ontology**

21 The SBO development approach follows conventional ontology development approaches in bioinformatics.
22 One of the principles being followed is that identifiers and meanings of terms in the CVs never change
23 and the terms are never deleted. Where some terms are deemed obsolete, the introduction of new terms
24 refine or supersede existing terms, but the existing identifiers are left in the CV. Thus, references never
25 end up pointing to nonexistent entries. In the case where synonymous terms are merged after agreement
26 that multiple terms are identical, the term identifiers are again left in the CV and they still refer to the
27 same concept as before. Out-of-date terms cached or hard-coded by an application remain usable in all
28 cases. (Moreover, machine-readable CV encodings and appropriate software design should **render possible**
29 **the development of** API libraries that automatically map older terms to newer terms as the CVs evolve.)

30 Therefore, a model is never in danger of ending up with SBO identifiers that cannot be dereferenced. If an
31 application finds an old model with a term **SBO:0000065**, it can be assured that it will be able to find this
32 term in SBO, even if it has been superseded by other, more preferred terms.

33 **5.4.2 Consistency of information**

34 If you have a means of linking (say) a reaction rate formula to a term in a CV, it is possible to have an
35 inconsistency between the formula in the SBML model and the one defined for the CV term. However, this
36 is not a new problem; it arises in other situations involving SBML models already. The guideline for these
37 situations is that the model must be self-contained and stand on its own. Therefore, in cases where they
38 differ, the definitions in the SBML model take precedence over the definitions referenced by the CV. In other
39 words, the model (and its MathML) is authoritative.

40 **5.4.3 Implications for network access**

41 A software tool does not need to have the ability to access the network or read the CV every time it encounters
42 a model or otherwise works with SBML. Since the SBO will likely stabilize and change infrequently once
43 a core set of terms is defined, applications can cache the controlled vocabulary, and not make network
44 accesses to the master SBO copy unless something forces them to (e.g., detecting a reference in a model

1 to an SBO term that the application does not recognize). Applications could have user preference settings
2 indicating how often the CV definitions should be refreshed (similar to how modern applications provide a
3 setting dictating how often they should check for new versions of themselves). Simple applications may go
4 further and hard code references to terms in SBO that have reached stability and community consensus.
5 SBO is available for download under different formats (<http://biomodels.net/SBO/>). Web services are also
6 available to provide programmatic access to the ontology.

7 **5.4.4 Implications for software tools**

8 **If a software tool does not** pay attention to the SBO annotations described here, one is faced with exactly the
9 situation that exists today: the SBML model must be interpreted as-is, without benefit of the information
10 added by the SBO terms. The purpose of introducing an ontology scheme and guidelines for its use is to give
11 tools enough information that they *could* perform added processing, if they were designed to take advantage
12 of that information.

6 A standard format for the annotation element

This section describes the recommended non-proprietary format for the content of annotation elements in SBML when (a) referring to controlled vocabulary terms and database identifiers which define and describe biological and biochemical entities, and (b) describing the creator of a model and its modification history. Such a structured format should facilitate the generation of models compliant with the MIRIAM guidelines for model curation (Le Novère et al., 2005).

This format should *not* be used to refer to SBO terms (Section 5), because SBO defines terms about mathematical modeling constructs and not the biological and biochemical entities that the mathematics represent.

The format described in this section is intended to be the form of one of the top-level elements that could reside in an annotation element attached to an SBML object derived from *Sbase*. The element is named `rdf:RDF`. The format described here is compliant with the constraints placed on the form of annotation elements described in Section 3.2.4. We refer readers to Section 3.2.4 for important information on the structure and organization of application-specific annotations; these are not described here.

6.1 Motivation

The SBML structures described elsewhere in this document do not have any biochemical or biological semantics. The format described in this section provides a scheme for linking SBML structures to external resources so that those structures can have such semantics. The motivation for the introduction of this scheme is similar to that given for the introduction of `sboTerm`; however, this scheme is significantly more flexible.

6.2 XML namespaces in the standard annotation

This format uses a restricted form of Dublin Core (Dublin Core Metadata Initiative, 2005) and BioModels qualifier elements (see <http://www.biomodels.net/index.php?s=Qualifiers>) embedded in RDF (W3C, 2004b). It uses a number of external XML standards and associated XML namespaces. Table 7 lists these namespaces and relevant documentation on those namespaces. The format constrains the order of elements in these namespaces beyond the constraints defined in the standard definitions for those namespaces. For each standard listed, the format only uses a subset of the possible syntax defined by the given standard. Thus it is possible for an annotation element to include XML that is compliant with those external standards but is not compliant with the format described here. Parsers wishing to support this format should be aware that a valid annotation element may contain an `rdf:RDF` element which is not compliant with the format described here. A parser should check that all aspects of the syntax defined here before assuming that the contained data is encoded in the format.

| | Namespace URI | Definition Document |
|---------|---|---|
| dc | http://purl.org/dc/elements/1.1/ | Powell and Johnston (2003) |
| rdf | http://www.w3.org/1999/02/22-rdf-syntax-ns# | W3C (2004a) |
| dcterms | http://purl.org/dc/terms/ | Kokkeliink and Schwänzl (2002) DCMI Usage Board (2005) |
| vcard | http://www.w3.org/2001/vcard-rdf/3.0# | Iannella (2001) |
| bqbiol | http://biomodels.net/biology-qualifiers/ | |
| bqmodel | http://biomodels.net/model-qualifiers/ | |

Table 7: The XML standards used in the SBML standard format for annotation. The namespace prefix are shown to indicate only the prefix used in the main text.

6.3 General syntax for the standard annotation

An outline of the format syntax is shown below.

```
<SBML_ELEMENT +++ metaid="SBML_META_ID" +++ >
+++
<annotation>
+++
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:dcterms="http://purl.org/dc/terms/"
  xmlns:vcard="http://www.w3.org/2001/vcard-rdf/3.0#"
  xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
  xmlns:bqmodel="http://biomodels.net/model-qualifiers/"
>
  <rdf:Description rdf:about="#SBML_META_ID">
    [MODEL_HISTORY]
    <RELATION_ELEMENT>
      <rdf:Bag>
        <rdf:li rdf:resource="URI" />
        ...
      </rdf:Bag>
    </RELATION_ELEMENT>
    ...
  </rdf:Description>
  +++
</rdf:RDF>
+++
</annotation>
+++
</SBML_ELEMENT>
```

The above outline shows the order of the elements. The capitalized identifiers refer to generic strings of a particular type: **SBML_ELEMENT** refers to any SBML element name that can contain an **annotation** element; **SBML_META_ID** is a XML ID string; **RELATION_ELEMENT** refers to element names in either the namespace <http://biomodels.net/biology-qualifiers/> or <http://biomodels.net/model-qualifiers/>; and **URI** is a URI. **[MODEL_HISTORY]** refers to an optional section described in Section 6.6 which can only be present within SBML model elements. ‘+++’ is a placeholder for either no content or valid XML syntax that is not defined by the standard annotation scheme but is consistent with the relevant standards for the enclosing elements. ‘...’ is a placeholder for zero or more elements of the same form as the immediately preceding element. The precise form of whitespace and the XML namespace prefix definitions is not constrained; however, the elements and attributes must be in the namespaces shown. The rest of this section describes the format formally in English.

In this format, the annotation of an element is located in a single **rdf:RDF** element contained within an SBML **annotation** element. The annotation element can contain other elements in any order as described in Section 3.2.4. The format described in this section only defines the form of the **rdf:RDF** element. The containing SBML *Sbase* element must have a **metaid** attribute value. (As this attribute is of the type ID its value must be unique to the entire SBML document.)

The first element of the **rdf:RDF** element must be an **rdf:Description** element with an **rdf:about** attribute. (This format doesn’t define the form of subsequent elements of the **rdf:RDF** element.) The value of the **rdf:about** attribute must be of the form **#<string>** where the string component is equal to the value of the **metaid** attribute of the containing SBML element.

The **rdf:Description** element can contain only an optional model history section (see Section 6.6) followed by a sequence of zero or more BioModels relation elements. The optional model history section can only be present within an SBML **Model** element. The specific type of the relation elements will vary depending on the relationship between the SBML component and referenced information or resource.

Although Section 6.5 describes the detailed semantics of each of the relation element types, the content of these elements follows the same form. The BioModels qualifiers relation elements must only contain a single

1 **rdf:Bag** element which in turn must only contain one or more **rdf:li** elements. The **rdf:li** elements must
2 only have a **rdf:resource** attribute containing a URI referring to an information resource (See Section 6.4).

3 Annotations in this format can be located at different depths within a model component.

4 6.4 Use of URIs

5 The format represents a set of relationships between the SBML element and the resources referred to by the
6 contained **rdf:resource** attribute values. The BioModels relation elements simply define the type of the
7 relationship.

8 For example, a **Species** element representing a protein could be annotated with a reference to the database
9 UniProt by the <http://www.uniprot.org/#P12999> resource identifier, identifying exactly the protein de-
10 scribed by the **Species** element. This identifier maps to a unique entry in UniProt which is never deleted
11 from the database. In the case of UniProt, this is the “accession” of the entry. When the entry is merged
12 with another one, both “accession” are conserved. Similarly in a controlled vocabulary resource, each term
13 is associated with a perennial identifier. The UniProt entry also possess an “entry name” (the Swiss-Prot
14 “identifier”), a “protein name”, “synonyms” etc. Only the “accession” is perennial and should be used.

15 The value of a **rdf:resource** attribute is a URI that both uniquely identifies the resource, and the data in
16 the resource. In this case the resource constraining the identifier precedes the '#' symbol and the term or
17 database identifier follows the '#' symbol. In the above example, the resource <http://www.uniprot.org/>
18 includes the entry P12999.

19 The value of the **rdf:resource** attribute is a URI, not a URL; as such, a URI does not have to reference a
20 physical web object but simply identifies a controlled vocabulary term or database object (a URI is a label
21 that, in this case, just happens to look like a URL). For instance, a true URL for an Internet resource such as
22 <http://www.uniprot.org/entry/P12999> might correspond to the URI <http://www.uniprot.org/#P12999>.

23 SBML does not specify how a parser is to interpret a URI. In the case of a transformation into a physical
24 URL, there could be several solutions. For example, the URI <http://www.geneontology.org/#GO:0007268>
25 can be translated into:

26 <http://www.ebi.ac.uk/ego/DisplayGoTerm?selected=GO:0007268>
27 <http://www.godatabase.org/cgi-bin/amigo/go.cgi?view=details&query=GO:0007268>
28 <http://www.informatics.jax.org/searches/GO.cgi?id=GO:0007268>

29 Similarly the URI <http://www.ec-code.org/#3.5.4.4> can refer to:

30 <http://www.ebi.ac.uk/intenz/query?cmd=SearchEC&ec=3.5.4.4>
31 <http://www.expasy.org/cgi-bin/nicezyme.pl?3.5.4.4>
32 <http://www.chem.qmul.ac.uk/iubmb/enzyme/EC3/5/4/4.html>
33 <http://www.genome.jp/dbget-bin/www.bget?ec:3.5.4.4>
34 etc.

35 To enable interoperability, the community has agreed on an initial set of standardized valid URI syntax rules
36 which may be used within the standard annotation format. This set of rules is not part of the SBML standard
37 but will grow independently from specific SBML Levels and Versions. As the set changes a given URI syntax
38 rule will not be modified although the physical resource associated with the rule may change. These URIs
39 will always be composed as **resource#id**. The MIRIAM resources [http://www.ebi.ac.uk/compneur-srv/
40 miriam-main/](http://www.ebi.ac.uk/compneur-srv/miriam-main/) lists URI syntaxes and possible physical links to controlled vocabulary and databases. Each
41 entry contains a list of SBML and relation elements in which the given URI can be appropriately embedded.
42 To enable consistent and thus useful links to external resources, the URI syntax rule set must have a consistent
43 view of the concepts represented by the different SBML elements for the purposes of this format. For example
44 as the rule set is designed to link SBML biological and biochemical resources the rule set assumes that a
45 **Species** element represents the concept of a biochemical entity type rather than mathematical symbol. The
46 URI rule list will evolve with the evolution of databases and resources. The annotation format described in
47 this section does not require a simple parser of this format to access this list.

6.5 Relation elements

To enable the format to encode different types of relationships between SBML elements and resources, different BioModel qualifier elements are used to enclose a set of `rdf:li` elements. The relation elements imply a specific relationship between the enclosing SBML element and the resources referenced by the `rdf:li` elements.

The detailed semantics (i.e. from the perspective of automatic parser) of the relation elements is defined by the URI list at <http://www.ebi.ac.uk/compneur-srv/miriam-main/>, and thus is outside the scope of SBML. The URI list generally assumes that the biological entity represented by the element is the concept linked to the reference resource.

Several relation elements with a given tag, enclosed in the same SBML element, each represent an alternative annotation to the SBML element. For example two `bqbiol:hasPart` elements within a `Species` SBML element represent two different sets of references to the parts making up the the chemical entity represented by the species. (The species is not made up of all the entities represented by all the references combined).

The complete list of the qualifier elements in the Biomodels qualifier namespaces is documented at <http://www.biomodels.net/index.php?s=Qualifiers>. The list is divided into two namespaces one for model qualifiers <http://biomodels.net/biology-qualifiers/> (prefix used here `bqbiol`) and the other for biological qualifiers <http://biomodels.net/model-qualifiers/> (prefix used here `bqmodel`). This list will only grow i.e no element will be removed from the list. The following is the list of elements at the time of writing:

- `bqmodel:is` The modeling object encoded by the SBML component is the subject of the referenced resource. For instance, this qualifier might be used to link the model to a model database.
- `bqmodel:isDescribedBy` The modeling object encoded by the SBML component is described by the referenced resource. This relation might be used to link SBML components to the literature that describes this model or this kinetic law.
- `bqbiol:is` The biological entity represented by the SBML component is the subject of the referenced resource. This relation might be used to link a reaction to its exact counterpart in (e.g.) KEGG or Reactome.
- `bqbiol:hasPart` The biological entity represented by the SBML component includes the subject of the referenced resource, either physically or logically. This relation might be used to link a complex to the description of its components.
- `bqbiol:isPartOf` The biological entity represented by the SBML component is a physical or logical part of the subject of the referenced resource. This relation might be used to link a component to the description of the complex it belongs to.
- `bqbiol:isVersionOf` The biological entity represented by the SBML component is a version or an instance of the subject of the referenced resource.
- `bqbiol:hasVersion` The subject of the referenced resource is a version or an instance of the biological entity represented by the SBML component.
- `bqbiol:isHomologTo` The biological entity represented by the SBML component is homologous to the subject of the referenced resource, i.e. they share a common ancestor.
- `bqbiol:isDescribedBy` The biological entity represented by the SBML component is described by the referenced resource. This relation should be used, for example, to link a species or a parameter to the literature that describes the quantity of the species or the value of the parameter.

6.6 Model history

When enclosed in an SBML `Model` element, the format described in previous sections can include additional elements to describe the history of the model. This history data must occur immediately before the first

BioModels relation elements. These additional elements encode information on the model creator and a sequence of dates recording changes to the model. The syntax for this section is outlined below.

```

3     <dc:creator rdf:parseType="Resource">
4         <rdf:Bag>
5             <rdf:li rdf:parseType="Resource">
6                 [[
7                     +++
8                     <vCard:N rdf:parseType="Resource">
9                         <vCard:Family>FAMILY_NAME</vCard:Family>
10                        <vCard:Given>GIVEN_NAME</vCard:Given>
11                    </vCard:N>
12                    +++
13                    [<vCard:EMAIL>EMAIL_ADDRESS</vCard:EMAIL>]
14                    +++
15                    [<vCard:ORG>
16                        <vCard:Orgname>ORGANIZATION_NAME</vCard:Orgname>
17                    </vCard:ORG>]
18                    +++
19                    ]]
20                </rdf:li>
21            ...
22        </rdf:Bag>
23    </dc:creator>
24    <dcterms:created rdf:parseType="Resource">
25        <dcterms:W3CDTF>DATE</dcterms:W3CDTF>
26    </dcterms:created>
27    <dcterms:modified rdf:parseType="Resource">
28        <dcterms:W3CDTF>DATE</dcterms:W3CDTF>
29    </dcterms:modified>

```

The order of elements is as shown above except that elements of the format contained between [[and]] can occur in any order. The capitalized identifiers refer to generic strings of a particular type: **FAMILY_NAME** is the family name of a person who created the model; **GIVEN_NAME** is the first name of the same person who created the model; **EMAIL_ADDRESS** is the email address of the same person who created the model; and **ORGANIZATION_NAME** is the name of the organization with which the same person who created the model is affiliated **DATE** is a date in W3C date format (Wolf and Wicksteed, 1998). **W3CDTF**, **N**, **ORG** and **EMAIL** are literal strings. The elements of the format contained between [and] are optional. ‘+++’ is a placeholder for either no content or valid XML syntax that is not defined by the standard annotation scheme but is consistent with the relevant standards for the enclosing elements. ‘...’ is a placeholder for zero or more elements of the same form as the immediately preceding element. The precise form of whitespace and the XML namespace prefix definitions is not constrained. The remaining text in this section describes the syntax formally in English.

The additional elements of the model history sub-format consist in sequence of a **dc:creator** element, a **dcterms:created** element and zero or more **dcterms:modified** elements. All these elements must have the attribute **rdf:parseType** set to **Resource**.

The **dc:creator** element describes the person who created the SBML encoding of the model and contains a single **rdf:Bag** element. The **rdf:Bag** element can contain any number of elements; however, the first element must be a **rdf:li** element. The **rdf:li** element can contain any number of elements in any order. The set of elements contained with the **rdf:li** element can include the following informative elements: **vCard:N**, **vCard:EMAIL** and **vCard:ORG**. The **vCard:N** contains the name of the creator and must consist of a sequence of two elements: **vCard:Family** and the **vCard:Given** whose content is the family (surname) and given (first) names of the creator respectively. The **vCard:N** must have the attribute **rdf:parseType** set to **Resource**. The content of the **vCard:EMAIL** element must be the email address of the creator. The content of the **vCard:ORG** element must contain a single **vCard:Orgname** element. The **vCard:Orgname** element must contain the name of an organization to which the creator is affiliated.

The **dcterms:created** and **dcterms:modified** elements must each contain a single **dcterms:W3CDTF** element whose content is a date in W3C date format (Wolf and Wicksteed, 1998) which is a a profile of (restricted form of) ISO 8601.

6.7 Examples

The following shows the annotation of a model with model creation data and links to external resources:

```
1 <model metaid="_180340" id="GMO" name="Goldbeter1991_MinMitOscil">
2   <annotation>
3     <rdf:RDF
4       xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
5       xmlns:dc="http://purl.org/dc/elements/1.1/"
6       xmlns:dcterms="http://purl.org/dc/terms/"
7       xmlns:vCard="http://www.w3.org/2001/vcard-rdf/3.0#"
8       xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
9       xmlns:bqmodel="http://biomodels.net/model-qualifiers/"
10     >
11       <rdf:Description rdf:about="#_180340">
12         <dc:creator rdf:parseType="Resource">
13           <rdf:Bag>
14             <rdf:li rdf:parseType="Resource">
15               <vCard:N rdf:parseType="Resource">
16                 <vCard:Family>Shapiro</vCard:Family>
17                 <vCard:Given>Bruce</vCard:Given>
18               </vCard:N>
19               <vCard:EMAIL>bshapiro@jpl.nasa.gov</vCard:EMAIL>
20               <vCard:ORG>
21                 <vCard:Orgname>NASA Jet Propulsion Laboratory</vCard:Orgname>
22               </vCard:ORG>
23             </rdf:li>
24           </rdf:Bag>
25         </dc:creator>
26         <dcterms:created rdf:parseType="Resource">
27           <dcterms:W3CDTF>2005-02-06T23:39:40+00:00</dcterms:W3CDTF>
28         </dcterms:created>
29         <dcterms:modified rdf:parseType="Resource">
30           <dcterms:W3CDTF>2005-09-13T13:24:56+00:00</dcterms:W3CDTF>
31         </dcterms:modified>
32         <bqmodel:is>
33           <rdf:Bag>
34             <rdf:li rdf:resource="http://www.ebi.ac.uk/biomodels/#BIOMD0000000003"/>
35           </rdf:Bag>
36         </bqmodel:is>
37         <bqmodel:isDescribedBy>
38           <rdf:Bag>
39             <rdf:li rdf:resource="http://www.pubmed.gov/#1833774"/>
40           </rdf:Bag>
41         </bqmodel:isDescribedBy>
42         <bqbiol:isVersionOf>
43           <rdf:Bag>
44             <rdf:li rdf:resource="http://www.genome.jp/kegg/pathway/#hsa04110"/>
45             <rdf:li rdf:resource="http://www.reactome.org/#69278"/>
46           </rdf:Bag>
47         </bqbiol:isVersionOf>
48       </rdf:Description>
49     </rdf:RDF>
50   </annotation>
51 </model>
```

The following example shows a [Reaction](#) structure annotated with a reference to its exact Reactome counterpart.

```
55 <reaction id="cdc2Phospho" metaid="jb007">
56   <annotation>
57     <rdf:RDF
58       xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
59       xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
60     >
61       <rdf:Description rdf:about="#jb007">
62         <bqbiol:is>
63           <rdf:Bag>
64             <rdf:li rdf:resource="http://www.reactome.org/#170156"/>
65           </rdf:Bag>
66         </bqbiol:is>
67       </rdf:Description>
68     </rdf:RDF>
69   </annotation>
70 </reaction>
```



```

1         </bqbiol:is>
2         </rdf:Description>
3     </rdf:RDF>
4 </annotation>
5 <listOfReactants>
6     <speciesReference species="cdc2"/>
7 </listOfReactants>
8 <listOfProducts>
9     <speciesReference species="cdc2-Y15P"/>
10 </listOfProducts>
11 <listOfModifiers>
12     <modifierSpeciesReference species="wee1"/>
13 </listOfModifiers>
14 </reaction>

```

The following example describes a species that represents a complex between the protein calmodulin and calcium ions:

```

17 <species id="Ca_calmodulin" metaid="cacam">
18 <annotation>
19 <rdf:RDF
20     xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
21     xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
22 >
23     <rdf:Description rdf:about="#cacam">
24         <bqbiol:hasPart>
25             <rdf:Bag>
26                 <rdf:li rdf:resource="http://www.uniprot.org/#P62158"/>
27                 <rdf:li rdf:resource="http://www.genome.jp/kegg/compound/#C00076"/>
28             </rdf:Bag>
29         </bqbiol:hasPart>
30     </rdf:Description>
31 </rdf:RDF>
32 </annotation>
33 </species>

```

The following example describes a species that represents either “Calcium/calmodulin-dependent protein kinase type II alpha chain” or “Calcium/calmodulin-dependent protein kinase type II beta chain”. This is the case, for example, in the somatic cytoplasm of striatal medium-size spiny neurons, where both are present but they cannot be functionally differentiated.

```

38 <species id="calcium_calmodulin" metaid="cacam">
39 <annotation>
40 <rdf:RDF
41     xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
42     xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
43 >
44     <rdf:Description rdf:about="#cacam">
45         <bqbiol:hasVersion>
46             <rdf:Bag>
47                 <rdf:li rdf:resource="http://www.uniprot.org/#Q9UQM7"/>
48                 <rdf:li rdf:resource="http://www.uniprot.org/#Q13554"/>
49             </rdf:Bag>
50         </bqbiol:hasVersion>
51     </rdf:Description>
52 </rdf:RDF>
53 </annotation>
54 </species>

```

The above approach should not be used to describe “any Calcium/calmodulin-dependent protein kinase type II chain”, because such an annotation requires referencing the products of other genes such as gamma or delta. All the known proteins could be enumerated, but such an approach would almost surely lead to inaccuracies because biological knowledge continues to evolve. Instead, the annotation should refer to generic information such as Ensembl family ENSF0000000194 “CALCIUM/CALMODULIN DEPENDENT KINASE TYPE II CHAIN” or PIR superfamily PIRSF000594 “Calcium/calmodulin-dependent protein kinase type II”.

The following two examples show how to use the qualifier `isVersionOf`. The first example is the relationship between a reaction and an EC code. An EC code describes an enzymatic activity and an enzymatic reaction involving a particular enzyme can be seen as an instance of this activity. For example, the following reaction represents the phosphorylation of a glutamate receptor by a complex calcium/calmodulin kinase II.

```

5      <reaction id="NMDAR_phosphorylation" metaid="thx1138">
6        <annotation>
7          <rdf:RDF
8            xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
9            xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
10         >
11          <rdf:Description rdf:about="#thx1138">
12            <bqbiol:isVersionOf>
13              <rdf:Bag>
14                <rdf:li rdf:resource="http://www.ec-code.org/#2.7.1.17"/>
15              </rdf:Bag>
16            </bqbiol:isVersionOf>
17          </rdf:Description>
18        </rdf:RDF>
19      </annotation>
20      <listOfReactants>
21        <speciesReference species="NMDAR"/>
22      </listOfReactants>
23      <listOfProducts>
24        <speciesReference species="P-NMDAR"/>
25      </listOfProducts>
26      <listOfModifiers>
27        <modifierSpeciesReference species="CaMKII"/>
28      </listOfModifiers>
29      <kineticLaw>
30        <math xmlns="http://www.w3.org/1998/Math/MathML">
31          <apply>
32            <times/>
33            <ci>CaMKII</ci>
34            <ci>kcat</ci>
35          <apply>
36            <divide/>
37            <ci>NMDAR</ci>
38          <apply>
39            </times>
40            <ci>NMDAR</ci>
41            <ci>Km</ci>
42          </apply>
43        </apply>
44      </math>
45      <listOfParameters>
46        <parameter id="kcat" value="1"/>
47        <parameter id="Km" value="5e-10"/>
48      </listOfParameters>
49    </kineticLaw>
50  </reaction>

```

The second example of the use of `isVersionOf` is the complex between Calcium/calmodulin-dependent protein kinase type II alpha chain and Calcium/calmodulin, that is only one of the “calcium- and calmodulin-dependent protein kinase complexes” described by the Gene Ontology term GO:0005954.

```

55      <species id="CaCaMKII" metaid="C8H10N402">
56        <annotation>
57          <rdf:RDF
58            xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
59            xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
60         >
61          <rdf:Description rdf:about="#C8H10N402">
62            <bqbiol:isVersionOf>
63              <rdf:Bag>
64                <rdf:li rdf:resource="http://www.geneontology.org/#GO:0005954"/>
65              </rdf:Bag>

```

```

1         </bqbiol:isVersionOf>
2         </rdf:Description>
3     </rdf:RDF>
4 </annotation>
5 </species>

```

6 The previous case is different from the following one, although they could seem similar at first sight.
7 The “Calcium/calmodulin-dependent protein kinase type II alpha chain” is a part of the above mentioned
8 “calcium- and calmodulin-dependent protein kinase complex”.

```

9     <species id="CaMKIIalpha" metaid="C10H14N2">
10     <annotation>
11     <rdf:RDF
12     xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
13     xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
14     >
15     <rdf:Description rdf:about="#C10H14N2">
16     <bqbiol:isPartOf>
17     <rdf:Bag>
18     <rdf:li rdf:resource="http://www.geneontology.org/#GO:0005954"/>
19     </rdf:Bag>
20     </bqbiol:isPartOf>
21     </rdf:Description>
22     </rdf:RDF>
23     </annotation>
24 </species>

```

25 It is possible describe a component with several alternative sets of qualified annotations. For example, the
26 following species represents a pool of GMP, GDP and GTP. We annotate it with the three corresponding
27 KEGG compound identifiers but also with the three corresponding ChEBI identifiers. The two alternative
28 annotations are encoded in separate **hasVersion** qualifier elements.

```

29 <species id="GXP" metaid="GXP">
30 <annotation>
31 <rdf:RDF
32 xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
33 xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
34 >
35 <rdf:Description rdf:about="#GXP">
36 <bqbiol:hasVersion>
37 <rdf:Bag>
38 <rdf:li rdf:resource="http://www.ebi.ac.uk/chebi/#CHEBI:17345"/>
39 <rdf:li rdf:resource="http://www.ebi.ac.uk/chebi/#CHEBI:17552"/>
40 <rdf:li rdf:resource="http://www.ebi.ac.uk/chebi/#CHEBI:17627"/>
41 </rdf:Bag>
42 </bqbiol:hasVersion>
43 <bqbiol:hasVersion>
44 <rdf:Bag>
45 <rdf:li rdf:resource="http://www.genome.jp/kegg/compound/#C00035"/>
46 <rdf:li rdf:resource="http://www.genome.jp/kegg/compound/#C00044"/>
47 <rdf:li rdf:resource="http://www.genome.jp/kegg/compound/#C00144"/>
48 </rdf:Bag>
49 </bqbiol:hasVersion>
50 </rdf:Description>
51 </rdf:RDF>
52 </annotation>
53 </species>

```

54 The following example presents a reaction being actually the combination of three different elementary
55 molecular reactions. We annotate it with the three corresponding KEGG reactions, but also with the three
56 corresponding enzymatic activities. Again the two **hasPart** elements represent two alternative annotations.
57 The process represented by the **Reaction** structure is composed of three parts, and not six parts.

```

58 <reaction id="adenineProd" metaid="adeprod">
59 <annotation>
60 <rdf:RDF
61     xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"

```

```

1      xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
2    >
3      <rdf:Description rdf:about="#adeprod">
4        <bqbiol:hasPart>
5          <rdf:Bag>
6            <rdf:li rdf:resource="http://www.ec-code.org/#2.5.1.22"/>
7            <rdf:li rdf:resource="http://www.ec-code.org/#3.2.2.16"/>
8            <rdf:li rdf:resource="http://www.ec-code.org/#4.1.1.50"/>
9          </rdf:Bag>
10         </bqbiol:hasPart>
11         <bqbiol:hasPart>
12           <rdf:Bag>
13             <rdf:li rdf:resource="http://www.genome.jp/kegg/reaction/#R00178"/>
14             <rdf:li rdf:resource="http://www.genome.jp/kegg/reaction/#R01401"/>
15             <rdf:li rdf:resource="http://www.genome.jp/kegg/reaction/#R02869"/>
16           </rdf:Bag>
17         </bqbiol:hasPart>
18       </rdf:Description>
19     </rdf:RDF>
20   </annotation>
21 </reaction>

```

It is possible to mix different URIs in a given set. The following example presents two alternative annotations of the human hemoglobin, the first with ChEBI heme and the second with KEGG heme.

```

24 <species id="heme" metaid="heme">
25   <annotation>
26     <rdf:RDF
27       xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
28       xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
29     >
30       <rdf:Description rdf:about="#heme">
31         <bqbiol:hasPart>
32           <rdf:Bag>
33             <rdf:li rdf:resource="http://www.uniprot.org/#P69905"/>
34             <rdf:li rdf:resource="http://www.uniprot.org/#P68871"/>
35             <rdf:li rdf:resource="http://www.ebi.ac.uk/chebi/#CHEBI:17627">
36           </rdf:Bag>
37         </bqbiol:hasPart>
38         <bqbiol:hasPart>
39           <rdf:Bag>
40             <rdf:li rdf:resource="http://www.uniprot.org/#P69905"/>
41             <rdf:li rdf:resource="http://www.uniprot.org/#P68871"/>
42             <rdf:li rdf:resource="http://www.genome.jp/kegg/compound/#C00032"/>
43           </rdf:Bag>
44         </bqbiol:hasPart>
45       </rdf:Description>
46     </rdf:RDF>
47   </annotation>
48 </species>

```

As formally defined above it is possible to use different qualifiers in the same annotation element. The following phosphorylation is annotated by its exact KEGG counterpart and by the generic GO term “phosphorylation”.

```

52 <reaction id="phosphorylation" metaid="phosphorylation">
53   <annotation>
54     <rdf:RDF
55       xmlns:bqbiol="http://biomodels.net/biology-qualifiers/"
56       xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
57     >
58       <rdf:Description rdf:about="#phosphorylation">
59         <bqbiol:is>
60           <rdf:Bag>
61             <rdf:li rdf:resource="http://www.genome.jp/kegg/reaction/#R03313" />
62           </rdf:Bag>
63         </bqbiol:is>
64         <bqbiol:isVersionOf>

```

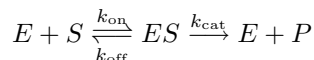
```
1         <rdf:Bag>
2           <rdf:li rdf:resource="http://www.geneontology.org/#GO:0016310" />
3         </rdf:Bag>
4       </bqbiol:isVersionOf>
5     </rdf:Description>
6   </rdf:RDF>
7 </annotation>
8 </reaction>
```

7 Example models expressed in XML using SBML

In this section, we present several examples of complete models encoded in XML using SBML Level 2.

7.1 A simple example application of SBML

Consider the following representation of an enzymatic reaction:



The following is the minimal SBML document encoding the model shown above:

```
<?xml version="1.0" encoding="UTF-8"?>
<sbml level="2" version="3" xmlns="http://www.sbml.org/sbml/level2/version3">
  <model name="EnzymaticReaction">
    <listOfUnitDefinitions>
      <unitDefinition id="per_second">
        <listOfUnits>
          <unit kind="second" exponent="-1"/>
        </listOfUnits>
      </unitDefinition>
      <unitDefinition id="litre_per_mole_per_second">
        <listOfUnits>
          <unit kind="mole" exponent="-1"/>
          <unit kind="litre" exponent="1"/>
          <unit kind="second" exponent="-1"/>
        </listOfUnits>
      </unitDefinition>
    </listOfUnitDefinitions>
    <listOfCompartments>
      <compartment id="cytosol" size="1e-14"/>
    </listOfCompartments>
    <listOfSpecies>
      <species compartment="cytosol" id="ES" initialAmount="0" name="ES"/>
      <species compartment="cytosol" id="P" initialAmount="0" name="P"/>
      <species compartment="cytosol" id="S" initialAmount="1e-20" name="S"/>
      <species compartment="cytosol" id="E" initialAmount="5e-21" name="E"/>
    </listOfSpecies>
    <listOfReactions>
      <reaction id="veq">
        <listOfReactants>
          <speciesReference species="E"/>
          <speciesReference species="S"/>
        </listOfReactants>
        <listOfProducts>
          <speciesReference species="ES"/>
        </listOfProducts>
        <kineticLaw>
          <math xmlns="http://www.w3.org/1998/Math/MathML">
            <apply>
              <times/>
              <ci>cytosol</ci>
              <apply>
                <minus/>
                <apply>
                  <times/>
                  <ci>kon</ci>
                  <ci>E</ci>
                  <ci>S</ci>
                </apply>
              </apply>
              <apply>
                <times/>
                <ci>koff</ci>
                <ci>ES</ci>
              </apply>
            </apply>
          </math>
        </kineticLaw>
      </reaction>
    </listOfReactions>
  </model>
</sbml>
```

```

1         <listOfParameters>
2             <parameter id="kon" value="1000000" units="litre_per_mole_per_second"/>
3             <parameter id="koff" value="0.2" units="per_second"/>
4         </listOfParameters>
5     </kineticLaw>
6 </reaction>
7 <reaction id="vcat" reversible="false">
8     <listOfReactants>
9         <speciesReference species="ES"/>
10    </listOfReactants>
11    <listOfProducts>
12        <speciesReference species="E"/>
13        <speciesReference species="P"/>
14    </listOfProducts>
15    <kineticLaw>
16        <math xmlns="http://www.w3.org/1998/Math/MathML">
17            <apply>
18                <times/>
19                <ci>cytosol</ci>
20                <ci>kcat</ci>
21                <ci>ES</ci>
22            </apply>
23        </math>
24        <listOfParameters>
25            <parameter id="kcat" value="0.1" units="per_second"/>
26        </listOfParameters>
27    </kineticLaw>
28 </reaction>
29 </listOfReactions>
30 </model>
31 </sbml>

```

In this example, the model has the identifier “EnzymaticReaction”. The model contains one compartment (with identifier “cytosol”), four species (with identifiers “ES”, “P”, “S”, and “E”), and two reactions (“veq” and “vcat”). The elements in the `listOfReactants` and `listOfProducts` in each reaction refer to the names of elements listed in the `listOfSpecies`. The correspondences between the various elements is explicitly stated by the `speciesReference` elements.

The model also features local parameter definitions in each reaction. In this case, the three parameters (“kon”, “koff”, “kcat”) all have unique identifiers and they could also have just as easily been declared global parameters in the model. Local parameters frequently become more useful in larger models, where it may become tedious to assign unique identifiers for all the different parameters.

7.2 Example involving units

The following model uses the units features of SBML Level 2. In this model, the default value of `substance` is changed to be mole units with a scale factor of -3 , or millimoles. This sets the default substance units in the model. The `volume` and `time` built-ins are left to their defaults, meaning `volume` is in litres and time is in seconds. The result is that, in this model, kinetic law formulas define rates in millimoles per second and the species identifiers in them represent concentration values in millimoles per litres. All the `species` elements set the initial amount of every given species to 1 millimole. The parameters “vm” and “km” are defined to be in millimoles per litres per second, and millimoles per litres, respectively.

```

49 <?xml version="1.0" encoding="UTF-8"?>
50 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3"
51     xmlns:xhtml="http://www.w3.org/1999/xhtml">
52     <model>
53         <listOfUnitDefinitions>
54             <unitDefinition id="substance">
55                 <listOfUnits>
56                     <unit kind="mole" scale="-3"/>
57                 </listOfUnits>
58             </unitDefinition>
59             <unitDefinition id="mmls">
60                 <listOfUnits>

```

```

1         <unit kind="mole" scale="-3"/>
2         <unit kind="litre" exponent="-1"/>
3         <unit kind="second" exponent="-1"/>
4     </listOfUnits>
5 </unitDefinition>
6 <unitDefinition id="mml">
7     <listOfUnits>
8         <unit kind="mole" scale="-3"/>
9         <unit kind="litre" exponent="-1"/>
10    </listOfUnits>
11 </unitDefinition>
12 </listOfUnitDefinitions>
13 <listOfCompartments>
14     <compartment id="cell" size="1"/>
15 </listOfCompartments>
16 <listOfSpecies>
17     <species id="x0" compartment="cell" initialConcentration="1"/>
18     <species id="x1" compartment="cell" initialConcentration="1"/>
19     <species id="s1" compartment="cell" initialConcentration="1"/>
20     <species id="s2" compartment="cell" initialConcentration="1"/>
21 </listOfSpecies>
22 <listOfParameters>
23     <parameter id="vm" value="2" units="mmls"/>
24     <parameter id="km" value="2" units="mml"/>
25 </listOfParameters>
26 <listOfReactions>
27     <reaction id="v1">
28         <listOfReactants>
29             <speciesReference species="x0"/>
30         </listOfReactants>
31         <listOfProducts>
32             <speciesReference species="s1"/>
33         </listOfProducts>
34         <kineticLaw>
35             <notes>
36                 <xhtml:p>((vm * s1)/(km + s1))*cell</xhtml:p>
37             </notes>
38             <math xmlns="http://www.w3.org/1998/Math/MathML">
39                 <apply>
40                     <times/>
41                     <apply>
42                         <divide/>
43                         <apply>
44                             <times/>
45                             <ci> vm </ci>
46                             <ci> s1 </ci>
47                         </apply>
48                         <plus/>
49                         <ci> km </ci>
50                         <ci> s1 </ci>
51                     </apply>
52                 </apply>
53             </math>
54             <ci> cell </ci>
55         </kineticLaw>
56     </reaction>
57 </listOfReactions>
58 <reaction id="v2">
59     <listOfReactants>
60         <speciesReference species="s1"/>
61     </listOfReactants>
62     <listOfProducts>
63         <speciesReference species="s2"/>
64     </listOfProducts>
65     <kineticLaw>
66         <notes>
67             <xhtml:p>((vm * s2)/(km + s2))*cell</xhtml:p>
68         </notes>
69     </kineticLaw>

```



```

1      <math xmlns="http://www.w3.org/1998/Math/MathML">
2          <apply>
3              <times/>
4              <apply>
5                  <divide/>
6                  <apply>
7                      <times/>
8                      <ci> vm </ci>
9                      <ci> s2 </ci>
10                 </apply>
11                <apply>
12                    <plus/>
13                    <ci> km </ci>
14                    <ci> s2 </ci>
15                </apply>
16            </apply>
17            <ci> cell </ci>
18        </apply>
19    </math>
20 </kineticLaw>
21 </reaction>
22 <reaction id="v3">
23     <listOfReactants>
24         <speciesReference species="s2"/>
25     </listOfReactants>
26     <listOfProducts>
27         <speciesReference species="x1"/>
28     </listOfProducts>
29     <kineticLaw>
30         <notes>
31             <xhtml:p>((vm * x1)/(km + x1))*cell</xhtml:p>
32         </notes>
33         <math xmlns="http://www.w3.org/1998/Math/MathML">
34             <apply>
35                 <times/>
36                 <apply>
37                     <divide/>
38                     <apply>
39                         <times/>
40                         <ci> vm </ci>
41                         <ci> x1 </ci>
42                     </apply>
43                     <plus/>
44                     <ci> km </ci>
45                     <ci> x1 </ci>
46                 </apply>
47             </apply>
48             <ci> cell </ci>
49         </math>
50     </kineticLaw>
51 </reaction>
52 </listOfReactions>
53 </model>
54 </sbml>

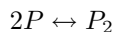
```

7.3 Example of a discrete version of a simple dimerization reaction

(Model contributed by Darren J. Wilkinson, Newcastle University, Newcastle upon Tyne, UK.)

This example illustrates subtle differences between models formulated for use in a continuous simulation framework (e.g., using differential equations) and those intended for a discrete simulation framework. The model shown here is suitable for use with a discrete stochastic simulation algorithm of the sort developed by Gillespie (1977). In such an approach, species are described in terms of molecular counts and simulation proceeds by computing the probability of the time and identity of the next reaction, then updating the species amounts appropriately.

1 The model involves a simple dimerization reaction for a protein named “P”:



3 The SBML representation is shown below. There are several important points to note. First, the species “P”
4 and “P2” declare they are always in discrete amounts by using the flag `hasOnlySubstanceUnits=“true”`.
5 This indicates that when the species identifiers appear in mathematical formulas, the units are *substance*, not
6 the default of *substance/size*. A second point is that, as a result, the corresponding “kinetic law” formulas
7 do not need volume corrections. In Gillespie’s approach, the constants in the rate expressions (here, “c1”
8 and “c2”) contain a contribution from the kinetic constants of the reaction and the size of the compartment
9 in which the reactions take place. Finally, it is worth noting the rate expression for the forward reaction is a
10 second-order mass-action reaction, but it is the *discrete* formulation of such a reaction rate (Gillespie, 1977).

```
11 <?xml version="1.0" encoding="UTF-8"?>
12 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
13   <model id="dimerization">
14     <listOfUnitDefinitions>
15       <unitDefinition id="substance">
16         <listOfUnits>
17           <unit kind="item" multiplier="1"/>
18         </listOfUnits>
19       </unitDefinition>
20       <unitDefinition id="per_second">
21         <listOfUnits>
22           <unit kind="second" exponent="-1"/>
23         </listOfUnits>
24       </unitDefinition>
25     </listOfUnitDefinitions>
26     <listOfCompartments>
27       <compartment id="Cell" size="1e-15"/>
28     </listOfCompartments>
29     <listOfSpecies>
30       <species id="P" compartment="Cell" initialAmount="301"
31         hasOnlySubstanceUnits="true"/>
32       <species id="P2" compartment="Cell" initialAmount="0"
33         hasOnlySubstanceUnits="true"/>
34     </listOfSpecies>
35     <listOfReactions>
36       <reaction id="Dimerization" reversible="false">
37         <listOfReactants>
38           <speciesReference species="P" stoichiometry="2"/>
39         </listOfReactants>
40         <listOfProducts>
41           <speciesReference species="P2"/>
42         </listOfProducts>
43         <kineticLaw>
44           <math xmlns="http://www.w3.org/1998/Math/MathML">
45             <apply>
46               <divide/>
47               <apply>
48                 <times/>
49                 <ci> c1 </ci>
50                 <ci> P </ci>
51               </apply>
52               <minus/>
53               <ci> P </ci>
54               <cn type="integer"> 1 </cn>
55             </apply>
56             </apply>
57             <cn type="integer"> 2 </cn>
58           </math>
59         </kineticLaw>
60       </reaction>
61     </listOfReactions>
62     <listOfParameters>
63       <parameter id="c1" value="0.00166" units="per_second"/>
64     </listOfParameters>
65   </model>
66 </sbml>
```

```

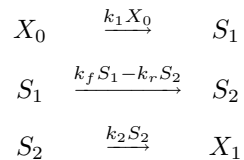
1      <reaction id="Dissociation" reversible="false">
2          <listOfReactants>
3              <speciesReference species="P2"/>
4          </listOfReactants>
5          <listOfProducts>
6              <speciesReference species="P" stoichiometry="2"/>
7          </listOfProducts>
8          <kineticLaw>
9              <math xmlns="http://www.w3.org/1998/Math/MathML">
10                 <apply>
11                     <times/>
12                     <ci> c2 </ci>
13                     <ci> P </ci>
14                 </apply>
15             </math>
16             <listOfParameters>
17                 <parameter id="c2" value="0.2" units="per_second"/>
18             </listOfParameters>
19         </kineticLaw>
20     </reaction>
21 </listOfReactions>
22 </model>
23 </sbml>

```

This example also illustrates the need to provide additional information in a model so that software tools using different mathematical frameworks can properly interpret it. In this case, a simulation tool designed for continuous ODE-based simulation would likely misinterpret the model (in particular the reaction rate formulas), unless it deduced that a discrete stochastic simulation was intended. One of the purposes of SBO annotations (Section 5) is to enable such interpretation without the need for deduction.

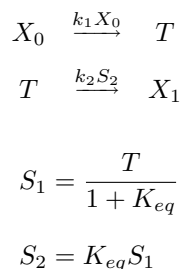
7.4 Example involving assignment rules

This section contains a model that simulates a system containing a fast reaction. This model uses rules to express the mathematics of the fast reaction explicitly rather than using the **fast** attribute on a reaction element. The system modeled is



$$k_1 = 0.1, \quad k_2 = 0.15, \quad k_f = K_{eq} 10000, \quad k_r = 10000, \quad K_{eq} = 2.5.$$

where X_0 , S_1 , S_1 , and S_2 are species in concentration units, and k_1 , k_2 , k_f , k_r , and K_{eq} are parameters. This system of reactions can be approximated with the following new system:



where T is a new species. The following example SBML model encodes the second system.

```

42     <?xml version="1.0" encoding="UTF-8"?>
43     <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3"
44         xmlns:math="http://www.w3.org/1998/Math/MathML">
45         <model>

```

```

1      <listOfUnitDefinitions>
2          <unitDefinition id="per_second">
3              <listOfUnits>
4                  <unit kind="second" exponent="-1"/>
5              </listOfUnits>
6          </unitDefinition>
7      </listOfUnitDefinitions>
8      <listOfCompartments>
9          <compartment id="cell" size="1"/>
10     </listOfCompartments>
11     <listOfSpecies>
12         <species id="X0" compartment="cell" initialConcentration="1"/>
13         <species id="X1" compartment="cell" initialConcentration="0"/>
14         <species id="T" compartment="cell" initialConcentration="0"/>
15         <species id="S1" compartment="cell" initialConcentration="0"/>
16         <species id="S2" compartment="cell" initialConcentration="0"/>
17     </listOfSpecies>
18     <listOfParameters>
19         <parameter id="Keq" value="2.5" units="dimensionless"/>
20     </listOfParameters>
21     <listOfRules>
22         <assignmentRule variable="S1">
23             <math xmlns="http://www.w3.org/1998/Math/MathML">
24                 <apply>
25                     <divide/>
26                     <ci> T </ci>
27                     <apply>
28                         <plus/>
29                         <cn> 1 </cn>
30                         <ci> Keq </ci>
31                     </apply>
32                 </apply>
33             </math>
34         </assignmentRule>
35         <assignmentRule variable="S2">
36             <math xmlns="http://www.w3.org/1998/Math/MathML">
37                 <apply>
38                     <times/>
39                     <ci> Keq </ci>
40                     <ci> S1 </ci>
41                 </apply>
42             </math>
43         </assignmentRule>
44     </listOfRules>
45     <listOfReactions>
46         <reaction id="in">
47             <listOfReactants>
48                 <speciesReference species="X0"/>
49             </listOfReactants>
50             <listOfProducts>
51                 <speciesReference species="T"/>
52             </listOfProducts>
53             <kineticLaw>
54                 <math xmlns="http://www.w3.org/1998/Math/MathML">
55                     <apply>
56                         <times/>
57                         <ci> k1 </ci>
58                         <ci> X0 </ci>
59                         <ci> cell </ci>
60                     </apply>
61                 </math>
62                 <listOfParameters>
63                     <parameter id="k1" value="0.1" units="per_second"/>
64                 </listOfParameters>
65             </kineticLaw>
66         </reaction>
67         <reaction id="out">
68             <listOfReactants>
69                 <speciesReference species="T"/>

```

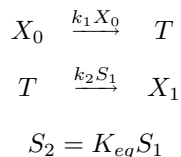
```

1         </listOfReactants>
2         <listOfProducts>
3             <speciesReference species="X1"/>
4         </listOfProducts>
5         <listOfModifiers>
6             <modifierSpeciesReference species="S2"/>
7         </listOfModifiers>
8         <kineticLaw>
9             <math xmlns="http://www.w3.org/1998/Math/MathML">
10                <apply>
11                    <times/>
12                    <ci> k2 </ci>
13                    <ci> S2 </ci>
14                    <ci> cell </ci>
15                </apply>
16            </math>
17            <listOfParameters>
18                <parameter id="k2" value="0.15" units="per_second"/>
19            </listOfParameters>
20        </kineticLaw>
21    </reaction>
22 </listOfReactions>
23 </model>
24 </sbml>

```

7.5 Example involving algebraic rules

This section contains an example model that contains an **AlgebraicRule** object. The model contains a different formulation of the fast reaction described in Section 7.4. The system described in Section 7.4 can be approximated with the following system:



with the constraint:

$$S_1 + S_2 - T = 0$$

The following example SBML model encodes this approximate form.

```

35 <?xml version="1.0" encoding="UTF-8"?>
36 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
37     <model>
38         <listOfUnitDefinitions>
39             <unitDefinition id="per_second">
40                 <listOfUnits>
41                     <unit kind="second" exponent="-1"/>
42                 </listOfUnits>
43             </unitDefinition>
44         </listOfUnitDefinitions>
45         <listOfCompartments>
46             <compartment id="cell" size="1"/>
47         </listOfCompartments>
48         <listOfSpecies>
49             <species id="X0" compartment="cell" initialConcentration="1"/>
50             <species id="X1" compartment="cell" initialConcentration="0"/>
51             <species id="T" compartment="cell" initialConcentration="0"/>
52             <species id="S1" compartment="cell" initialConcentration="0"/>
53             <species id="S2" compartment="cell" initialConcentration="0"/>
54         </listOfSpecies>
55         <listOfParameters>
56             <parameter id="Keq" value="2.5" units="dimensionless"/>
57         </listOfParameters>

```

```

1      <listOfRules>
2          <assignmentRule variable="S2">
3              <math xmlns="http://www.w3.org/1998/Math/MathML">
4                  <apply>
5                      <times/>
6                      <ci> Keq </ci>
7                      <ci> S1 </ci>
8                  </apply>
9              </math>
10         </assignmentRule>
11         <algebraicRule>
12             <math xmlns="http://www.w3.org/1998/Math/MathML">
13                 <apply>
14                     <minus/>
15                     <apply>
16                         <plus/>
17                             <ci> S2 </ci>
18                             <ci> S1 </ci>
19                         </apply>
20                     <ci> T </ci>
21                 </apply>
22             </math>
23         </algebraicRule>
24     </listOfRules>
25     <listOfReactions>
26         <reaction id="in">
27             <listOfReactants>
28                 <speciesReference species="X0"/>
29             </listOfReactants>
30             <listOfProducts>
31                 <speciesReference species="T"/>
32             </listOfProducts>
33             <kineticLaw>
34                 <math xmlns="http://www.w3.org/1998/Math/MathML">
35                     <apply>
36                         <times/>
37                         <ci> k1 </ci>
38                         <ci> X0 </ci>
39                         <ci> cell </ci>
40                     </apply>
41                 </math>
42                 <listOfParameters>
43                     <parameter id="k1" value="0.1" units="per_second"/>
44                 </listOfParameters>
45             </kineticLaw>
46         </reaction>
47         <reaction id="out">
48             <listOfReactants>
49                 <speciesReference species="T"/>
50             </listOfReactants>
51             <listOfProducts>
52                 <speciesReference species="X1"/>
53             </listOfProducts>
54             <listOfModifiers>
55                 <modifierSpeciesReference species="S2"/>
56             </listOfModifiers>
57             <kineticLaw>
58                 <math xmlns="http://www.w3.org/1998/Math/MathML">
59                     <apply>
60                         <times/>
61                         <ci> k2 </ci>
62                         <ci> S2 </ci>
63                         <ci> cell </ci>
64                     </apply>
65                 </math>
66                 <listOfParameters>
67                     <parameter id="k2" value="0.15" units="per_second"/>
68                 </listOfParameters>
69             </kineticLaw>

```

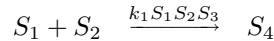
```

1         </reaction>
2     </listOfReactions>
3 </model>
4 </sbml>

```

7.6 Example with combinations of boundaryCondition and constant values on Species with RateRule objects

In this section, we discuss a model that includes four species, each with a different combination of values for their boundaryCondition and constant attributes. The model represents a hypothetical system containing one reaction,



where S_3 is a species that catalyzes the conversion of species S_1 and S_2 into S_4 . S_1 and S_2 are on the boundary of the system (i.e., S_1 and S_2 are reactants but their values are not determined by a kinetic law). The value of S_1 in the system is determined over time by the rate rule:

$$\frac{dS_1}{dt} = k_2$$

The values of constant parameters in the system are:

$$S_2 = 1, \quad S_3 = 2, \quad k_1 = 0.5, \quad k_2 = 0.1$$

and the initial values of species are:

$$S_1 = 0, \quad S_4 = 0$$

The value of S_1 varies over time so in SBML S_1 has a constant attribute with a default value of “false”. The values of S_2 and S_3 are fixed so in SBML they have a constant attribute values of “true”. S_3 only occurs as a modifier so the value of its boundaryCondition attribute can default to “false”. S_4 is a product whose value is determined by a kinetic law and therefore in the SBML representation has “false” (the default) for both its boundaryCondition and constant attributes.

The following is the SBML rendition of the model shown above:

```

25 <?xml version="1.0" encoding="UTF-8"?>
26 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
27   <model id="BoundaryCondExampleModel">
28     <listOfUnitDefinitions>
29       <unitDefinition id="mole_per_litre_per_second">
30         <listOfUnits>
31           <unit kind="mole" />
32           <unit kind="litre" exponent="-1"/>
33           <unit kind="second" exponent="-1"/>
34         </listOfUnits>
35       </unitDefinition>
36       <unitDefinition id="litre_sq_per_mole_sq_per_second">
37         <listOfUnits>
38           <unit kind="mole" exponent="-2"/>
39           <unit kind="litre" exponent="2"/>
40           <unit kind="second" exponent="-1"/>
41         </listOfUnits>
42       </unitDefinition>
43     </listOfUnitDefinitions>
44     <listOfCompartments>
45       <compartment id="compartmentOne" size="1"/>
46     </listOfCompartments>
47     <listOfSpecies>
48       <species id="S1" initialConcentration="0" compartment="compartmentOne"
49         boundaryCondition="true"/>
50       <species id="S2" initialConcentration="1" compartment="compartmentOne"
51         boundaryCondition="true" constant="true"/>
52       <species id="S3" initialConcentration="3" compartment="compartmentOne"

```

```

1          constant="true"/>
2      <species id="S4" initialConcentration="0" compartment="compartmentOne"/>
3  </listOfSpecies>
4  <listOfParameters>
5      <parameter id="k1" value="0.5" units="litre_sq_per_mole_sq_per_second"/>
6      <parameter id="k2" value="0.1" units="mole_per_litre_per_second"/>
7  </listOfParameters>
8  <listOfRules>
9      <rateRule variable="S1">
10         <math xmlns="http://www.w3.org/1998/Math/MathML">
11             <ci> k2 </ci>
12         </math>
13     </rateRule>
14 </listOfRules>
15 <listOfReactions>
16     <reaction id="reaction_1" reversible="false">
17         <listOfReactants>
18             <speciesReference species="S1"/>
19             <speciesReference species="S2"/>
20         </listOfReactants>
21         <listOfProducts>
22             <speciesReference species="S4"/>
23         </listOfProducts>
24         <listOfModifiers>
25             <modifierSpeciesReference species="S3"/>
26         </listOfModifiers>
27         <kineticLaw>
28             <math xmlns="http://www.w3.org/1998/Math/MathML">
29                 <apply>
30                     <times/>
31                     <ci> k1 </ci>
32                     <ci> S1 </ci>
33                     <ci> S2 </ci>
34                     <ci> S3 </ci>
35                     <ci> compartmentOne </ci>
36                 </apply>
37             </math>
38         </kineticLaw>
39     </reaction>
40 </listOfReactions>
41 </model>
42 </sbml>

```

7.7 Example of translation from a multi-compartmental model to ODEs

This section contains a model with 2 compartments and 4 reactions. The model is derived from Lotka-Volterra, with the addition of a reversible transport step. When observed in a time-course simulation, three of this model's species display damped oscillations.

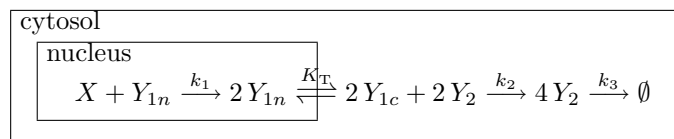


Figure 29: A example multi-compartmental model.

Figure 29 illustrates the arrangement of compartments and reactions in the model `LotkaVolterra_transport`. The text of the SBML representation of the model is shown below, and it is followed by its complete translation into ordinary differential equations. In this SBML model, the reaction equations are in substance per time units. The reactions have also been simplified to reduce common stoichiometric factors. The species variables are in concentration units; their initial quantities are declared using the attribute `initialAmount` on the `species` definitions, but since the attribute `hasOnlySubstanceUnits` is *not* set to true, the identifiers of

the species represent their concentrations when those identifiers appear in mathematical expressions elsewhere in the model. Note that the species whose identifier is “X” is a boundary condition, as indicated by the attribute `boundaryCondition=“true”` in its definition. The attribute `speciesType=“Y”` in the definitions of “Y1n” and “Y1c” indicates that these species are pools of the same participant, but located in different compartments.

```

6      <?xml version="1.0" encoding="UTF-8"?>
7      <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
8          <model name="LotkaVolterra_transport">
9              <listOfUnitDefinitions>
10                 <unitDefinition id="per_second">
11                     <listOfUnits>
12                         <unit kind="second" exponent="-1"/>
13                     </listOfUnits>
14                 </unitDefinition>
15                 <unitDefinition id="litre_per_mole_per_second">
16                     <listOfUnits>
17                         <unit kind="mole" exponent="-1"/>
18                         <unit kind="litre" exponent="1"/>
19                         <unit kind="second" exponent="-1"/>
20                     </listOfUnits>
21                 </unitDefinition>
22             </listOfUnitDefinitions>
23             <listOfSpeciesTypes>
24                 <speciesType id="Y1"/>
25             </listOfSpeciesTypes>
26             <listOfCompartments>
27                 <compartment id="cytoplasm" size="5"/>
28                 <compartment id="nucleus" outside="cytoplasm" size="1"/>
29             </listOfCompartments>
30             <listOfSpecies>
31                 <species id="X" compartment="nucleus" initialAmount="1" constant="true"
32                     boundaryCondition="true"/>
33                 <species id="Y1n" compartment="nucleus" speciesType="Y1" initialAmount="1"/>
34                 <species id="Y1c" compartment="cytoplasm" speciesType="Y1" initialAmount="0"/>
35                 <species id="Y2" compartment="cytoplasm" initialAmount="1"/>
36             </listOfSpecies>
37             <listOfParameters>
38                 <parameter id="k1" value="2500" units="litre_per_mole_per_second"/>
39                 <parameter id="k2" value="2500" units="litre_per_mole_per_second"/>
40                 <parameter id="kT" value="25000" units="per_second"/>
41                 <parameter id="k3" value="2500" units="per_second"/>
42             </listOfParameters>
43             <listOfReactions>
44                 <reaction id="production" reversible="false">
45                     <listOfReactants>
46                         <speciesReference species="X"/>
47                         <speciesReference species="Y1n"/>
48                     </listOfReactants>
49                     <listOfProducts>
50                         <speciesReference species="Y1n"/>
51                         <speciesReference species="Y1n"/>
52                     </listOfProducts>
53                     <kineticLaw>
54                         <math xmlns="http://www.w3.org/1998/Math/MathML">
55                             <apply>
56                                 <times/>
57                                 <ci>nucleus</ci>
58                                 <ci>k1</ci>
59                                 <ci>X</ci>
60                                 <ci>Y1n</ci>
61                             </apply>
62                         </math>
63                     </kineticLaw>
64                 </reaction>
65                 <reaction id="transport" reversible="true">
66                     <listOfReactants>
67                         <speciesReference species="Y1n"/>

```

```

1      </listOfReactants>
2      <listOfProducts>
3          <speciesReference species="Y1c"/>
4      </listOfProducts>
5      <kineticLaw>
6          <math xmlns="http://www.w3.org/1998/Math/MathML">
7              <apply>
8                  <times/>
9                  <ci>cytoplasm</ci>
10                 <ci>KT</ci>
11                 <apply>
12                     <minus/>
13                     <ci>Y1n</ci>
14                     <ci>Y1c</ci>
15                 </apply>
16             </apply>
17         </math>
18     </kineticLaw>
19 </reaction>
20 <reaction id="transformation" reversible="false">
21     <listOfReactants>
22         <speciesReference species="Y1c"/>
23         <speciesReference species="Y2"/>
24     </listOfReactants>
25     <listOfProducts>
26         <speciesReference species="Y2" stoichiometry="2"/>
27     </listOfProducts>
28     <kineticLaw>
29         <math xmlns="http://www.w3.org/1998/Math/MathML">
30             <apply>
31                 <times/>
32                 <ci>cytoplasm</ci>
33                 <ci>k2</ci>
34                 <ci>Y1c</ci>
35                 <ci>Y2</ci>
36             </apply>
37         </math>
38     </kineticLaw>
39 </reaction>
40 <reaction id="degradation" reversible="false">
41     <listOfReactants>
42         <speciesReference species="Y2"/>
43     </listOfReactants>
44     <kineticLaw>
45         <math xmlns="http://www.w3.org/1998/Math/MathML">
46             <apply>
47                 <times/>
48                 <ci>cytoplasm</ci>
49                 <ci>k3</ci>
50                 <ci>Y2</ci>
51             </apply>
52         </math>
53     </kineticLaw>
54 </reaction>
55 </listOfReactions>
56 </model>
57 </sbml>

```

The ODE translation of this model is as follows. First, we give the values of the constant parameters:

$$k_1 = 2500, \quad k_2 = 2500, \quad K_3 = 25000, \quad K_T = 2500$$

Next, here are the initial conditions of the variables. The species variables X , Y_{1n} , Y_{1c} , and Y_2 in the following equations are all in terms of concentrations, and we use V_n to represent for the size of compartment “nucleus” and V_c to represent for the size of compartment “cytoplasm”:

$$V_n = 1, \quad V_c = 5, \quad X = 1, \quad Y_{1n} = 1, \quad Y_{1c} = 0, \quad Y_2 = 1$$

1 And finally, here are the differential equations:

$$2 \quad \frac{dX}{dt} = 0$$

$$3 \quad V_n \frac{dY_{1n}}{dt} = k_1 \cdot X \cdot Y_{1n} \cdot V_n - K_T \cdot (Y_{1n} - Y_{1c}) \cdot V_c \quad \text{reactions production and transport}$$

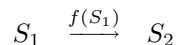
$$4 \quad V_c \frac{dY_{1c}}{dt} = K_T \cdot (Y_{1n} - Y_{1c}) \cdot V_c - k_2 \cdot Y_{1c} \cdot Y_2 \cdot V_c \quad \text{reactions transport and transformation}$$

$$5 \quad V_c \frac{dY_2}{dt} = k_2 \cdot Y_{1c} \cdot Y_2 \cdot V_c - k_3 \cdot Y_2 \cdot V_c \quad \text{reactions transformation and degradation}$$

7 As formulated here, this example assumes constant volumes. If the sizes of the compartments “cytoplasm”
8 or “nucleus” could change during simulation, then it would be preferable to use a different approach to
9 constructing the differential equations. In this alternative approach, the ODEs would compute substance
10 change rather than concentration change, and the concentration values would be computed using separate
11 equations. This approach is used in Section 4.13.6.

12 7.8 Example involving function definitions

13 This section contains a model that uses the function definition feature of SBML. Consider the following
14 hypothetical system:



16 where

$$17 \quad f(x) = 2 \times x$$

18 The following is the XML document that encodes the model shown above:

```
19 <?xml version="1.0" encoding="UTF-8"?>
20 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
21   <model id="Example">
22     <listOfFunctionDefinitions>
23       <functionDefinition id="f">
24         <math xmlns="http://www.w3.org/1998/Math/MathML">
25           <lambda>
26             <bvar>
27               <ci> x </ci>
28             </bvar>
29             <apply>
30               <times/>
31               <ci> x </ci>
32               <cn> 2 </cn>
33             </apply>
34           </lambda>
35         </math>
36       </functionDefinition>
37     </listOfFunctionDefinitions>
38     <listOfCompartments>
39       <compartment id="compartmentOne" size="1"/>
40     </listOfCompartments>
41     <listOfSpecies>
42       <species id="S1" initialConcentration="1" compartment="compartmentOne"/>
43       <species id="S2" initialConcentration="0" compartment="compartmentOne"/>
44     </listOfSpecies>
45     <listOfParameters>
46       <parameter id="t" value = "1" units="second"/>
47     </listOfParameters>
48     <listOfReactions>
49       <reaction id="reaction_1" reversible="false">
50         <listOfReactants>
```

```

1         <speciesReference species="S1"/>
2     </listOfReactants>
3     <listOfProducts>
4         <speciesReference species="S2"/>
5     </listOfProducts>
6     <kineticLaw>
7         <math xmlns="http://www.w3.org/1998/Math/MathML">
8             <apply>
9                 <divide/>
10                <apply>
11                    <times/>
12                    <apply>
13                        <ci> f </ci>
14                        <ci> S1 </ci>
15                    </apply>
16                    <ci> compartmentOne </ci>
17                </apply>
18                <ci> t</ci>
19            </apply>
20        </math>
21    </kineticLaw>
22 </reaction>
23 </listOfReactions>
24 </model>
25 </sbml>

```

7.9 Example involving *delay* functions

The following is a simple model illustrating the use of *delay* to represent a gene that suppresses its own expression. The model can be expressed in a single rule:

$$\frac{dP}{dt} = \frac{1}{1 + m(P_{delayed})^q} - P$$

where

$P_{delayed}$ is *delay*(P, Δ_t) or P at $t - \Delta_t$
 P is protein concentration
 τ is the response time
 m is a multiplier or equilibrium constant
 q is the Hill coefficient

and the species quantities are in concentration units. The text of an SBML encoding of this model is given below:

```

34 <?xml version="1.0" encoding="UTF-8"?>
35 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
36     <model>
37         <listOfCompartments>
38             <compartment id="cell" size="1"/>
39         </listOfCompartments>
40         <listOfSpecies>
41             <species id="P" compartment="cell" initialConcentration="0"/>
42         </listOfSpecies>
43         <listOfParameters>
44             <parameter id="tau" value="1" units="second"/>
45             <parameter id="m" value="0.5" units="dimensionless"/>
46             <parameter id="q" value="1" units="dimensionless"/>
47             <parameter id="delta_t" value="1" units="second"/>
48         </listOfParameters>
49         <listOfRules>
50             <rateRule variable="P">
51                 <math xmlns="http://www.w3.org/1998/Math/MathML">
52                     <apply>
53                         <divide/>
54                     </apply>

```

```

1      <minus/>
2      <apply>
3        <divide/>
4        <cn> 1 </cn>
5        <apply>
6          <plus/>
7          <cn> 1 </cn>
8          <apply>
9            <times/>
10           <ci> m </ci>
11           <apply>
12             <power/>
13             <apply>
14               <csymbol encoding="text"
15                 definitionURL="http://www.sbml.org/sbml/symbols/delay">
16                 delay
17               </csymbol>
18               <ci> P </ci>
19               <ci> delta_t </ci>
20             </apply>
21             <ci> q </ci>
22           </apply>
23         </apply>
24       </apply>
25     </apply>
26     <ci> P </ci>
27   </apply>
28   <ci> tau </ci>
29 </apply>
30 </math>
31 </rateRule>
32 </listOfRules>
33 </model>
34 </sbml>

```

7.10 Example involving events

This section presents a simple model system that demonstrates the use of events in SBML. Consider a system with two genes, G_1 and G_2 . G_1 is initially on and G_2 is initially off. When turned on, the two genes lead to the production of two products, P_1 and P_2 , respectively, at a fixed rate. When P_1 reaches a given concentration, G_2 switches on. This system can be represented mathematically as follows:

$$\frac{dP_1}{dt} = k_1(G_1 - P_1)$$

$$\frac{dP_2}{dt} = k_2(G_2 - P_2)$$

$$G_2 = \begin{cases} 0 & \text{when } P_1 \leq \tau, \\ 1 & \text{when } P_1 > \tau. \end{cases}$$

The initial values are:

$$G_1 = 1, \quad G_2 = 0, \quad \tau = 0.25, \quad P_1 = 0, \quad P_2 = 0, \quad k_1 = k_2 = 1.$$

The SBML Level 2 representation of this as follows:

```

46 <?xml version="1.0" encoding="UTF-8"?>
47 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3"
48   xmlns:math="http://www.w3.org/1998/Math/MathML">
49   <model>
50     <listOfUnitDefinitions>
51       <unitDefinition id="per_second">
52         <listOfUnits>
53           <unit kind="second" exponent="-1"/>

```

```

1         </listOfUnits>
2     </unitDefinition>
3     <unitDefinition id="concentration">
4         <listOfUnits>
5             <unit kind="mole"/>
6             <unit kind="litre" exponent="-1"/>
7         </listOfUnits>
8     </unitDefinition>
9 </listOfUnitDefinitions>
10 <listOfCompartments>
11     <compartment id="cell" size="1"/>
12 </listOfCompartments>
13 <listOfSpecies>
14     <species id="P1" compartment="cell" initialConcentration="0"/>
15     <species id="P2" compartment="cell" initialConcentration="0"/>
16 </listOfSpecies>
17 <listOfParameters>
18     <parameter id="k1" value="1" units="per_second"/>
19     <parameter id="k2" value="1" units="per_second" />
20     <parameter id="tau" value="0.25" units="concentration"/>
21     <parameter id="G1" value="1" units="concentration" constant="false"/>
22     <parameter id="G2" value="0" units="concentration" constant="false"/>
23 </listOfParameters>
24 <listOfRules>
25     <rateRule variable="P1">
26         <math:math>
27             <math:apply>
28                 <math:times/>
29                 <math:ci> k1 </math:ci>
30             <math:apply>
31                 <math:minus/>
32                 <math:ci> G1 </math:ci>
33                 <math:ci> P1 </math:ci>
34             </math:apply>
35         </math:math>
36     </rateRule>
37     <rateRule variable="P2">
38         <math:math>
39             <math:apply>
40                 <math:times/>
41                 <math:ci> k2 </math:ci>
42             <math:apply>
43                 <math:minus/>
44                 <math:ci> G2 </math:ci>
45                 <math:ci> P2 </math:ci>
46             </math:apply>
47         </math:math>
48     </rateRule>
49 </listOfRules>
50 <listOfEvents>
51     <event>
52         <trigger>
53             <math:math>
54                 <math:apply>
55                     <math:gt/>
56                     <math:ci> P1 </math:ci>
57                     <math:ci> tau </math:ci>
58                 </math:apply>
59             </math:math>
60         </trigger>
61         <listOfEventAssignments>
62             <eventAssignment variable="G2">
63                 <math:math>
64                     <math:apply>
65                         <math:cn> 1 </math:cn>
66                     </math:apply>
67                 </math:math>
68             </eventAssignment>
69         </listOfEventAssignments>

```

```

1         </eventAssignment>
2     </listOfEventAssignments>
3 </event>
4 <event>
5     <trigger>
6         <math:math>
7             <math:apply>
8                 <math:leq/>
9                 <math:ci> P1 </math:ci>
10                <math:ci> tau </math:ci>
11            </math:apply>
12        </math:math>
13    </trigger>
14    <listOfEventAssignments>
15        <eventAssignment variable="G2">
16            <math:math>
17                <math:apply>
18                    <math:cn> 0 </math:cn>
19                </math:apply>
20            </math:math>
21        </eventAssignment>
22    </listOfEventAssignments>
23 </event>
24 </listOfEvents>
25 </model>
26 </sbml>

```

7.11 Example involving two-dimensional compartments

The following example is a model that uses a two-dimensional compartment. It is a fragment of a larger model of calcium regulation across the plasma membrane of a cell. The model includes a calcium influx channel, “Ca_channel”, and a calcium-extruding PMCA pump, “Ca_Pump”. It also includes two cytosolic proteins that buffer calcium via the “CalciumCalbindin_gt_BoundCytosol” and “CalciumBuffer_gt_BoundCytosol” reactions. Finally, the rate expressions in this model do not include explicit factors of the compartment volumes; instead, the various rate constants are assume to include any necessary corrections for volume.

```

34 <?xml version="1.0" encoding="UTF-8"?>
35 <sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
36     <model id="facilitated_ca_diffusion">
37         <listOfUnitDefinitions>
38             <unitDefinition id="substance">
39                 <listOfUnits>
40                     <unit kind="mole" scale="-6"/>
41                 </listOfUnits>
42             </unitDefinition>
43             <unitDefinition id="area">
44                 <listOfUnits>
45                     <unit kind="metre" scale="-6" exponent="2"/>
46                 </listOfUnits>
47             </unitDefinition>
48             <unitDefinition id="per_second">
49                 <listOfUnits>
50                     <unit kind="second" exponent="-1"/>
51                 </listOfUnits>
52             </unitDefinition>
53             <unitDefinition id="litre_per_mole_per_second">
54                 <listOfUnits>
55                     <unit kind="mole" exponent="-1" scale="-6"/>
56                     <unit kind="litre" exponent="1"/>
57                     <unit kind="second" exponent="-1"/>
58                 </listOfUnits>
59             </unitDefinition>
60             <unitDefinition id="subs_per_vol">
61                 <listOfUnits>
62                     <unit kind="mole" exponent="1" scale="-6"/>
63                     <unit kind="litre" exponent="-1"/>
64                 </listOfUnits>

```

```

1      </unitDefinition>
2  </listOfUnitDefinitions>
3  <listOfCompartments>
4      <compartment id="Extracellular"
5          spatialDimensions="3" size="1"/>
6      <compartment id="PlasmaMembrane"
7          outside="Extracellular" spatialDimensions="2" size="1"/>
8      <compartment id="Cytosol"
9          outside="PlasmaMembrane" spatialDimensions="3" size="1"/>
10 </listOfCompartments>
11 <listOfSpecies>
12 <species id="CaBPB_C"
13     compartment="Cytosol" initialConcentration="47.17"/>
14 <species id="B_C"
15     compartment="Cytosol" initialConcentration="396.04"/>
16 <species id="CaB_C"
17     compartment="Cytosol" initialConcentration="3.96"/>
18 <species id="Ca_C"
19     name="Ca" compartment="Cytosol" initialConcentration="0.1"/>
20 <species id="Ca_EC"
21     name="Ca" compartment="Extracellular" initialConcentration="1000"/>
22 <species id="CaCh_PM"
23     compartment="PlasmaMembrane" initialConcentration="1"/>
24 <species id="CaPump_PM"
25     compartment="PlasmaMembrane" initialConcentration="1"/>
26 <species id="CaBP_C"
27     compartment="Cytosol" initialConcentration="202.83"/>
28 </listOfSpecies>
29 <listOfReactions>
30 <reaction id="CalciumCalbindin_gt_BoundCytosol" fast="true">
31     <listOfReactants>
32         <speciesReference species="CaBP_C"/>
33         <speciesReference species="Ca_C"/>
34     </listOfReactants>
35     <listOfProducts>
36         <speciesReference species="CaBPB_C"/>
37     </listOfProducts>
38     <kineticLaw>
39         <notes>
40             <p xmlns="http://www.w3.org/1999/xhtml">
41                 ((Kf_CalciumCalbindin_BoundCytosol * CaBP_C) * Ca_C) -
42                 (Kr_CalciumCalbindin_BoundCytosol * CaBPB_C)
43             </p>
44         </notes>
45         <math xmlns="http://www.w3.org/1998/Math/MathML">
46             <apply>
47                 <times/>
48                 <ci> Cytosol </ci>
49                 <apply>
50                     <minus/>
51                     <apply>
52                         <times/>
53                         <ci> Kf_CalciumCalbindin_BoundCytosol </ci>
54                         <ci> CaBP_C </ci>
55                         <ci> Ca_C </ci>
56                     </apply>
57                     <apply>
58                         <times/>
59                         <ci> Kr_CalciumCalbindin_BoundCytosol </ci>
60                         <ci> CaBPB_C </ci>
61                     </apply>
62                 </apply>
63             </math>
64         </listOfParameters>
65         <parameter id="Kf_CalciumCalbindin_BoundCytosol" value="20.0"
66             units="litre_per_mole_per_second"/>
67         <parameter id="Kr_CalciumCalbindin_BoundCytosol" value="8.6"
68             units="per_second"/>
69     </kineticLaw>
31 </listOfReactions>

```



```

1         </listOfParameters>
2     </kineticLaw>
3 </reaction>
4 <reaction id="CalciumBuffer_gt_BoundCytosol" fast="true">
5     <listOfReactants>
6         <speciesReference species="Ca_C"/>
7         <speciesReference species="B_C"/>
8     </listOfReactants>
9     <listOfProducts>
10        <speciesReference species="CaB_C"/>
11    </listOfProducts>
12    <kineticLaw>
13        <notes>
14            <p xmlns="http://www.w3.org/1999/xhtml">
15                ((Kf_CalciumBuffer_BoundCytosol * Ca_C * B_C) -
16                 (Kr_CalciumBuffer_BoundCytosol * CaB_C))
17            </p>
18        </notes>
19        <math xmlns="http://www.w3.org/1998/Math/MathML">
20            <apply>
21                <times/>
22                <ci> Cytosol</ci>
23            <apply>
24                <minus/>
25                <apply>
26                    <times/>
27                    <ci> Kf_CalciumBuffer_BoundCytosol </ci>
28                    <ci> Ca_C </ci>
29                    <ci> B_C </ci>
30                </apply>
31                <apply>
32                    <times/>
33                    <ci> Kr_CalciumBuffer_BoundCytosol </ci>
34                    <ci> CaB_C </ci>
35                </apply>
36            </apply>
37        </math>
38        <listOfParameters>
39            <parameter id="Kf_CalciumBuffer_BoundCytosol" value="0.1"
40                units="litre_per_mole_per_second"/>
41            <parameter id="Kr_CalciumBuffer_BoundCytosol" value="1.0"
42                units="per_second"/>
43        </listOfParameters>
44    </kineticLaw>
45 </reaction>
46 <reaction id="Ca_Pump">
47     <listOfReactants>
48         <speciesReference species="Ca_C"/>
49     </listOfReactants>
50     <listOfProducts>
51         <speciesReference species="Ca_EC"/>
52     </listOfProducts>
53     <listOfModifiers>
54         <modifierSpeciesReference species="CaPump_PM"/>
55     </listOfModifiers>
56     <kineticLaw>
57         <notes>
58             <p xmlns="http://www.w3.org/1999/xhtml">
59                 ((Vmax * kP * ((Ca_C - Ca_Rest) / (Ca_C + kP)) /
60                  (Ca_Rest + kP)) * CaPump_PM)
61             </p>
62         </notes>
63         <math xmlns="http://www.w3.org/1998/Math/MathML">
64             <apply>
65                 <times/>
66                 <ci> PlasmaMembrane</ci>
67             <apply>
68                 <divide/>

```

```

1          <apply>
2            <times/>
3            <ci> Vmax </ci>
4            <ci> kP </ci>
5            <ci> CaPump_PM </ci>
6            <apply>
7              <minus/>
8              <ci> Ca_C </ci>
9              <ci> Ca_Rest </ci>
10           </apply>
11        </apply>
12      <apply>
13        <times/>
14        <apply>
15          <plus/>
16          <ci> Ca_C </ci>
17          <ci> kP </ci>
18        </apply>
19        <apply>
20          <plus/>
21          <ci> Ca_Rest </ci>
22          <ci> kP </ci>
23        </apply>
24      </apply>
25    </apply>
26  </math>
27 </math>
28 <listOfParameters>
29   <parameter id="Vmax" value="4000" units="per_second"/>
30   <parameter id="kP" value="0.25" units="subs_per_vol"/>
31   <parameter id="Ca_Rest" value="0.1" units="subs_per_vol"/>
32 </listOfParameters>
33 </kineticLaw>
34 </reaction>
35 <reaction id="Ca_channel">
36   <listOfReactants>
37     <speciesReference species="Ca_EC"/>
38   </listOfReactants>
39   <listOfProducts>
40     <speciesReference species="Ca_C"/>
41   </listOfProducts>
42   <listOfModifiers>
43     <modifierSpeciesReference species="CaCh_PM"/>
44   </listOfModifiers>
45   <kineticLaw>
46     <notes>
47       <p xmlns="http://www.w3.org/1999/xhtml">
48         
$$J_0 * K_c * (Ca_{EC} - Ca_C) / (K_c + Ca_C) * CaCh_{PM}$$

49       </p>
50     </notes>
51     <math xmlns="http://www.w3.org/1998/Math/MathML">
52       <apply>
53         <times/>
54         <ci> PlasmaMembrane </ci>
55         <apply>
56           <divide/>
57           <apply>
58             <times/>
59             <ci> CaCh_PM </ci>
60             <ci> J0 </ci>
61             <ci> Kc </ci>
62           </apply>
63           <minus/>
64           <ci> Ca_EC </ci>
65           <ci> Ca_C </ci>
66         </apply>
67       </apply>
68     </math>
69     <plus/>

```

```
1             <ci> Kc </ci>
2             <ci> Ca_C </ci>
3         </apply>
4     </apply>
5 </math>
6 </math>
7 <listOfParameters>
8     <parameter id="J0" value="0.014" units="litre_per_mole_per_second"/>
9     <parameter id="Kc" value="0.5" units="subs_per_vol"/>
10 </listOfParameters>
11 </kineticLaw>
12 </reaction>
13 </listOfReactions>
14 </model>
15 </sbml>
```

8 Discussion

The volume of data now emerging from molecular biotechnology leave little doubt that extensive computer-based modeling, simulation and analysis will be critical to understanding and interpreting the data (Abbott, 1999; Gilman, 2000; Popel and Winslow, 1998; Smaglik, 2000). This has led to an explosion in the development of computer tools by many research groups across the world. The explosive rate of progress is exciting, but the rapid growth of the field is accompanied by problems and pressing needs.

One problem is that simulation models and results often cannot be directly compared, shared or re-used, because the tools developed by different groups often are not compatible with each other. As the field of systems biology matures, researchers increasingly need to communicate their results as computational models rather than box-and-arrow diagrams. They also need to reuse published and curated models as library components in order to succeed with large-scale efforts (e.g., the Alliance for Cellular Signaling; Gilman, 2000; Smaglik, 2000). These needs require that models implemented in one software package be portable to other software packages, to maximize public understanding and to allow building up libraries of curated computational models.

We offer SBML to the systems biology community as a suggested format for exchanging models between simulation/analysis tools. SBML is an open model representation language oriented specifically towards representing systems of biochemical reactions.

Our vision for SBML is to create an open standard that will enable different software tools to exchange computational models. SBML is not static; we continue to develop and experiment with it, and we interact with other groups who seek to develop similar markup languages. We plan on continuing to evolve SBML with the help of the systems biology community to make SBML increasingly more powerful, flexible and useful.

8.1 Future enhancements: SBML Level 3 and beyond

Many people have expressed a desire to see additional capabilities added to SBML. The following summarizes additional features that are under consideration to be included in SBML Level 3; additional information is available in the wiki on the SBML project website (<http://sbml.org>).

- *Arrays*. This will enable the creation of arrays of components (species, reactions, compartments and submodels).
- *Connections*. This will be a mechanism for describing the connections between items in an array.
- *Geometry*. This will enable the encoding of the spatial characteristics of models including the geometry of compartments, the diffusion properties of species and the specification of different species concentrations across different regions of a cell.
- *Model Composition*. This will enable a large model to be built up out of instances of other models. It will also allow the reuse of model components and the creation of several instances of the same model.
- *Multistate and Complex Species*. This will allow the straight-forward construction of models involving species with a large number of states or species composed of subcomponents. The representation scheme would be designed to contain the combinatorial explosion of objects that often results from these types of models.
- *Diagrams*. This feature will allow components to be annotated with data to enable the display of the model in a diagram.
- *Dynamic Structure*. This will enable model structure to vary during simulation. One aspect of this allowing rules and reactions to have their effect conditional on the state of the model system. For example in SBML Level 2 it is possible to create a rule with the effect:

$$\frac{ds}{dt} = \begin{cases} 0 & \text{if } s > 0 \\ y & \text{otherwise} \end{cases}$$

1 Dynamic restructuring would enable the expression of the following example:

$$2 \quad \text{if } s > 0 \quad \frac{ds}{dt} = y$$

3 where s is not determined by the rule when $s \leq 0$.

- 4 • *Tie-breaking algorithm.* This will include a controlled vocabulary and associated **attributes** on models
5 to indicate the simultaneous event tie-breaking algorithm required to correctly simulate the model.
- 6 • *Distributions.* This will provide a means of specifying random variables and statistical distribution of
7 values.

Acknowledgments

This document benefited from repeated reviews and feedback by many people. We give special thanks to James Schaff, Martin Ginkel, Sven Sahle, Ralph Gauges, Ion Moraru, Rainer Machné, Robert Phair, Ben Bornstein, Bruce Shapiro, Maria Schilstra, and Herbert Sauro for critical discussions of many issues, including the implementation of software supporting SBML and the simulation of models.

The development of SBML was originally funded entirely by the Japan Science and Technology Agency (JST) under the ERATO Kitano Symbiotic Systems Project during the years 2000–2003. The principal investigators were Hiroaki Kitano and John C. Doyle. The original SBML Team was lead by Hamid Bolouri and consisted of Hamid Bolouri, Andrew Finney, Herbert Sauro, and Michael Hucka.

We gratefully acknowledge sponsorship from many funding agencies. Support for the continued development of SBML and associated software, meetings and activities comes from the following sources: the National Institute of General Medical Sciences (USA) via grant numbers GM070923 and GM077671; the National Human Genome Research Institute (USA); the International Joint Research Program of NEDO (Japan); the JST ERATO-SORST Program (Japan); the Japanese Ministry of Agriculture; the Japanese Ministry of Education, Culture, Sports, Science and Technology; the BBSRC e-Science Initiative (UK); the DARPA IPTO Bio-Computation Program (USA); and the Air Force Office of Scientific Research (USA). Additional support has been or continues to be provided by the California Institute of Technology (USA), the University of Hertfordshire (UK), the Molecular Sciences Institute (USA), and the Systems Biology Institute (Japan).

SBML was first conceived at the JST/ERATO-sponsored *First Workshop on Software Platforms for Systems Biology*, held in April, 2000, at the California Institute of Technology in Pasadena, California, USA. The participants collectively decided to begin developing a common XML-based declarative language for representing models. A draft version of the Systems Biology Markup Language was developed by the Caltech ERATO team and delivered to all collaborators in August, 2000. This draft version underwent extensive discussion over mailing lists and then again during the *Second Workshop on Software Platforms for Systems Biology* held in Tokyo, Japan, November 2000. A revised version of SBML was issued by the Caltech ERATO team in December, 2000, and after further discussions over mailing lists and in meetings, we produced a specification for SBML Level 1 (Hucka et al., 2001).

SBML Level 2 was conceived at the *5th Workshop on Software Platforms for Systems Biology*, held in July 2002, at the University of Hertfordshire, UK. The participants collectively decided to revise the form of SBML in SBML Level 2. The first draft of the Level 2 Version 1 document was released in August 2002. The final set of features in SBML Level 2 Version 1 was finalized in May 2003 at the *7th Workshop on Software Platforms for Systems Biology* in Ft. Lauderdale, Florida. SBML Level 2 Version 2 was largely finalized after the 2005 SBML Forum meeting in Boston and a final document was issued in September 2006.

SBML Level 2 Version 3 was finalized after the 2006 SBML Forum meeting in Yokohama, Japan, and the 2007 SBML Hackathon in Newcastle, UK. It was developed with contributions from so many people constituting the worldwide *SBML Forum* that we regret it has become infeasible to list individuals by name. We are grateful to everyone on the sbml-discuss@caltech.edu and libsbml-discuss@caltech.edu mailing lists, the creators of CellML (Hedley et al., 2001), the members of the DARPA Bio-SPICE project, and the authors of the following software systems: BALSAM, BASIS, BIOCHAM, BioCharon, ByoDyn, BioCyc, BioGrid, BioModels, BioNetGen, BioPathwise, Bio Sketch Pad, BioSens, BioSPICE Dashboard, BioSpreadsheet, BioTapestry, BioUML, BSTLab, CADLIVE, CellDesigner, Cellerator, CellML2SBML, Cellware, CL-SBML, CLEML, COPASI, Cyto-Sim, Cytoscape, DBsolve, Dizzy, E-CELL, ecellJ, ESS, FluxAnalyzer, Fluxor, Gepasi, Gillespie2, HSMB, HybridSBML, INSILICO discovery, JACOBIAN, Jarnac, JDesigner, JigCell, JSim, JWS Online, KEGG2SBML, Kineticon, libSBML, MathSBML, MIRIAM, MesoRD, MetaFluxNet, MMT2, Modesto, Molecuizer, Monod, Narrator, NetBuilder, Oscill8, PANTHER Pathway, PathArt, Pathway Analyzer, PathwayLab, PathScout, Pathway Tools, PathwayBuilder, PATIKAwab, PaVESy, PET, PNK, PottersWheel, Reactome, ProcessDB, PROTON, pysbml, PySCeS, RSBML, runSBML, SABIO-RK, SBML ODE Solver, SBML-PET, SBMLeditor, SBMLmerge, SBMLR, SBMLSim, SBMLToolbox, SBliD, SBO, SBToolbox, SBW, SCiPath, Sigmoid, SigPath, SigTran, SIMBA, SimBiology, Simpathica, SimPheny, SimWiz, SloppyCell, SmartCell, SRS Pathway Editor, StochSim, StochKit, STOCKS, TERANODE Suite, Trellis, VANTED, Virtual Cell, WebCell, WinSCAMP, Xholon, and XPPAUT.

A XML Schema for SBML

The following is an XML Schema definition for SBML Level 2 Version 3, using the W3C Recommendation for XML Schema version 1.0 of 2 May 2001 (Biron and Malhotra, 2000; Fallside, 2000; Thompson et al., 2000). This Schema does not define all aspects of SBML Level 2: an SBML document validated by this schema is not necessarily a valid SBML Level 2 document. Appendix B contains a schema for the SBML MathML subset. Appendix C contains a list of the remaining checks required to validate a model in addition to making it consistent with these two schemas.

Note to implementors: the following schema is self-contained and makes reference to the official XML Schema for MathML hosted at the W3. However, for use in software systems, it is more efficient to store the MathML subset Schema of Appendix B in a file on a user's local disk, and change the `schemaLocation` value on line 27 below to refer to this local copy of the MathML subset Schema. Doing so will avoid requiring a network access every time this SBML Schema is used.

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema targetNamespace="http://www.sbml.org/sbml/level2/version3"
  xmlns="http://www.sbml.org/sbml/level2/version3"
  xmlns:mml="http://www.w3.org/1998/Math/MathML"
  xmlns:xlink="http://www.w3.org/1999/xlink"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified"
  version="$ $Id: sbml.xsd,v 1.3 2007/03/14 22:23:43 mhucka Exp $ $">
  <xsd:import namespace="http://www.w3.org/1998/Math/MathML"
    schemaLocation="http://www.w3.org/Math/XMLSchema/mathml2/mathml2.xsd"/>
  <xsd:annotation>
    <xsd:documentation>
      * Filename : sbml.xsd
      * Description: XML Schema for SBML Level 2 Version 3.
      * Author(s) : Michael Hucka
      * Revision : $ $Id: sbml.xsd,v 1.3 2007/03/14 22:23:43 mhucka Exp $ $
      * $ $Source: /cvsroot/sbml/specifications/sbml-level-2/version-3/schema/sbml.xsd,v $ $
      *
      * Copyright 2007 California Institute of Technology.
      * Copyright 2003-2006 California Institute of Technology, the Japan Science
      * and Technology Corporation, and the University of Hertfordshire.
      *
      * This software is licensed according to the terms described in the file
      * named "LICENSE.txt" included with this distribution and available
      * online at http://sbml.org/xml-schemas/LICENSE.txt
    </xsd:documentation>
  </xsd:annotation>
  <!--The definition of new primitive types follows.-->
  <xsd:simpleType name="SID">
    <xsd:annotation>
      <xsd:documentation>The type SID is used throughout SBML as the
        type of the 'id' attributes on model elements.</xsd:documentation>
    </xsd:annotation>
    <xsd:restriction base="xsd:string">
      <xsd:pattern value="(?![a-z]|[A-Z])(?![a-z]|[A-Z])[0-9]*/">
    </xsd:restriction>
  </xsd:simpleType>
  <xsd:simpleType name="UnitSID">
    <xsd:annotation>
      <xsd:documentation>The type UnitSID is used to refer to units.</xsd:documentation>
    </xsd:annotation>
    <xsd:union>
      <xsd:simpleType>
        <xsd:restriction base="SID" />
      </xsd:simpleType>
      <xsd:simpleType>
        <xsd:restriction base="SID">
          <xsd:enumeration value="ampere"/>
          <xsd:enumeration value="becquerel"/>
          <xsd:enumeration value="candela"/>
          <xsd:enumeration value="coulomb"/>
          <xsd:enumeration value="dimensionless"/>
          <xsd:enumeration value="farad"/>
          <xsd:enumeration value="gram"/>
          <xsd:enumeration value="gray"/>
          <xsd:enumeration value="henry"/>
          <xsd:enumeration value="hertz"/>
        </xsd:restriction>
      </xsd:simpleType>
    </xsd:union>
  </xsd:simpleType>

```

```

1      <xsd:enumeration value="item"/>
2      <xsd:enumeration value="joule"/>
3      <xsd:enumeration value="katal"/>
4      <xsd:enumeration value="kelvin"/>
5      <xsd:enumeration value="kilogram"/>
6      <xsd:enumeration value="litre"/>
7      <xsd:enumeration value="lumen"/>
8      <xsd:enumeration value="lux"/>
9      <xsd:enumeration value="metre"/>
10     <xsd:enumeration value="mole"/>
11     <xsd:enumeration value="newton"/>
12     <xsd:enumeration value="ohm"/>
13     <xsd:enumeration value="pascal"/>
14     <xsd:enumeration value="radian"/>
15     <xsd:enumeration value="second"/>
16     <xsd:enumeration value="siemens"/>
17     <xsd:enumeration value="sievert"/>
18     <xsd:enumeration value="steradian"/>
19     <xsd:enumeration value="tesla"/>
20     <xsd:enumeration value="volt"/>
21     <xsd:enumeration value="watt"/>
22     <xsd:enumeration value="weber"/>
23   </xsd:restriction>
24 </xsd:simpleType>
25 <xsd:simpleType>
26   <xsd:restriction base="SID">
27     <xsd:enumeration value="substance"/>
28     <xsd:enumeration value="volume"/>
29     <xsd:enumeration value="area"/>
30     <xsd:enumeration value="length"/>
31     <xsd:enumeration value="time"/>
32   </xsd:restriction>
33 </xsd:simpleType>
34 </xsd:union>
35 </xsd:simpleType>
36 <xsd:simpleType name="SBOTerm">
37   <xsd:annotation>
38     <xsd:documentation>The data type for sboTerm attribute values.</xsd:documentation>
39   </xsd:annotation>
40   <xsd:restriction base="xsd:string">
41     <xsd:pattern value="(SB0:){[0-9]7}" />
42   </xsd:restriction>
43 </xsd:simpleType>
44 <!--The definition of SBase follows.-->
45 <xsd:complexType name="SBase" abstract="true">
46   <xsd:annotation>
47     <xsd:documentation>The SBase type is the base type of all main
48       components in SBML. It supports attaching metadata, notes and
49       annotations to components.</xsd:documentation>
50   </xsd:annotation>
51   <xsd:sequence>
52     <xsd:element name="notes" minOccurs="0">
53       <xsd:complexType>
54         <xsd:sequence>
55           <xsd:any namespace="http://www.w3.org/1999/xhtml"
56             processContents="skip"
57             minOccurs="0" maxOccurs="unbounded" />
58         </xsd:sequence>
59       </xsd:complexType>
60     </xsd:element>
61     <xsd:element name="annotation" minOccurs="0">
62       <xsd:complexType>
63         <xsd:sequence>
64           <xsd:any processContents="skip"
65             minOccurs="0" maxOccurs="unbounded" />
66         </xsd:sequence>
67       </xsd:complexType>
68     </xsd:element>
69   </xsd:sequence>
70   <xsd:attribute name="metaid" type="xsd:ID" use="optional"/>
71   <xsd:attribute name="sboTerm" type="SBOTerm" use="optional"/>
72 </xsd:complexType>
73 <!--The definition of main SBML classes follows.-->
74 <xsd:complexType name="FunctionDefinition">
75   <xsd:complexContent>
76     <xsd:extension base="SBase">
77       <xsd:sequence>
78         <xsd:element ref="mml:math"/>
79       </xsd:sequence>
80     </xsd:extension>
81   </xsd:complexContent>
82   <xsd:attribute name="id" type="SID" use="required"/>

```



```

1         <xsd:attribute name="name" type="xsd:string" use="optional"/>
2     </xsd:extension>
3 </xsd:complexType>
4 </xsd:complexType>
5 <xsd:complexType name="Unit">
6     <xsd:complexContent>
7         <xsd:extension base="SBase">
8             <xsd:attribute name="kind" type="UnitSID" use="required"/>
9             <xsd:attribute name="exponent" type="xsd:int" default="1"/>
10            <xsd:attribute name="scale" type="xsd:int" default="0"/>
11            <xsd:attribute name="multiplier" type="xsd:double" default="1"/>
12        </xsd:extension>
13    </xsd:complexContent>
14 </xsd:complexType>
15 <xsd:complexType name="ListOfUnits">
16     <xsd:complexContent>
17         <xsd:extension base="SBase">
18             <xsd:sequence>
19                 <xsd:element name="unit" type="Unit" maxOccurs="unbounded"/>
20             </xsd:sequence>
21         </xsd:extension>
22    </xsd:complexContent>
23 </xsd:complexType>
24 <xsd:complexType name="UnitDefinition">
25     <xsd:complexContent>
26         <xsd:extension base="SBase">
27             <xsd:sequence>
28                 <xsd:element name="listOfUnits" type="ListOfUnits"/>
29             </xsd:sequence>
30             <xsd:attribute name="id" type="UnitSID" use="required"/>
31             <xsd:attribute name="name" type="xsd:string" use="optional"/>
32         </xsd:extension>
33    </xsd:complexContent>
34 </xsd:complexType>
35 <xsd:complexType name="CompartmentType">
36     <xsd:complexContent>
37         <xsd:extension base="SBase">
38             <xsd:attribute name="id" type="SID" use="required"/>
39             <xsd:attribute name="name" type="xsd:string" use="optional"/>
40         </xsd:extension>
41    </xsd:complexContent>
42 </xsd:complexType>
43 <xsd:complexType name="SpeciesType">
44     <xsd:complexContent>
45         <xsd:extension base="SBase">
46             <xsd:attribute name="id" type="SID" use="required"/>
47             <xsd:attribute name="name" type="xsd:string" use="optional"/>
48         </xsd:extension>
49    </xsd:complexContent>
50 </xsd:complexType>
51 <xsd:complexType name="Compartment">
52     <xsd:complexContent>
53         <xsd:extension base="SBase">
54             <xsd:attribute name="id" type="SID" use="required"/>
55             <xsd:attribute name="name" type="xsd:string" use="optional"/>
56             <xsd:attribute name="compartmentType" type="SID" use="optional"/>
57             <xsd:attribute name="spatialDimensions" use="optional" default="3">
58                 <xsd:simpleType>
59                     <xsd:restriction base="xsd:int">
60                         <xsd:minInclusive value="0"/>
61                         <xsd:maxInclusive value="3"/>
62                     </xsd:restriction>
63                 </xsd:simpleType>
64             </xsd:attribute>
65             <xsd:attribute name="size" type="xsd:double" use="optional"/>
66             <xsd:attribute name="units" type="UnitSID" use="optional"/>
67             <xsd:attribute name="outside" type="SID" use="optional"/>
68             <xsd:attribute name="constant" type="xsd:boolean"
69                 use="optional" default="true"/>
70         </xsd:extension>
71    </xsd:complexContent>
72 </xsd:complexType>
73 <xsd:complexType name="Species">
74     <xsd:complexContent>
75         <xsd:extension base="SBase">
76             <xsd:attribute name="id" type="SID" use="required"/>
77             <xsd:attribute name="name" type="xsd:string" use="optional"/>
78             <xsd:attribute name="speciesType" type="SID" use="optional"/>
79             <xsd:attribute name="compartment" type="SID" use="required"/>
80             <xsd:attribute name="initialAmount" type="xsd:double" use="optional"/>

```

```

1         <xsd:attribute name="initialConcentration" type="xsd:double" use="optional"/>
2         <xsd:attribute name="substanceUnits" type="UnitSid" use="optional"/>
3         <xsd:attribute name="hasOnlySubstanceUnits" type="xsd:boolean"
4             use="optional" default="false"/>
5         <xsd:attribute name="boundaryCondition" type="xsd:boolean"
6             use="optional" default="false"/>
7         <xsd:attribute name="charge" type="xsd:int" use="optional"/>
8         <xsd:attribute name="constant" type="xsd:boolean"
9             use="optional" default="false"/>
10        </xsd:extension>
11    </xsd:complexContent>
12 </xsd:complexType>
13 <xsd:complexType name="Parameter">
14     <xsd:complexContent>
15         <xsd:extension base="SBase">
16             <xsd:attribute name="id" type="Sid" use="required"/>
17             <xsd:attribute name="name" type="xsd:string" use="optional"/>
18             <xsd:attribute name="value" type="xsd:double" use="optional"/>
19             <xsd:attribute name="units" type="UnitSid" use="optional"/>
20             <xsd:attribute name="constant" type="xsd:boolean"
21                 use="optional" default="true"/>
22         </xsd:extension>
23     </xsd:complexContent>
24 </xsd:complexType>
25 <xsd:complexType name="ListOfParameters">
26     <xsd:complexContent>
27         <xsd:extension base="SBase">
28             <xsd:sequence>
29                 <xsd:element name="parameter" type="Parameter" maxOccurs="unbounded"/>
30             </xsd:sequence>
31         </xsd:extension>
32     </xsd:complexContent>
33 </xsd:complexType>
34 <xsd:complexType name="InitialAssignment">
35     <xsd:complexContent>
36         <xsd:extension base="SBase">
37             <xsd:sequence>
38                 <xsd:element ref="mml:math"/>
39             </xsd:sequence>
40             <xsd:attribute name="symbol" type="Sid" use="required"/>
41         </xsd:extension>
42     </xsd:complexContent>
43 </xsd:complexType>
44 <xsd:complexType name="Rule" abstract="true">
45     <xsd:complexContent>
46         <xsd:extension base="SBase">
47             <xsd:sequence>
48                 <xsd:element ref="mml:math"/>
49             </xsd:sequence>
50         </xsd:extension>
51     </xsd:complexContent>
52 </xsd:complexType>
53 <xsd:complexType name="AlgebraicRule">
54     <xsd:complexContent>
55         <xsd:extension base="Rule"/>
56     </xsd:complexContent>
57 </xsd:complexType>
58 <xsd:complexType name="AssignmentRule">
59     <xsd:complexContent>
60         <xsd:extension base="Rule">
61             <xsd:attribute name="variable" type="Sid" use="required"/>
62         </xsd:extension>
63     </xsd:complexContent>
64 </xsd:complexType>
65 <xsd:complexType name="RateRule">
66     <xsd:complexContent>
67         <xsd:extension base="Rule">
68             <xsd:attribute name="variable" type="Sid" use="required"/>
69         </xsd:extension>
70     </xsd:complexContent>
71 </xsd:complexType>
72 <xsd:complexType name="Constraint">
73     <xsd:complexContent>
74         <xsd:extension base="SBase">
75             <xsd:sequence>
76                 <xsd:element ref="mml:math"/>
77                 <xsd:element name="message" minOccurs="0">
78                     <xsd:complexType>
79                         <xsd:sequence>
80                             <xsd:any namespace="http://www.w3.org/1999/xhtml"

```

```

1                                     processContents="skip"
2                                     minOccurs="0" maxOccurs="unbounded"/>
3                                 </xsd:sequence>
4                             </xsd:complexType>
5                         </xsd:element>
6                     </xsd:sequence>
7                 </xsd:extension>
8             </xsd:complexContent>
9         </xsd:complexType>
10    <xsd:complexType name="KineticLaw">
11        <xsd:complexContent>
12            <xsd:extension base="SBase">
13                <xsd:sequence>
14                    <xsd:element ref="mml:math"/>
15                    <xsd:element name="listOfParameters" type="ListOfParameters" minOccurs="0"/>
16                </xsd:sequence>
17            </xsd:extension>
18        </xsd:complexContent>
19    </xsd:complexType>
20    <xsd:complexType name="SimpleSpeciesReference" abstract="true">
21        <xsd:complexContent>
22            <xsd:extension base="SBase">
23                <xsd:attribute name="id" type="SId" use="optional"/>
24                <xsd:attribute name="name" type="xsd:string" use="optional"/>
25                <xsd:attribute name="species" type="SId" use="required"/>
26            </xsd:extension>
27        </xsd:complexContent>
28    </xsd:complexType>
29    <xsd:complexType name="ModifierSpeciesReference">
30        <xsd:complexContent>
31            <xsd:extension base="SimpleSpeciesReference"/>
32        </xsd:complexContent>
33    </xsd:complexType>
34    <xsd:complexType name="ListOfModifierSpeciesReferences">
35        <xsd:complexContent>
36            <xsd:extension base="SBase">
37                <xsd:sequence>
38                    <xsd:element name="modifierSpeciesReference"
39                                type="ModifierSpeciesReference"
40                                maxOccurs="unbounded"/>
41                </xsd:sequence>
42            </xsd:extension>
43        </xsd:complexContent>
44    </xsd:complexType>
45    <xsd:complexType name="StoichiometryMath">
46        <xsd:complexContent>
47            <xsd:extension base="SBase">
48                <xsd:sequence>
49                    <xsd:element ref="mml:math"/>
50                </xsd:sequence>
51            </xsd:extension>
52        </xsd:complexContent>
53    </xsd:complexType>
54    <xsd:complexType name="SpeciesReference">
55        <xsd:complexContent>
56            <xsd:extension base="SimpleSpeciesReference">
57                <xsd:sequence>
58                    <xsd:element name="stoichiometryMath"
59                                type="StoichiometryMath"
60                                minOccurs="0"/>
61                </xsd:sequence>
62                <xsd:attribute name="stoichiometry" type="xsd:double"
63                            use="optional" default="1"/>
64            </xsd:extension>
65        </xsd:complexContent>
66    </xsd:complexType>
67    <xsd:complexType name="ListOfSpeciesReferences">
68        <xsd:complexContent>
69            <xsd:extension base="SBase">
70                <xsd:sequence>
71                    <xsd:element name="speciesReference"
72                                type="SpeciesReference"
73                                maxOccurs="unbounded"/>
74                </xsd:sequence>
75            </xsd:extension>
76        </xsd:complexContent>
77    </xsd:complexType>
78    <xsd:complexType name="Reaction">
79        <xsd:complexContent>
80            <xsd:extension base="SBase">

```

```

1         <xsd:sequence>
2             <xsd:element name="listOfReactants" type="ListOfSpeciesReferences"
3                 minOccurs="0"/>
4             <xsd:element name="listOfProducts" type="ListOfSpeciesReferences"
5                 minOccurs="0"/>
6             <xsd:element name="listOfModifiers" type="ListOfModifierSpeciesReferences"
7                 minOccurs="0"/>
8             <xsd:element name="kineticLaw" type="KineticLaw"
9                 minOccurs="0"/>
10        </xsd:sequence>
11        <xsd:attribute name="id" type="SId" use="required"/>
12        <xsd:attribute name="name" type="xsd:string" use="optional"/>
13        <xsd:attribute name="reversible" type="xsd:boolean"
14            use="optional" default="true"/>
15        <xsd:attribute name="fast" type="xsd:boolean"
16            use="optional" default="false"/>
17    </xsd:extension>
18 </xsd:complexContent>
19 </xsd:complexType>
20 <xsd:complexType name="EventAssignment">
21     <xsd:complexContent>
22         <xsd:extension base="SBase">
23             <xsd:sequence>
24                 <xsd:element ref="mml:math"/>
25             </xsd:sequence>
26             <xsd:attribute name="variable" type="SId" use="required"/>
27         </xsd:extension>
28     </xsd:complexContent>
29 </xsd:complexType>
30 <xsd:complexType name="ListOfEventAssignments">
31     <xsd:complexContent>
32         <xsd:extension base="SBase">
33             <xsd:sequence>
34                 <xsd:element name="eventAssignment"
35                     type="EventAssignment"
36                     maxOccurs="unbounded"/>
37             </xsd:sequence>
38         </xsd:extension>
39     </xsd:complexContent>
40 </xsd:complexType>
41 <xsd:complexType name="Trigger">
42     <xsd:complexContent>
43         <xsd:extension base="SBase">
44             <xsd:sequence>
45                 <xsd:element ref="mml:math"/>
46             </xsd:sequence>
47         </xsd:extension>
48     </xsd:complexContent>
49 </xsd:complexType>
50 <xsd:complexType name="Delay">
51     <xsd:complexContent>
52         <xsd:extension base="SBase">
53             <xsd:sequence>
54                 <xsd:element ref="mml:math"/>
55             </xsd:sequence>
56         </xsd:extension>
57     </xsd:complexContent>
58 </xsd:complexType>
59 <xsd:complexType name="Event">
60     <xsd:complexContent>
61         <xsd:extension base="SBase">
62             <xsd:sequence>
63                 <xsd:element name="trigger" type="Trigger"/>
64                 <xsd:element name="delay" type="Delay" minOccurs="0"/>
65                 <xsd:element name="listOfEventAssignments" type="ListOfEventAssignments"/>
66             </xsd:sequence>
67             <xsd:attribute name="id" type="SId" use="optional"/>
68             <xsd:attribute name="name" type="xsd:string" use="optional"/>
69         </xsd:extension>
70     </xsd:complexContent>
71 </xsd:complexType>
72 <xsd:complexType name="Model">
73     <xsd:complexContent>
74         <xsd:extension base="SBase">
75             <xsd:sequence>
76                 <xsd:element name="listOfFunctionDefinitions" minOccurs="0">
77                     <xsd:complexType>
78                         <xsd:complexContent>
79                             <xsd:extension base="SBase">
80                                 <xsd:sequence>

```

```

1         <xsd:element name="functionDefinition"
2             type="FunctionDefinition"
3             maxOccurs="unbounded"/>
4     </xsd:sequence>
5 </xsd:extension>
6 </xsd:complexContent>
7 </xsd:complexType>
8 </xsd:element>
9 <xsd:element name="listOfUnitDefinitions" minOccurs="0">
10 <xsd:complexType>
11 <xsd:complexContent>
12 <xsd:extension base="SBase">
13 <xsd:sequence>
14 <xsd:element name="unitDefinition"
15     type="UnitDefinition"
16     maxOccurs="unbounded"/>
17 </xsd:sequence>
18 </xsd:extension>
19 </xsd:complexContent>
20 </xsd:complexType>
21 </xsd:element>
22 <xsd:element name="listOfCompartmentTypes" minOccurs="0">
23 <xsd:complexType>
24 <xsd:complexContent>
25 <xsd:extension base="SBase">
26 <xsd:sequence>
27 <xsd:element name="compartmentType"
28     type="CompartmentType"
29     maxOccurs="unbounded"/>
30 </xsd:sequence>
31 </xsd:extension>
32 </xsd:complexContent>
33 </xsd:complexType>
34 </xsd:element>
35 <xsd:element name="listOfSpeciesTypes" minOccurs="0">
36 <xsd:complexType>
37 <xsd:complexContent>
38 <xsd:extension base="SBase">
39 <xsd:sequence>
40 <xsd:element name="speciesType"
41     type="SpeciesType"
42     maxOccurs="unbounded"/>
43 </xsd:sequence>
44 </xsd:extension>
45 </xsd:complexContent>
46 </xsd:complexType>
47 </xsd:element>
48 <xsd:element name="listOfCompartments" minOccurs="0">
49 <xsd:complexType>
50 <xsd:complexContent>
51 <xsd:extension base="SBase">
52 <xsd:sequence>
53 <xsd:element name="compartment"
54     type="Compartment"
55     maxOccurs="unbounded"/>
56 </xsd:sequence>
57 </xsd:extension>
58 </xsd:complexContent>
59 </xsd:complexType>
60 </xsd:element>
61 <xsd:element name="listOfSpecies" minOccurs="0">
62 <xsd:complexType>
63 <xsd:complexContent>
64 <xsd:extension base="SBase">
65 <xsd:sequence>
66 <xsd:element name="species"
67     type="Species"
68     maxOccurs="unbounded"/>
69 </xsd:sequence>
70 </xsd:extension>
71 </xsd:complexContent>
72 </xsd:complexType>
73 </xsd:element>
74 <xsd:element name="listOfParameters" type="ListOfParameters" minOccurs="0"/>
75 <xsd:element name="listOfInitialAssignments" minOccurs="0">
76 <xsd:complexType>
77 <xsd:complexContent>
78 <xsd:extension base="SBase">
79 <xsd:sequence>
80 <xsd:element name="initialAssignment"

```

```

1                                     type="InitialAssignment"
2                                     maxOccurs="unbounded"/>
3                                     </xsd:sequence>
4                                 </xsd:extension>
5                                 </xsd:complexContent>
6                                 </xsd:complexType>
7                             </xsd:element>
8                             <xsd:element name="listOfRules" minOccurs="0">
9                                 <xsd:complexType>
10                                    <xsd:complexContent>
11                                        <xsd:extension base="SBase">
12                                            <xsd:choice maxOccurs="unbounded">
13                                                <xsd:element name="algebraicRule"
14                                                    type="AlgebraicRule"
15                                                    minOccurs="0"/>
16                                                <xsd:element name="assignmentRule"
17                                                    type="AssignmentRule"
18                                                    minOccurs="0"/>
19                                                <xsd:element name="rateRule"
20                                                    type="RateRule"
21                                                    minOccurs="0"/>
22                                            </xsd:choice>
23                                        </xsd:extension>
24                                    </xsd:complexContent>
25                                </xsd:complexType>
26                            </xsd:element>
27                            <xsd:element name="listOfConstraints" minOccurs="0">
28                                <xsd:complexType>
29                                    <xsd:complexContent>
30                                        <xsd:extension base="SBase">
31                                            <xsd:sequence>
32                                                <xsd:element name="constraint"
33                                                    type="Constraint"
34                                                    maxOccurs="unbounded"/>
35                                            </xsd:sequence>
36                                        </xsd:extension>
37                                    </xsd:complexContent>
38                                </xsd:complexType>
39                            </xsd:element>
40                            <xsd:element name="listOfReactions" minOccurs="0">
41                                <xsd:complexType>
42                                    <xsd:complexContent>
43                                        <xsd:extension base="SBase">
44                                            <xsd:sequence>
45                                                <xsd:element name="reaction"
46                                                    type="Reaction"
47                                                    maxOccurs="unbounded"/>
48                                            </xsd:sequence>
49                                        </xsd:extension>
50                                    </xsd:complexContent>
51                                </xsd:complexType>
52                            </xsd:element>
53                            <xsd:element name="listOfEvents" minOccurs="0">
54                                <xsd:complexType>
55                                    <xsd:complexContent>
56                                        <xsd:extension base="SBase">
57                                            <xsd:sequence>
58                                                <xsd:element name="event"
59                                                    type="Event"
60                                                    maxOccurs="unbounded"/>
61                                            </xsd:sequence>
62                                        </xsd:extension>
63                                    </xsd:complexContent>
64                                </xsd:complexType>
65                            </xsd:element>
66                        </xsd:sequence>
67                        <xsd:attribute name="id" type="SId" use="optional"/>
68                        <xsd:attribute name="name" type="xsd:string" use="optional"/>
69                    </xsd:extension>
70                </xsd:complexContent>
71            </xsd:complexType>
72            <!-- The following is the type definition for the top-level element in an SBML document.-->
73            <xsd:complexType name="Sbml">
74                <xsd:complexContent>
75                    <xsd:extension base="SBase">
76                        <xsd:sequence>
77                            <xsd:element name="model" type="Model"/>
78                        </xsd:sequence>
79                        <xsd:attribute name="level" type="xsd:positiveInteger"
80                            use="required" fixed="2"/>

```

```
1         <xsd:attribute name="version" type="xsd:positiveInteger"  
2             use="required" fixed="3"/>  
3     </xsd:extension>  
4 </xsd:complexContent>  
5 </xsd:complexType>  
6 <!--The following is the (only) top-level element allowed in an SBML document.-->  
7 <xsd:element name="sbml" type="Sbml"/>  
8 <!--The end.-->  
9 </xsd:schema>
```

B XML Schema for MathML subset

The following XML schema defines the syntax of the MathML syntax that is used in [SBML Level 2](#).

```
1
2
3
4 <?xml version="1.0" encoding="UTF-8"?>
5 <!--
6 * Filename      : sbml-mathml.xsd
7 * Description   : Schema for the MathML subset used by SBML L2.
8 * Author(s)    : Andrew Finney, Michael Hucka
9 * Organization : SBML Team <sbml-team@caltech.edu>
10 * Revision     : $ Id: sbml-mathml.xsd,v 1.3 2006/07/20 23:46:31 mhucka Exp $
11 * Source      : $ Source: /cvsroot/sbml/specifications/sbml-mathml/sbml-mathml.xsd,v $
12 *
13 * Copyright 2003-2006 California Institute of Technology, the Japan Science
14 * and Technology Corporation, and the University of Hertfordshire.
15 *
16 * This software is licensed according to the terms described in the file
17 * named "LICENSE.txt" included with this distribution and available
18 * online at http://sbml.org/xml-schemas/LICENSE.txt
19 *
20 * Summary:
21 *
22 * This is a reduced version of the XML Schema for MathML 2.0. It
23 * corresponds to the subset of MathML 2.0 used in SBML Level 2, and
24 * should be used by validating XML parsers instead of the actual
25 * MathML XML Schema when validating SBML files. To accomplish that,
26 * changed the value of the attribute 'schemaLocation' in the SBML
27 * XML Schema file (sbml.xsd) to refer to a copy of this Schema file
28 * on your computer's local hard disk.
29 -->
30 <xs:schema xmlns="http://www.w3.org/1998/Math/MathML"
31           xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
32           xmlns:xs="http://www.w3.org/2001/XMLSchema"
33           targetNamespace="http://www.w3.org/1998/Math/MathML"
34           elementFormDefault="qualified"
35           attributeFormDefault="unqualified">
36   <xs:attributeGroup name="MathAttributes">
37     <xs:attribute name="class" type="xs:NMTOKENS" use="optional"/>
38     <xs:attribute name="style" type="xs:string" use="optional"/>
39     <xs:attribute name="id" type="xs:ID" use="optional"/>
40   </xs:attributeGroup>
41   <xs:complexType name="MathBase">
42     <xs:attributeGroup ref="MathAttributes"/>
43   </xs:complexType>
44   <xs:attributeGroup name="CnAttributes">
45     <xs:attribute name="type">
46       <xs:simpleType>
47         <xs:restriction base="xs:NMTOKEN">
48           <xs:enumeration value="e-notation"/>
49           <xs:enumeration value="integer"/>
50           <xs:enumeration value="rational"/>
51           <xs:enumeration value="real"/>
52         </xs:restriction>
53       </xs:simpleType>
54     </xs:attribute>
55     <xs:attributeGroup ref="MathAttributes"/>
56   </xs:attributeGroup>
57   <xs:complexType name="SepType"/>
58   <xs:complexType name="Cn" mixed="true">
59     <xs:choice minOccurs="0">
60       <xs:element name="sep" type="SepType"/>
61     </xs:choice>
62     <xs:attributeGroup ref="CnAttributes"/>
63   </xs:complexType>
64   <xs:complexType name="Ci">
65     <xs:simpleContent>
66       <xs:extension base="xs:string">
67         <xs:attributeGroup ref="MathAttributes"/>
68       </xs:extension>
69     </xs:simpleContent>
70   </xs:complexType>
71   <xs:simpleType name="CsymbolURI">
72     <xs:restriction base="xs:string">
73       <xs:enumeration value="http://www.sbml.org/sbml/symbols/time"/>
74       <xs:enumeration value="http://www.sbml.org/sbml/symbols/delay"/>
75     </xs:restriction>
76   </xs:simpleType>
77   <xs:complexType name="Csymbol">
78     <xs:simpleContent
```



```

1      <xs:extension base="xs:string">
2          <xs:attribute name="encoding" use="required" fixed="text"/>
3          <xs:attribute name="definitionURL" type="CsymbolURI" use="required"/>
4          <xs:attributeGroup ref="MathAttributes"/>
5      </xs:extension>
6  </xs:simpleContent>
7 </xs:complexType>
8 <xs:complexType name="NodeContainer">
9     <xs:complexContent>
10        <xs:extension base="MathBase">
11            <xs:group ref="Node"/>
12        </xs:extension>
13    </xs:complexContent>
14 </xs:complexType>
15 <xs:complexType name="Apply">
16     <xs:complexContent>
17        <xs:extension base="MathBase">
18            <xs:sequence>
19                <xs:choice>
20                    <xs:element name="ci" type="Ci"/>
21                    <xs:element name="csymbol" type="Csymbol"/>
22                    <xs:element name="eq" type="MathBase"/>
23                    <xs:element name="neq" type="MathBase"/>
24                    <xs:element name="gt" type="MathBase"/>
25                    <xs:element name="lt" type="MathBase"/>
26                    <xs:element name="geq" type="MathBase"/>
27                    <xs:element name="leq" type="MathBase"/>
28                    <xs:element name="plus" type="MathBase"/>
29                    <xs:element name="minus" type="MathBase"/>
30                    <xs:element name="times" type="MathBase"/>
31                    <xs:element name="divide" type="MathBase"/>
32                    <xs:element name="power" type="MathBase"/>
33                    <xs:sequence>
34                        <xs:element name="root" type="MathBase"/>
35                        <xs:element name="degree" type="NodeContainer" minOccurs="0"/>
36                    </xs:sequence>
37                    <xs:element name="abs" type="MathBase"/>
38                    <xs:element name="exp" type="MathBase"/>
39                    <xs:element name="ln" type="MathBase"/>
40                    <xs:sequence>
41                        <xs:element name="log" type="MathBase"/>
42                        <xs:element name="logbase" type="NodeContainer" minOccurs="0"/>
43                    </xs:sequence>
44                    <xs:element name="floor" type="MathBase"/>
45                    <xs:element name="ceiling" type="MathBase"/>
46                    <xs:element name="factorial" type="MathBase"/>
47                    <xs:element name="and" type="MathBase"/>
48                    <xs:element name="or" type="MathBase"/>
49                    <xs:element name="xor" type="MathBase"/>
50                    <xs:element name="not" type="MathBase"/>
51                    <xs:element name="sin" type="MathBase"/>
52                    <xs:element name="cos" type="MathBase"/>
53                    <xs:element name="tan" type="MathBase"/>
54                    <xs:element name="sec" type="MathBase"/>
55                    <xs:element name="csc" type="MathBase"/>
56                    <xs:element name="cot" type="MathBase"/>
57                    <xs:element name="sinh" type="MathBase"/>
58                    <xs:element name="cosh" type="MathBase"/>
59                    <xs:element name="tanh" type="MathBase"/>
60                    <xs:element name="sech" type="MathBase"/>
61                    <xs:element name="csch" type="MathBase"/>
62                    <xs:element name="coth" type="MathBase"/>
63                    <xs:element name="arcsin" type="MathBase"/>
64                    <xs:element name="arccos" type="MathBase"/>
65                    <xs:element name="arctan" type="MathBase"/>
66                    <xs:element name="arcsec" type="MathBase"/>
67                    <xs:element name="arccsc" type="MathBase"/>
68                    <xs:element name="arccot" type="MathBase"/>
69                    <xs:element name="arcsinh" type="MathBase"/>
70                    <xs:element name="arccosh" type="MathBase"/>
71                    <xs:element name="arctanh" type="MathBase"/>
72                    <xs:element name="arcsech" type="MathBase"/>
73                    <xs:element name="arccsch" type="MathBase"/>
74                    <xs:element name="arccoth" type="MathBase"/>
75                </xs:choice>
76                <xs:group ref="Node" maxOccurs="unbounded"/>
77            </xs:sequence>
78        </xs:extension>
79    </xs:complexContent>
80 </xs:complexType>

```

```

1 <xs:complexType name="Piece">
2   <xs:complexContent>
3     <xs:extension base="MathBase">
4       <xs:group ref="Node" minOccurs="2" maxOccurs="2"/>
5     </xs:extension>
6   </xs:complexContent>
7 </xs:complexType>
8 <xs:complexType name="Otherwise">
9   <xs:complexContent>
10     <xs:extension base="MathBase">
11       <xs:group ref="Node"/>
12     </xs:extension>
13   </xs:complexContent>
14 </xs:complexType>
15 <xs:complexType name="Piecewise">
16   <xs:complexContent>
17     <xs:extension base="MathBase">
18       <xs:sequence>
19         <xs:element name="piece" type="Piece" minOccurs="0" maxOccurs="unbounded"/>
20         <xs:element name="otherwise" type="Otherwise" minOccurs="0" maxOccurs="1"/>
21       </xs:sequence>
22     </xs:extension>
23   </xs:complexContent>
24 </xs:complexType>
25 <xs:attributeGroup name="AnnotationAttributes">
26   <xs:attributeGroup ref="MathAttributes"/>
27   <xs:attribute name="encoding" type="xs:string" use="required"/>
28 </xs:attributeGroup>
29 <xs:complexType name="Annotation">
30   <xs:simpleContent>
31     <xs:extension base="xs:string">
32       <xs:attributeGroup ref="AnnotationAttributes"/>
33     </xs:extension>
34   </xs:simpleContent>
35 </xs:complexType>
36 <xs:complexType name="Annotation-xml">
37   <xs:sequence maxOccurs="unbounded">
38     <xs:any processContents="skip"/>
39   </xs:sequence>
40   <xs:attributeGroup ref="AnnotationAttributes"/>
41 </xs:complexType>
42 <xs:complexType name="Semantics">
43   <xs:complexContent>
44     <xs:extension base="MathBase">
45       <xs:sequence>
46         <xs:group ref="Node"/>
47         <xs:sequence maxOccurs="unbounded">
48           <xs:choice>
49             <xs:element name="annotation" type="Annotation"/>
50             <xs:element name="annotation-xml" type="Annotation-xml"/>
51           </xs:choice>
52         </xs:sequence>
53       </xs:sequence>
54     </xs:extension>
55   </xs:complexContent>
56 </xs:complexType>
57 <xs:group name="Node">
58   <xs:choice>
59     <xs:element name="apply" type="Apply"/>
60     <xs:element name="cn" type="Cn"/>
61     <xs:element name="ci" type="Ci"/>
62     <xs:element name="csymbol" type="Csymbol"/>
63     <xs:element name="true" type="MathBase"/>
64     <xs:element name="false" type="MathBase"/>
65     <xs:element name="notanumber" type="MathBase"/>
66     <xs:element name="pi" type="MathBase"/>
67     <xs:element name="infinity" type="MathBase"/>
68     <xs:element name="exponentiale" type="MathBase"/>
69     <xs:element name="semantics" type="Semantics"/>
70     <xs:element name="piecewise" type="Piecewise"/>
71   </xs:choice>
72 </xs:group>
73 <xs:complexType name="Bvar">
74   <xs:complexContent>
75     <xs:extension base="MathBase">
76       <xs:sequence>
77         <xs:element name="ci" type="Ci"/>
78       </xs:sequence>
79     </xs:extension>
80   </xs:complexContent>

```

```

1      </xs:complexType>
2      <xs:complexType name="Lambda">
3          <xs:complexContent>
4              <xs:extension base="MathBase">
5                  <xs:sequence>
6                      <xs:element name="bvar" type="Bvar" minOccurs="0" maxOccurs="unbounded"/>
7                      <xs:group ref="Node"/>
8                  </xs:sequence>
9              </xs:extension>
10         </xs:complexContent>
11     </xs:complexType>
12     <xs:complexType name="Math">
13         <xs:complexContent>
14             <xs:extension base="MathBase">
15                 <xs:choice>
16                     <xs:group ref="Node"/>
17                     <xs:element name="lambda" type="Lambda"/>
18                 </xs:choice>
19             </xs:extension>
20         </xs:complexContent>
21     </xs:complexType>
22     <xs:element name="math" type="Math">
23         <!--This is the top-level element for a 'math' container in SBML.-->
24     </xs:element>
25 </xs:schema>

```

C Validation rules for SBML

This section contains a summary of all the conditions that should be true of a model, in addition to consistency with the XML Schemas given in Appendixes A and B, for that model to be considered valid SBML.

General XML validation

10101. An SBML XML file must use UTF-8 as the character encoding. More precisely, the **encoding** attribute of the XML declaration at the beginning of the XML data stream cannot have a value other than “UTF-8”. An example valid declaration is `<?xml version="1.0" encoding="UTF-8"?>`. (References: L2V2 Section 4.1; L2V3 Section 4.1.)
10102. An SBML XML document must not contain undefined elements or attributes in the SBML namespace. Documents containing unknown elements or attributes placed in the SBML namespace do not conform to the SBML Level 2 specification. (References: L2V2 Section 4.1; L2V3 Section 4.1.)
10103. An SBML XML document must conform to the XML Schema for the corresponding SBML Level, Version and Release. The XML Schema for SBML defines the basic SBML object structure, the data types used by those objects, and the order in which the objects may appear in an SBML document. (References: SBML L2V2 Section 4.1; L2V3 Section 4.1.)

General MathML validation

10201. All MathML content in SBML must appear within a **math** element, and the **math** element must be either explicitly or implicitly in the XML namespace “<http://www.w3.org/1998/Math/MathML>”. (References: L2V2 Section 3.5; L2V3 Section 3.4.)
10202. The only permitted MathML 2.0 elements in SBML Level 2 are the following: **cn**, **ci**, **csymbol**, **sep**, **apply**, **piecewise**, **piece**, **otherwise**, **eq**, **neq**, **gt**, **lt**, **geq**, **leq**, **plus**, **minus**, **times**, **divide**, **power**, **root**, **abs**, **exp**, **ln**, **log**, **floor**, **ceiling**, **factorial**, **and**, **or**, **xor**, **not**, **degree**, **bvar**, **logbase**, **sin**, **cos**, **tan**, **sec**, **csc**, **cot**, **sinh**, **cosh**, **tanh**, **sech**, **csch**, **coth**, **arcsin**, **arccos**, **arctan**, **arcsec**, **arccsc**, **arccot**, **arcsinh**, **arccosh**, **arctanh**, **arcsech**, **arccsch**, **arccoth**, **true**, **false**, **notanumber**, **pi**, **infinity**, **exponentiale**, **semantics**, **annotation**, and **annotation-xml**. (References: L2V2 Section 3.5.1; L2V3 Section 3.4.1.)
10203. In the SBML subset of MathML 2.0, the MathML attribute **encoding** is only permitted on **csymbol**. No other MathML elements may have the **encoding** attribute. (References: L2V2 Section 3.5.1; L2V3 Section 3.4.1.)
10204. In the SBML subset of MathML 2.0, the MathML attribute **definitionURL** is only permitted on **csymbol**. No other MathML elements may have a **definitionURL** attribute. (References: L2V2 Section 3.5.1; L2V3 Section 3.4.1.)
10205. In SBML Level 2, the only values permitted for the **definitionURL** attribute on a **csymbol** element are “<http://www.sbml.org/sbml/symbols/time>” and “<http://www.sbml.org/sbml/symbols/delay>”. (References: L2V2 Section 3.5.5; L2V3 Section 3.4.6.)
10206. In the SBML subset of MathML 2.0, the MathML attribute **type** is only permitted on the **cn** construct. No other MathML elements may have a **type** attribute. (References: L2V2 Section 3.5.1; L2V3 Section 3.4.1.)
10207. The only permitted values for the **type** attribute on MathML **cn** elements are “**e-notation**”, “**real**”, “**integer**”, and “**rational**”. (References: L2V2 Section 3.5.2; L2V3 Section 3.4.2.)
10208. MathML **lambda** elements are only permitted as the first element inside the **math** element of a **FunctionDefinition** or as the first element of a **semantics** element immediately inside the **math** element of a **FunctionDefinition**; they may not be used elsewhere in an SBML model. (References: L2V2 Sections 3.5.1 and 4.3; L2V3 Sections 3.4.1 and 4.3.2.)

- 1 10209. The arguments of the MathML logical operators **and**, **or**, **xor**, and **not** must have boolean values.
2 (References: L2V2 Section 3.5.8; L2V3 Section 3.4.9.)
- 3 10210. The arguments to the following MathML constructs must evaluate to be **numbers (i.e., MathML,
4 real, integer, rational, or “e-notation” numbers, or the time or delay csymbol)**: **plus**, **minus**,
5 **times**, **divide**, **power**, **root**, **abs**, **exp**, **ln**, **log**, **floor**, **ceiling**, **factorial**, **sin**, **cos**, **tan**, **sec**,
6 **csc**, **cot**, **sinh**, **cosh**, **tanh**, **sech**, **csch**, **coth**, **arcsin**, **arccos**, **arctan**, **arcsec**, **arccsc**, **arccot**,
7 **arcsinh**, **arccosh**, **arctanh**, **arcsech**, **arccsch**, **arccoth**. (References: L2V2 Section 3.5.8; L2V3
8 Section 3.4.9.)
- 9 10211. The values of all arguments to **eq** and **neq** operators should have the same type (either all boolean
10 or all numerical). (References: L2V2 Section 3.5.8; L2V3 Section 3.4.9.)
- 11 10212. The types of values within **piecewise** operators should all be consistent: the set of expressions that
12 make up the first arguments of the **piece** and **otherwise** operators within the same **piecewise** oper-
13 ator should all return values of the same type. (References: L2V2 Section 3.5.8; L2V3 Section 3.4.9.)
- 14 10213. The second argument of a MathML **piece** operator must have a boolean value. (References: L2V2
15 Section 3.5.8; L2V3 Section 3.4.9.)
- 16 10214. Outside of a **FunctionDefinition**, if a **ci** element is the first element within a MathML **apply**, then the
17 **ci**’s value can only be chosen from the set of identifiers of **FunctionDefinition** instances defined in the
18 enclosing **SBML Model** instance. (References: L2V2 Section 4.3.2; L2V3 Section 4.3.2.)
- 19 10215. Outside of a **FunctionDefinition**, if a **ci** element is not the first element within a MathML **apply**, then
20 the **ci**’s value can only be chosen from the set of identifiers of **Species**, **Compartment**, **Parameter** or
21 **Reaction** objects defined in the SBML model. (References: L2V2 Section 4.3.2; L2V3 Section 3.4.3.)
- 22 10216. The **id** value of a **Parameter** defined within a **KineticLaw** can only be used in **ci** elements within the
23 MathML content of that same **KineticLaw**; the identifier is not visible to other parts of the model.
24 (References: L2V2 Sections 3.4.1, 3.5.3 and 4.13.5; L2V3 Sections 3.3.1, 3.4.3 and 4.13.5.)
- 25 10217. The MathML formulas in the following elements must yield **numerical values (that is, MathML
26 real, integer or “e-notation” numbers, or the time or delay csymbol)**: **math** in **KineticLaw**,
27 **stoichiometryMath** in **SpeciesReference**, **math** in **InitialAssignment**, **math** in **AssignmentRule**, **math**
28 in **RateRule**, **math** in **AlgebraicRule**, **math** in **Event Delay**, and **math** in **EventAssignment**. (References:
29 L2V2 Section 4.10, 4.11, 4.13 and 4.14; L2V3 Sections 4.10, 4.11, 4.13 and 4.14.)
- 30 10218. A MathML operator must be supplied the number of arguments appropriate for that operator.
31 (References: L2V2 Section 3.5.1; L2V3 Section 3.4.1.)

32 General identifier validation

- 33 10301. The value of the **id** attribute on every instance of the following classes of objects in a model must
34 be unique across the set of all **id** values in a model: **Model**, **FunctionDefinition**, **CompartmentType**,
35 **SpeciesType**, **Compartment**, **Species**, **Reaction**, **SpeciesReference**, **ModifierSpeciesReference**, **Event**,
36 and model-wide **Parameters**. Exception: the identifiers of instances of **UnitDefinition**, and parameters
37 defined inside a reaction, are treated separately. (References: L2V1 Section 3.5; L2V2 Section 3.4.1;
38 L2V3 Section 3.3)
- 39 10302. The value of the **id** attribute of every **UnitDefinition** must be unique across the set of all **UnitDefinitions**
40 in the entire model. (References: L2V1 Section 3.5 and 4.4; L2V2 Sections 3.4.1 and 4.4; L2V3
41 Sections 3.3 and 4.4.)
- 42 10303. The value of the **id** attribute of each parameter defined locally within a **KineticLaw** must be unique
43 across the set of all such parameter definitions in that **KineticLaw**. (References: L2V1 Sections 3.4.1
44 and 4.13.9; L2V2 Sections 3.4.1 and 4.13.5; L2V3 Sections 3.3.1 and 4.13.5.)

- 1 10304. The value of the **variable attribute** in all **AssignmentRule** and **RateRule** definitions must be unique
2 across the set of all such rule definitions in a model. (References: L2V1 Section 4.8.4; L2V2 Sec-
3 tion 4.11; **L2V3 Section 4.11.3**.)
- 4 10305. In each **Event**, the value of the **variable attribute** within every **EventAssignment** definition must
5 be unique across the set of all **EventAssignments** within that **Event**. **In other words, a single Event**
6 **cannot make more than one assignment to a given identifier**. (References: L2V1 erratum 17; L2V2
7 Section 4.14.2; **L2V3 Section 4.14.4**.)
- 8 10306. An identifier used as the value of **variable** in an **EventAssignment** cannot also appear as the value
9 of **variable** in an **AssignmentRule**. (References: L2V1 Section 4.10.5; L2V2 Section 4.14.2; **L2V3**
10 **Section 4.14.4**.)
- 11 10307. Every **metaid attribute** value must be unique across the set of all **metaid** values in a model. (Refer-
12 ences: L2V2 Sections 3.1.6 and 3.3.1; **L2V3 Sections 3.2.1 and 3.1.6**.)
- 13 10308. The value of a **sboTerm** attribute must have the data type **SBOTerm**, which is a string consisting of
14 the characters ‘S’, ‘B’, ‘O’, ‘:’, followed by exactly seven digits. (References: L2V2 Section 3.1.9;
15 **L2V3 Section 3.1.9**.)
- 16 10309. The syntax of **metaid attribute** values must conform to the syntax of the XML type **ID**. (References:
17 L2V2 Sections 3.1.6 and 3.3.1; **L2V3 Sections 3.2.1 and 3.1.6**.)
- 18 10310. The syntax of **id attribute** values must conform to the syntax of the SBML type **SIID**. (References:
19 L2V2 Section 3.1.7; **L2V3 Section 3.1.7**.)
- 20 10311. **The syntax of unit identifiers (i.e., the values of the id attribute on UnitDefinition, the units attribute**
21 **on Compartment, the units attribute on Parameter, and the substanceUnits attribute on Species)**
22 **must conform to the syntax of the SBML type UnitSIID**. (References: L2V3 Section 3.1.8.)

23 **General annotation validation**

- 24 10401. Every top-level element within an **annotation** element must have a namespace declared. (References:
25 L2V2 Section 3.3.3; **L2V3 Section 3.2.4**.)
- 26 10402. There cannot be more than one top-level element using a given namespace inside a given **annotation**
27 element. (References: L2V2 Section 3.3.3; **L2V3 Section 3.2.4**.)
- 28 10403. Top-level elements within an **annotation** element cannot use any SBML namespace, whether ex-
29 plicitly (by declaring the namespace to be one of the URIs “<http://www.sbml.org/sbml/level1>”,
30 “<http://www.sbml.org/sbml/level2>”, “<http://www.sbml.org/sbml/level2/version2>”, or
31 “<http://www.sbml.org/sbml/level2/version3>”), or implicitly (by failing to declare any names-
32 pace). (References: L2V2 Section 3.3.3; **L2V3 Section 3.2.4**.)

33 **General unit validation**

- 34 10501. The units of the expressions used as arguments to a function call must match the units expected for
35 the arguments of that function. (References: L2V2 Section 3.5; **L2V3 Section 3.4**.)
- 36 10511. When the **variable** in an **AssignmentRule** refers to a **Compartment**, the units of the rule’s right-hand
37 side must be consistent with the units of that compartment’s size. (References: L2V2 Section 4.11.3;
38 **L2V3 Section 4.11.3**.)
- 39 10512. When the **variable** in an **AssignmentRule** refers to a **Species**, the units of the rule’s right-hand side
40 must be consistent with the units of the species’ quantity. (References: L2V2 Section 4.11.3; **L2V3**
41 **Section 4.11.3**.)

- 1 10513. When the **variable** in an **AssignmentRule** refers to a **Parameter**, the units of the rule's right-hand side
2 must be consistent with the units declared for that parameter. (References: L2V2 Section 4.11.3;
3 **L2V3 Section 4.11.3**.)
- 4 10521. When the **variable** in an **InitialAssignment** refers to a **Compartment**, the units of the **InitialAssign-**
5 **ment's math** expression must be consistent with the units of that compartment's size. (References:
6 L2V2 Section 4.10.4; **L2V3 Section 4.10**.)
- 7 10522. When the **variable** in an **InitialAssignment** refers to a **Species**, the units of the **InitialAssignment's**
8 **math** expression must be consistent with the units of that species' quantity. (References: L2V2
9 Section 4.10.4; **L2V3 Section 4.11.3**.)
- 10 10523. When the **variable** in an **InitialAssignment** refers to a **Parameter**, the units of the **InitialAssignment's**
11 **math** expression must be consistent with the units declared for that parameter. (References: L2V2
12 Section 4.10.4; **L2V3 Section 4.11.3**.)
- 13 10531. When the **variable** in a **RateRule** definition refers to a **Compartment**, the units of the rule's right-
14 hand side must be of the form *x per time*, where *x* is either the **units** in that **Compartment** definition,
15 or (in the absence of explicit units declared for the compartment size) the default units for that
16 compartment, and *time* refers to the units of time for the model. (References: L2V2 Section 4.11.4;
17 **L2V3 Section 4.11.4**.)
- 18 10532. When the **variable** in a **RateRule** definition refers to a **Species**, the units of the rule's right-hand
19 side must be of the form *x per time*, where *x* is the units of that species' quantity, and *time* refers to
20 the units of time for the model. (References: L2V2 Section 4.11.4; **L2V3 Section 4.11.4**.)
- 21 10533. When the **variable** in a **RateRule** definition refers to a **Parameter**, the units of the rule's right-hand
22 side must be of the form *x per time*, where *x* is the **units** in that **Parameter** definition, and *time*
23 refers to the units of time for the model. (References: L2V2 Section 4.11.4; **L2V3 Section 4.11.4**.)
- 24 10541. The units of the **math** formula in a **KineticLaw** definition must be the equivalent of *substance per time*.
25 (References: L2V2 Section 4.13.5; **L2V3 Section 4.13.5**.)
- 26 10551. The units of the mathematical formula in the **delay** of an **Event** must correspond to the model's
27 overall units of time. (**L2V3 Section 4.14.3**.)
- 28 10561. When the **variable** attribute of an **EventAssignment** contains the identifier of a **Compartment** in
29 the model, the units of the mathematical expression in the **EventAssignment's math** expression must
30 be consistent with the units of that compartment's size. (References: L2V2 Section 4.14.2; L2V3
31 Section **4.14.4**.)
- 32 10562. When the **variable** attribute of an **EventAssignment** contains the identifier of a **Species** in the model,
33 the units of the mathematical expression in the **EventAssignment's math** expression must be consistent
34 with the units of that species' quantity. (References: L2V2 Section 4.14.2; L2V3 Section **4.14.4**.)
- 35 10563. When the **variable** attribute of an **EventAssignment** contains the identifier of a **Parameter** in the
36 model, the units of the mathematical expression in the **EventAssignment's math** expression must
37 be consistent with the units declared for that parameter. (References: L2V2 Section 4.14.2; L2V3
38 Section **4.14.4**.)

39 **General model validation**

- 40 10601. A system of equations created from an SBML model must not be overdetermined. (References: L2V2
41 Section 4.11.5; **L2V3 Section 4.11.5**.)

General SBO validation

10701. The value of the **sboTerm** attribute on a **Model** must be an SBO identifier referring to a modeling framework defined in SBO. That is, the value must be a term derived from **SBO:0000004**, “modeling framework”. (References: L2V2 Section 4.2.1; L2V3 Section 4.2.2.)
10702. The value of the **sboTerm** attribute on a **FunctionDefinition** must be an SBO identifier referring to a mathematical expression. That is, the value must be a term derived from **SBO:0000064**, “mathematical expression”. (References: L2V2 Section 4.3.3; L2V3 Section 4.3.3.)
10703. The value of the **sboTerm** attribute on a **Parameter** must be an SBO identifier referring to a quantitative parameter. That is, the value must be a term derived from **SBO:0000002**, “quantitative parameter”. (References: L2V2 Section 4.9.5; L2V3 Section 4.9.5.)
10704. The value of the **sboTerm** attribute on an **InitialAssignment** must be an SBO identifier referring to a mathematical expression. That is, the value must be a term derived from **SBO:0000064**, “mathematical expression”. (References: L2V2 Section 4.10.3; L2V3 Section 4.10.3.)
10705. The value of the **sboTerm** attribute on a **AlgebraicRule**, **RateRule** or **AssignmentRule** must be an SBO identifier referring to a mathematical expression. That is, the value must be a term derived from **SBO:0000064**, “mathematical expression”. (References: L2V2 Section 4.11.1; L2V3 Section 4.11.1.)
10706. The value of the **sboTerm** attribute on a **Constraint** must be an SBO identifier referring to a mathematical expression. That is, the value must be a term derived from **SBO:0000064**, “mathematical expression”. (References: L2V2 Section 4.12.3; L2V3 Section 4.12.3.)
10707. The value of the **sboTerm** attribute on a **Reaction** must be an SBO identifier referring to a modeling framework. That is, the value must be a term derived from **SBO:0000231**, “event”. (References: L2V2 Section 4.13.1; L2V3 Section 4.13.1.)
10708. The value of the **sboTerm** attribute on a **SpeciesReference** or a **ModifierSpeciesReference** must be an SBO identifier referring to a participant functional role. That is, the value must be a term derived from **SBO:0000003**, “participant functional type”. The appropriate term depends on whether the entity is a reactant, product or modifier. If a reactant, the term should be chosen from the **SBO:0000010**, “reactant” hierarchy; if a product, the term should be chosen from the **SBO:0000011**, “product” hierarchy; and if a modifier, the term should be chosen from the **SBO:0000019**, “modifier” hierarchy. (References: L2V2 Section 4.13.2; L2V3 Sections 4.13.1 and 5.)
10709. The value of the **sboTerm** attribute on a **KineticLaw** must be an SBO identifier referring to a rate law. That is, the value must be a term derived from **SBO:0000001**, “rate law”. (References: L2V2 Section 4.13.5; L2V3 Section 4.13.1.)
10710. The value of the **sboTerm** attribute on an **Event** must be an SBO identifier referring to a mathematical expression. That is, the value must be a term derived from **SBO:0000231**, “event”. (References: L2V2 Section 4.14.1; L2V3 Section 4.14.1.)
10711. The value of the **sboTerm** attribute on an **EventAssignment** must be an SBO identifier referring to a mathematical expression. That is, the value must be a term derived from **SBO:0000064**, “mathematical expression”. (References: L2V2 Section 4.14.2; L2V3 Section 4.14.1.)
10712. The value of the **sboTerm** attribute on a **Compartment** must be an SBO identifier referring to a participant physical type. That is, the value must be a term derived from **SBO:0000236**, “participant physical type”. (References: L2V3 Section 5.)
10713. The value of the **sboTerm** attribute on a **Species** must be an SBO identifier referring to a participant physical type. That is, the value must be a term derived from **SBO:0000236**, “participant physical type”. (References: L2V3 Section 5.)

- 1 10714. The value of the **sboTerm** attribute on a **CompartmentType** must be an SBO identifier referring to a
2 participant physical type. That is, the value must be a term derived from SBO:0000236, “participant
3 physical type”. (References: L2V3 Section 5.)
- 4 10715. The value of the **sboTerm** attribute on a **SpeciesType** must be an SBO identifier referring to a
5 participant physical type. That is, the value must be a term derived from SBO:0000236, “participant
6 physical type”. (References: L2V3 Section 5.)
- 7 10716. The value of the **sboTerm** attribute on a **Trigger** must be an SBO identifier referring to a mathematical
8 expression. That is, the value must be a term derived from SBO:0000064, “mathematical expression”.
9 (References: L2V3 Section 5.)
- 10 10717. The value of the **sboTerm** attribute on a **Delay** must be an SBO identifier referring to a mathematical
11 expression. That is, the value must be a term derived from SBO:0000064, “mathematical expression”.
12 (References: L2V3 Section 5.)

13 General notes validation

- 14 10801. The contents of the **notes** element must be explicitly placed in the XHTML XML namespace.
15 (References: L2V3 Section 3.2.3.)
- 16 10802. The contents of the **notes** element must not contain an XML declaration (i.e., a string of the form
17 “<?xml version="1.0" encoding="UTF-8"?>” or similar). (References: L2V2 Section 3.3.2; L2V3
18 Section 3.2.3.)
- 19 10803. The contents of the **notes** element must not contain an XML DOCTYPE declaration (i.e., a string
20 beginning with the characters “<!DOCTYPE”). (References: L2V2 Section 3.3.2; L2V3 Section 3.2.3.)
- 21 10804. The XHTML content inside a **notes** element can only take one of the following general forms: (1)
22 a complete XHTML document beginning with the element <html> and ending with </html>; (2)
23 the “body” portion of a document beginning with the element <body> and ending with </body>; or
24 (3) XHTML content that is permitted within a <body> ... </body> elements. (References: L2V2
25 Section 3.3.2; L2V3 Section 3.2.3.)

26 SBML container validation

- 27 20101. The **sbml** container element must declare the XML Namespace for SBML, and this declaration must
28 be consistent with the values of the **level** and **version** attributes on the **sbml** element. (References:
29 L2V2 Section 4.1; L2V3 Section 4.1.)
- 30 20102. The **sbml** container element must declare the SBML Level using the attribute **level**, and this dec-
31 laration must be consistent with the XML Namespace declared for the **sbml** element. (References:
32 L2V2 Section 4.1; L2V3 Section 4.1.)
- 33 20103. The **sbml** container element must declare the SBML Version using the attribute **version**, and this
34 declaration must be consistent with the XML Namespace declared for the **sbml** element. (References:
35 L2V2 Section 4.1; L2V3 Section 4.1.)

36 Model validation

- 37 20201. An SBML document must contain a **Model** definition. (References: L2V1, L2V2 and L2V3 Sec-
38 tion 4.1).
- 39 20202. The order of subelements within a **Model** object instance must be the following, with each ele-
40 ment optional: **listOfFunctionDefinitions**, **listOfUnitDefinitions**, **listOfCompartmentTypes**,
41 **listOfSpeciesTypes**, **listOfCompartments**, **listOfSpecies**, **listOfParameters**,
42 **listOfInitialAssignments**, **listOfRules**, **listOfConstraints**, **listOfReactions**, **listOfEvents**.
43 (References: L2V2 Section 4.2; L2V3 Section 4.2.)

- 1 20203. The `listOf_` container elements in a **Model** instance are optional, but if present, the elements must
 2 not be empty. Specifically, if any of the following are present in a **Model**, they must not be empty:
 3 `listOfFunctionDefinitions`, `listOfUnitDefinitions`, `listOfCompartmentTypes`, `listOfSpecies-`
 4 `Types`, `listOfCompartments`, `listOfSpecies`, `listOfParameters`, `listOfInitialAssignments`,
 5 `listOfRules`, `listOfConstraints`, `listOfReactions` and `listOfEvents`. (References: This is re-
 6 quirement stemming from the XML Schema used for SBML; **L2V3 Section 4.2**.)
- 7 20204. If a model defines any **Species**, then the model must also define at least one **Compartment**. This is
 8 an implication of the fact that the `compartment` attribute on **Species** is not optional. (References:
 9 L2V1 Section 4.5; L2V2 Section 4.8.3; **L2V3 Section 4.8.3**.)

10 FunctionDefinition validation

- 11 20301. The top-level element within `math` in a **FunctionDefinition** must be **one and only one MathML lambda**
 12 element or a MathML `semantics` element containing one and only one `lambda` element. (References:
 13 L2V1 Section 4.3.2; L2V2 Section 4.3.2; **L2V3 Section 4.3.2**.)
- 14 20302. Inside the `lambda` of a **FunctionDefinition**, if a `ci` element is the first element within a MathML `apply`,
 15 then the `ci`'s value can only be chosen from the set of identifiers of other SBML **FunctionDefinitions**
 16 defined prior to that point in the SBML model. In other words, forward references to user-defined
 17 functions are not permitted. (References: L2V2 Section 3.5.3; **L2V3 Section 3.4.3**.)
- 18 20303. Inside the `lambda` of a **FunctionDefinition**, the identifier of that **FunctionDefinition** cannot appear as
 19 the value of a `ci` element. SBML functions are not permitted to be recursive. (References: L2V2
 20 L2V2 Section 3.5.3 and 4.3.2; **L2V3 Sections 3.4.3 and 4.3.2**.)
- 21 20304. Inside the `lambda` of a **FunctionDefinition**, if a `ci` element is not the first element within a MathML
 22 `apply`, then the `ci`'s value can only be **an identifier provided as the** value of a `bvar` element declared
 23 in that `lambda`. **This restriction also applies to the `csymbol` for time.** In other words, all model
 24 **quantities and variables** referenced inside a function definition must be passed as arguments to that
 25 function. (References: L2V2 L2V2 Section 3.5.3 and 4.3.2; **L2V3 Sections 3.4.3 and 4.3.2**.)
- 26 20305. The value type returned by a **FunctionDefinition**'s `lambda` must be either boolean or **numerical**. (Ref-
 27 erences: L2V2 Section 3.5.8; **L2V3 Section 3.4.9**.)

28 Unit and UnitDefinition validation

- 29 20401. The value of the `id` attribute in a **UnitDefinition** must be **of type UnitSId** and not be identical to
 30 any unit predefined in SBML. **That is, the identifier must not be the same as any of the following**
 31 predefined units: “ampere”, “gram”, “katal”, “metre”, “second”, “watt”, “becquerel”, “gray”,
 32 “kelvin”, “mole”, “siemens”, “weber”, “candela”, “henry”, “kilogram”, “newton”, “sievert”,
 33 “coulomb”, “hertz”, “litre”, “ohm”, “steradian”, “dimensionless”, “item”, “lumen”, “pascal”,
 34 “tesla”, “farad”, “joule”, “lux”, “radian”, or “volt”. (References: L2V1 erratum 14; L2V2
 35 Section 4.4.2; **L2V3 Section 4.4.2**.)
- 36 20402. Redefinitions of the built-in unit **substance** must be based on the unit `mole`, `item`, `gram`, `kilogram`,
 37 or `dimensionless`. More formally, a **UnitDefinition** for **substance** must simplify to a single **Unit** in
 38 which either (a) the `kind` attribute has a value of “`mole`”, “`item`”, “`gram`”, or “`kilogram`”, and the
 39 `exponent` attribute has a value of “1”, or (b) the `kind` attribute has a value of “`dimensionless`”
 40 with any `exponent` attribute value. (References: L2V1 Section 4.4.3; L2V2 Section 4.4.3; **L2V3**
 41 **Section 4.4.3**.)
- 42 20403. Redefinitions of the built-in unit **length** must be based on the unit `metre` or `dimensionless`. More
 43 formally, a **UnitDefinition** for **length** must simplify to a single **Unit** in which either (a) the `kind`
 44 `attribute` has a value of “`metre`” and the `exponent` attribute has a value of “1”, or (b) the `kind`
 45 `attribute` has a value of “`dimensionless`” with any `exponent` value. (References: L2V1 Section
 46 4.4.3; L2V2 Section 4.4.3; **L2V3 Section 4.4.3**.)

- 1 20404. Redefinitions of the built-in unit **area** must be based on squared **metres** or **dimensionless**. More
2 formally, a **UnitDefinition** for **area** must simplify to a single **Unit** in which either (a) the **kind attribute**
3 has a value of “**metre**” and the **exponent attribute** has a value of “**2**”, or (b) the **kind attribute**
4 has a value of “**dimensionless**” with any **exponent** value. (References: L2V1 Section 4.4.3; L2V2
5 Section 4.4.3; **L2V3 Section 4.4.3**.)
- 6 20405. Redefinitions of the built-in unit **time** must be based on **the unit second**. More formally, a **UnitDef-**
7 **inition** for **time** must simplify to a single **Unit** in which either (a) the **kind attribute** has a value of
8 “**second**” and the **exponent attribute** has a value of “**1**”, or (b) the **kind attribute** has a value of
9 “**dimensionless**” with any **exponent** value. (References: L2V1 Section 4.4.3; L2V2 Section 4.4.3;
10 **L2V3 Section 4.4.3**.)
- 11 20406. Redefinitions of the built-in unit **volume** must be based on **litre**, **metre** or **dimensionless**. More
12 formally, a **UnitDefinition** for **volume** must simplify to a single **Unit** in which the **kind attribute** value
13 is either “**litre**”, “**metre**”, or “**dimensionless**”. Additional constraints apply if the kind is “**litre**”
14 or “**metre**”. (References: L2V1 Section 4.4.3; L2V2 Section 4.4.3; **L2V3 Section 4.4.3**.)
- 15 20407. If a **UnitDefinition** for **volume** simplifies to a **Unit** in which the **kind attribute** value is “**litre**”, then its
16 **exponent attribute** value must be “**1**”. (References: L2V1 Section 4.4.3; L2V2 Section 4.4.3; **L2V3**
17 **Section 4.4.3**.)
- 18 20408. If a **UnitDefinition** for **volume** simplifies to a **Unit** in which the **kind attribute** value is “**metre**”, then its
19 **exponent attribute** value must be “**3**”. (References: L2V1 Section 4.4.3; L2V2 Section 4.4.3; **L2V3**
20 **Section 4.4.3**.)
- 21 20409. The **listOfUnits** container in a **UnitDefinition** cannot be empty. (References: L2V2 Section 4.4;
22 **L2V3 Section 4.4**.)
- 23 20410. The value of the **kind attribute** of a **Unit** can only be one of the **predefined base units** in SBML; that
24 is, the SBML unit system is not hierarchical and user-defined units cannot be defined using other
25 user-defined units. (References: L2V2 Section 4.4.1; **L2V3 Section 4.4.2**.)
- 26 20411. (Rule does not apply in SBML Level 2 Version 3.)
- 27 20412. (Rule does not apply in SBML Level 2 Version 3.)

28 **Compartment validation**

- 29 20501. The size of a **Compartment** must not be set if the compartment’s **spatialDimensions attribute** has
30 value **0**. (References: L2V1 Section 4.5.3; L2V2 Section 4.7.4; **L2V3 Section 4.7.5**.)
- 31 20502. If a **Compartment** definition has a **spatialDimensions** value of “**0**”, then its **units attribute** must not
32 be set. If the compartment has no dimensions, then no units can be associated with a non-existent
33 size. (References: L2V1 Section 4.5.4; L2V2 Section 4.7.5; **L2V3 Section 4.7.5**.)
- 34 20503. If a **Compartment** definition has a **spatialDimensions** value of “**0**”, then its **constant attribute** value
35 must either default to or be set to “**true**”. If the compartment has no dimensions, then its size can
36 never change. (References: L2V1 Section 4.5.5; L2V2 Section 4.7.4; **L2V3 Section 4.7.6**.)
- 37 20504. The **outside attribute** value of a **Compartment** must be the identifier of another **Compartment** defined
38 in the model. (References: L2V1 Section 4.5.6; L2V2 Section 4.7.7; **L2V3 Section 4.7.7**.)
- 39 20505. A **Compartment** may not enclose itself through a chain of references involving the **outside attribute**.
40 This means that a compartment cannot have its own identifier as the value of **outside**, nor can it point
41 to another compartment whose **outside attribute** points directly or indirectly to the compartment.
42 (References: L2V1 erratum 11; L2V2 Section 4.7.7; **L2V3 Section 4.7.7**.)

- 1 20506. The **outside** attribute value of a given **Compartment** instance cannot be a compartment having a
2 **spatialDimensions** value of “0”, unless both compartments have **spatialDimensions**=“0”. Simply
3 put, a zero-dimensional compartment cannot enclose compartments that have anything other than
4 zero dimensions themselves. (References: L2V2 Section 4.7.7; L2V3 Section 4.7.7.)
- 5 20507. The value of the **units** attribute on a **Compartment** having **spatialDimensions** of “1” must be either
6 “length”, “metre”, “dimensionless”, or the identifier of a **UnitDefinition** based on either **metre** (with
7 **exponent** equal to “1”) or **dimensionless**. (References: L2V1 Section 4.5.4; L2V2 Section 4.7.5;
8 L2V3 Section 4.7.5.)
- 9 20508. The value of the **units** attribute on a **Compartment** having **spatialDimensions** of “2” must be
10 either “area”, “dimensionless”, or the identifier of a **UnitDefinition** based on either **metre** (with
11 **exponent** equal to “2”) or **dimensionless**. (References: L2V1 Section 4.5.4; L2V2 Section 4.7.5;
12 L2V3 Section 4.7.5.)
- 13 20509. The value of the **units** attribute on a **Compartment** having **spatialDimensions** of “3” must be
14 either “volume”, “litre”, or the identifier of a **UnitDefinition** based on either **litre**, **metre** (with
15 **exponent** equal to “3”), or **dimensionless**. (References: L2V1 Section 4.5.4; L2V2 Section 4.7.5;
16 L2V3 Section 4.7.5.)
- 17 20510. If the **compartmentType** attribute is given a value in a **Compartment** definition, it must contain the
18 identifier of an existing **CompartmentType** defined in the model. (References: L2V2 Section 4.7.2;
19 L2V3 Section 4.7.2.)

20 Species validation

- 21 20601. The value of **compartment** in a **Species** definition must be the identifier of an existing **Compartment**
22 defined in the model. (References: L2V1 Section 4.6.2; L2V2 Section 4.8.3; L2V3 Section 4.8.3.)
- 23 20602. (Rule does not apply in SBML Level 2 Version 3.)
- 24 20603. (Rule does not apply in SBML Level 2 Version 3.)
- 25 20604. If a **Species** located in a **Compartment** whose **spatialDimensions** is set to “0”, then that **Species**
26 definition cannot set **initialConcentration**. (References: L2V1 Section 4.6.3; L2V2 Section 4.8.4;
27 L2V3 Section 4.8.4.)
- 28 20605. (Rule does not apply in SBML Level 2 Version 3.)
- 29 20606. (Rule does not apply in SBML Level 2 Version 3.)
- 30 20607. (Rule does not apply in SBML Level 2 Version 3.)
- 31 20608. The value of a **Species**’s **substanceUnits** attribute can only be one of the following: “substance”,
32 “mole”, “item”, “gram”, “kilogram”, “dimensionless”, or the identifier of a **UnitDefinition** de-
33 rived from “mole” (with an **exponent** of “1”), “item” (with an **exponent** of “1”), “gram” (with
34 an **exponent** of “1”), “kilogram” (with an **exponent** of “1”), or “dimensionless”. (References:
35 L2V1 Section 4.6.4; L2V2 Section 4.8.5; L2V3 Section 4.8.5.)
- 36 20609. A **Species** cannot set values for both **initialConcentration** and **initialAmount** because they are
37 mutually exclusive. (References: L2V1 Section 4.6.3; L2V2 Section 4.8.4; L2V3 Section 4.8.4.)
- 38 20610. A **Species**’ quantity cannot be determined simultaneously by both reactions and rules. More formally,
39 if the identifier of a **Species** definition having **boundaryCondition**=“false” and **constant**=“false”
40 is referenced by a **SpeciesReference** anywhere in a model, then this identifier cannot also appear as
41 the value of a **variable** in an **AssignmentRule** or a **RateRule**. (References: L2V1 Section 4.6.5; L2V2
42 Section 4.8.6; L2V3 Section 4.8.6.)

- 1 20611. A **Species** having `boundaryCondition="false"` cannot appear as a reactant or product in any reaction if that **Species** also has `constant="true"`. (References: L2V1 Section 4.6.5; L2V2 Section 4.8.6; L2V3 Section 4.8.6.)
- 2
- 3
- 4 20612. The value of `speciesType` in a **Species** definition must be the identifier of an existing **SpeciesType**. (References: L2V2 Section 4.8.2; L2V3 Section 4.8.2.)
- 5
- 6 20613. There cannot be more than one species of a given **SpeciesType** in the same compartment of a model. More formally, for any given compartment, there cannot be more than one **Species** definition in which both of the following hold simultaneously: (i) the **Species**' `compartment` value is set to that compartment's identifier and (ii) the **Species**' `speciesType` is set the same value as the `speciesType` of another **Species** that also sets its `compartment` to that compartment identifier. (References: L2V2 Section 4.8.2; L2V3 Section 4.8.2.)
- 7
- 8
- 9
- 10
- 11
- 12 20614. The `compartment` attribute in a **Species** is mandatory. A species definition in a model must include a value for this attribute. (References: L2V2 Section 4.8.3; L2V3 Section 4.8.3.)
- 13

14 Parameter validation

- 15 20701. The `units` in a **Parameter** definition must be a value chosen from among the following: a built-in unit (e.g., "substance", "time", etc.), the identifier of a **UnitDefinition** in the model, or one of the predefined base units in SBML (e.g., "litre", "mole", "metre", etc.). (References: L2V1 Section 4.7.3; L2V2 Section 4.9.3; L2V3 Section 4.9.3.)
- 16
- 17
- 18

19 InitialAssignment validation

- 20 20801. The value of `symbol` in an **InitialAssignment** definition must be the identifier of an existing **Compartment**, **Species**, or **Parameter** defined in the model. (References: L2V2 Section 4.10.1; L2V3 Section 4.10.)
- 21
- 22
- 23 20802. A given identifier cannot appear as the value of more than one **InitialAssignment**'s `symbol` attribute across the set of all **InitialAssignments** in a model. (References: L2V2 Section 4.10; L2V3 Section 4.10.)
- 24
- 25 20803. The identifier given as the value of a `symbol` attribute in any **InitialAssignment** definition cannot also appear as the value of the `variable` attribute in an **AssignmentRule**. In other words, a model cannot simultaneously define both an initial assignment and an assignment rule for the same species, compartment or parameter in a model. (References: L2V2 Section 4.10.4; L2V3 Section 4.10.)
- 26
- 27
- 28

29 AssignmentRule and RateRule validation

- 30 20901. The value of an **AssignmentRule**'s `variable` must be the identifier of an existing **Compartment**, **Species**, or globally-defined **Parameter**. (References: L2V1 Section 4.8.2; L2V2 Section 4.11.3 L2V3 Section 4.11.3.)
- 31
- 32
- 33 20902. The value of a **RateRule**'s `variable` must be the identifier of an existing **Compartment**, **Species**, or globally-defined **Parameter**. (References: L2V1 Section 4.8.3; L2V2 Section 4.11.4; L2V3 Section 4.11.4.)
- 34
- 35
- 36 20903. Any **Compartment**, **Species** or **Parameter** whose identifier is the value of a `variable` attribute in an **AssignmentRule**, must have a value of "false" for `constant`. (References: L2V1 Section 4.8.4; L2V2 Section 4.11.3; L2V3 Section 4.11.3.)
- 37
- 38
- 39 20904. Any **Compartment**, **Species** or **Parameter** whose identifier is the value of a `variable` attribute in a **RateRule**, must have a value of "false" for `constant`. (References: L2V1 Section 4.8.4; L2V2 Section 4.11.4; L2V3 Section 4.11.4.)
- 40
- 41
- 42 20905. (Rule removed because it was effectively a duplicate of 10304.)

1 20906. There must not be circular dependencies in the combined set of **InitialAssignment**, **AssignmentRule**
2 and **KineticLaw** definitions in a model. Each of these constructs has the effect of assigning a value
3 to an identifier (i.e., the identifier given in the **attribute** symbol in **InitialAssignment**, the **attribute**
4 **variable** in **AssignmentRule**, and the **attribute id** on the **KineticLaw**'s enclosing **Reaction**). Each of
5 these constructs computes the value using a mathematical formula. The formula for a given identifier
6 cannot make reference to a second identifier whose own definition depends directly or indirectly on
7 the first identifier. (References: L2V2 Section 4.11.5; L2V3 Section 4.11.5.)

8 Constraint validation

9 21001. A **Constraint** **math** expression must evaluate to a value of type **boolean**. (References: L2V2 Sec-
10 tion 4.12.1; L2V3 Section 4.12.)

11 21002. The order of subelements within **Constraint** must be the following: **math**, **message**. The **message**
12 element is optional, but if present, must follow the **math** element. (References: L2V2 Section 4.12;
13 L2V3 Section 4.12.)

14 21003. The contents of the **message** element in a **Constraint** must be explicitly placed in the XHTML XML
15 namespace. (References: L2V3 Section 4.12.2.)

16 21004. The contents of the **message** element must not contain an XML declaration (i.e., a string of the form
17 “<?xml version="1.0" encoding="UTF-8"?>” or similar). (References: L2V3 Section 4.12.2.)

18 21005. The contents of the **message** element must not contain an XML DOCTYPE declaration (i.e., a string
19 beginning with the characters “<!DOCTYPE”). (References: L2V3 Section 4.12.2.)

20 21006. The XHTML content inside a **Constraint**'s **message** element can only take one of the following general
21 forms: (1) a complete XHTML document beginning with the element **<html>** and ending with
22 **</xhtml>**; (2) the “body” portion of a document beginning with the element **<body>** and ending
23 with **</body>**; or (3) XHTML content that is permitted within a **<body> ... </body>** elements.
24 (References: L2V3 Section 4.12.2.)

25 Reaction validation

26 21101. A **Reaction** definition must contain at least one **SpeciesReference**, either in its **listOfReactants**
27 or its **listOfProducts**. A reaction without any reactant or product species is not permitted, re-
28 gardless of whether the reaction has any modifier species. (References: L2V2 Section 4.13.1; L2V3
29 Section 4.13.3.)

30 21102. The order of subelements within **Reaction** must be the following (where every one is optional):
31 **listOfReactants**, **listOfProducts**, **listOfModifiers**, **kineticLaw**. (References: L2V2 Section 4.13;
32 L2V3 Section 4.13.)

33 21103. The following containers are all optional in a **Reaction**, but if any present is, it must not be empty:
34 **listOfReactants**, **listOfProducts**, **listOfModifiers**, **kineticLaw**. (References: L2V2 Section 4.13;
35 L2V3 Section 4.13.)

36 21104. The list of reactants (**listOfReactants**) and list of products (**listOfProducts**) in a **Reaction** can
37 only contain **speciesReference** elements. (References: L2V1 Section 4.9; L2V2 Section 4.13; L2V3
38 Section 4.13.)

39 21105. The list of modifiers (**listOfModifiers**) in a **Reaction** can only contain **modifierSpeciesReference**
40 elements. (References: L2V1 Section 4.9; L2V2 Section 4.13; L2V3 Section 4.13.)

41 SpeciesReference and ModifierSpeciesReference validation

42 21111. The value of a **SpeciesReference** **species** **attribute** must be the identifier of an existing **Species** in
43 the model. (References: L2V1 Section 4.9.5; L2V2 Section 4.13.2; L2V3 Section 4.13.3.)

1 21112. (Rule removed because it was effectively a duplicate of 20611.)

2 21113. A **SpeciesReference** must not have a value for both **stoichiometry** and **stoichiometryMath**; they
3 are mutually exclusive. (References: L2V1 Section 4.9.5; L2V2 Section 4.13.3; L2V3 Section 4.13.3.)

4 **KineticLaw validation**

5 21121. All species referenced in the **KineticLaw** formula of a given reaction must first be declared using
6 **SpeciesReference** or **ModifierSpeciesReference**. More formally, if a **Species** identifier appears in a
7 **ci** element of a **Reaction**'s **KineticLaw** formula, that same identifier must also appear in at least
8 one **SpeciesReference** or **ModifierSpeciesReference** in the **Reaction** definition. (References: L2V2
9 Section 4.13.5; L2V3 Section 4.13.5.)

10 21122. The order of subelements within **KineticLaw** must be the following: **math**, **listOfParameters**. The
11 **listOfParameters** is optional, but if present, must follow **math**. (References: L2V2 Section 4.13.5;
12 L2V3 Section 4.13.5.)

13 21123. If present, the **listOfParameters** in a **KineticLaw** must not be an empty list. (References: L2V2
14 Section 4.13.5; L2V3 Section 4.13.)

15 21124. The **constant** attribute on a **Parameter** local to a **KineticLaw** cannot have a value other than "true".
16 The values of parameters local to **KineticLaw** definitions cannot be changed, and therefore they are
17 always constant. (References: L2V2 Section 4.9.4; L2V3 Section 4.9.4.)

18 21125. (Rule does not apply in SBML Level 2 Version 3.)

19 21126. (Rule does not apply in SBML Level 2 Version 3.)

20 **StoichiometryMath validation**

21 21131. All species referenced in the **StoichiometryMath** formula of a given reaction must first be declared
22 using **SpeciesReference** or **ModifierSpeciesReference**. More formally, if a **Species** identifier appears
23 in a **ci** element of a **Reaction**'s **StoichiometryMath** formula, that same identifier must also appear in
24 at least one **SpeciesReference** or **ModifierSpeciesReference** in the **Reaction** definition. (References:
25 L2V2 Sections 4.13.2 and 4.13.4; L2V3 Sections 4.13.2 and 4.13.4.)

26 **Event validation**

27 21201. An **Event** object must have a **trigger**. (References: L2V1 Section 4.10.2; L2V2 Section 4.14.1; L2V3
28 Section 4.14.2.)

29 21202. An **Event** **trigger** expression must evaluate to a value of type **boolean**. (References: L2V1 Section
30 4.10.2; L2V2 Section 4.14.1; L2V3 Section 4.14.2.)

31 21203. An **Event** object must have at least one **EventAssignment** object in its **listOfEventAssignments**.
32 (References: L2V1 Section 4.10.5; L2V2 Section 4.14; L2V3 Section 4.14.)

33 21204. (Rule does not apply in SBML Level 2 Version 3.)

34 21205. The order of subelements within an **Event** object instance must be the following: **trigger**, **delay**,
35 **listOfEventAssignments**. The **delay** element is optional, but if present, must follow **trigger**.
36 (References: L2V2 Section 4.14; L2V3 Section 4.14.)

37 **EventAssignment validation**

38 21211. The value of **variable** in an **EventAssignment** can only be the identifier of a **Compartment**, **Species**,
39 or model-wide **Parameter** definition. (References: L2V1 Section 4.10.5; L2V2 Section 4.14.2; L2V3
40 Section 4.14.4.)

1 21212. Any **Compartment**, **Species** or **Parameter** definition whose identifier is used as the value of **variable**
2 in an **EventAssignment** must have a value of “false” for its **constant attribute**. (References: L2V1
3 Section 4.10.5; L2V2 Section 4.14.2; L2V3 Section 4.14.4.)

D A method for assessing whether an SBML model is overdetermined

As explained in Section 4.11.5, an SBML model must not be overdetermined. It is possible to use purely static analysis to assess this condition for the system of equations implied by a model, by constructing a bipartite graph of the model's variables and equations and then searching for a maximal matching (Chartrand, 1977). A efficient algorithm for finding a maximal matching is described by Hopcroft and Karp (1973). In this appendix, we provide a concrete application to SBML of the general approach described in Section 4.11.5. The approach is defined in terms of the ordinary differential equations (ODEs) implied by an SBML model; despite our use of a differential equation framework for this explanation, it should be understood that this use of ODEs has no implication about the framework actually used to simulate the model.

Definition of the method

First, we treat both assignment rules and formulas in KineticLaw objects as SBML-style algebraic equations. (Both are essentially assignment statements, and assignment statements can be trivially converted to algebraic equations with zero on the left-hand side.) We also assume that an ODE is constructed for each species determined by one or more Reaction's KineticLaw math expressions. Finally, we assume that the model has already been determined to be valid in all other respects (e.g., there are no undefined variables in the equations), and what remains is to evaluate whether it is overdetermined.

We construct the bipartite graph for a given SBML model as follows:

1. For each of the following in the model, create one vertex representing an equation:
 - Every Species object having boundaryCondition="false", constant="false", and which is referenced as a reactant or product in one or more Reaction objects containing KineticLaw objects
 - Every AssignmentRule object
 - Every RateRule object
 - Every AlgebraicRule object
 - Every KineticLaw object
2. For each of the following in the model, create one vertex representing a variable:
 - Every Species object having constant="false"
 - Every Compartment object having constant="false"
 - Every global Parameter having constant="false"
 - Every Reaction object
3. For each of the following, create one edge:
 - Every vertex created in step 2(a) to that species' equation vertex created in step 1(a)
 - Every vertex created in step 1(b) to the particular vertex created in steps 2(a)–2(d) that represents the variable referenced by the variable attribute of the rule
 - Every vertex created in step 1(c) to the particular vertex created in steps 2(a)–2(d) that represents the variable referenced by the variable attribute of the rule
 - Every vertex created in step 1(e) to the particular vertex created in step 2(d) that is the Reaction object containing that particular KineticLaw object
 - Every vertex created in steps 2(a)–2(d) representing an identifier appearing as the content of a MathML ci element within an expression of an AssignmentRule, to the vertex for that particular AssignmentRule created in step 1(b)
 - Every vertex created in steps 2(a)–2(d) representing an identifier appearing as the content of a MathML ci element within an expression of an AlgebraicRule, to the vertex for that particular AlgebraicRule created in step 1(d)
 - Every vertex created in steps 2(a)–2(c) representing an identifier appearing as the content of a MathML ci element within an expression of a KineticLaw, to the vertex for that particular KineticLaw created in step 1(e)

Example application of the method

What follows is an example of applying the method above to the SBML model shown below:

```
<?xml version="1.0" encoding="UTF-8"?>
<sbml xmlns="http://www.sbml.org/sbml/level2/version3" level="2" version="3">
  <model id="example">
    <listOfUnitDefinitions>
      <unitDefinition id="per_second">
        <listOfUnits>
          <unit kind="second" exponent="-1"/>
        </listOfUnits>
      </unitDefinition>
    </listOfUnitDefinitions>
    <listOfCompartments>
      <compartment id="C" size="1"/>
    </listOfCompartments>
    <listOfSpecies>
      <species id="S1" compartment="C" initialConcentration="1"/>
      <species id="S2" compartment="C" initialConcentration="0"/>
    </listOfSpecies>
    <listOfRules>
      <algebraicRule>
        <math xmlns="http://www.w3.org/1998/Math/MathML">
          <apply>
            <plus/>
            <ci> S1 </ci>
            <ci> S2 </ci>
          </apply>
        </math>
      </algebraicRule>
    </listOfRules>
    <listOfReactions>
      <reaction id="R">
        <listOfReactants>
          <speciesReference species="S1"/>
        </listOfReactants>
        <listOfProducts>
          <speciesReference species="S2"/>
        </listOfProducts>
        <kineticLaw>
          <math xmlns="http://www.w3.org/1998/Math/MathML">
            <apply>
              <times/>
              <ci> C </ci>
              <ci> k1 </ci>
              <ci> S1 </ci>
            </apply>
          </math>
          <listOfParameters>
            <parameter id="k1" value="0.1" units="per_second"/>
          </listOfParameters>
        </kineticLaw>
      </reaction>
    </listOfReactions>
  </model>
</sbml>
```

For the model above, we create *equation* vertices as follows:

1. [Corresponding to step 1(a) in Section D.] Every **Species** object having `boundaryCondition="false"`, `constant="false"`, and which is referenced as a reactant or product in one or more **Reaction** objects containing **KineticLaw** objects. This generates two vertices, for "S1" and "S2".
2. [Corresponding to step 1(b) in Section D.] Every **AlgebraicRule** object. This generates one vertex, for the model's lone algebraic rule (call it "rule").
3. [Corresponding to step 1(e) in Section D.] Every **KineticLaw** object. This generates one vertex, for the lone kinetic law in the model (call it "law").

1 We create *variable* vertices for the following:

- 2 1. [Corresponding to step 2(a) in Section D.] Every **Species** object having `constant="false"`. This
3 generates two vertices, for "S1" and "S2".
- 4 2. [Corresponding to step 2(b) in Section D.] Every **Compartment** object having `constant="false"`. This
5 generates one vertex, for "C".
- 6 3. [Corresponding to step 2(d) in Section D.] Every **Reaction** object. This generates one vertex, for "R".

7 Note that it is not necessary to include parameters declared within **KineticLaw** objects because they are local
8 to a particular reaction and cannot be affected by rules. With the steps above, we have the following set of
9 graph nodes:

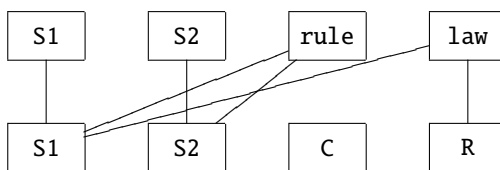
Vertices for equations



Vertices for variables

10 Next, we create edges following the procedure described above. Doing so results in the following graph:

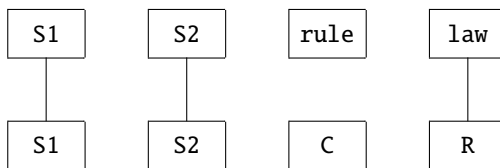
Vertices for equations



Vertices for variables

11 The algorithm of [Hopcroft and Karp \(1973\)](#) can now be applied to search for a maximal matching of the
12 bipartite graph. A maximal matching is a graph in which each vertex is connected to at most one other
13 vertex and the maximum possible number of connections have been made. Doing so here results in the
14 following:
15
16

Vertices for equations



Vertices for variables

17 If the maximal matching of the bipartite graph leaves any equation vertex unconnected, then the model is
18 considered overdetermined. That is the case for the example shown here, because the equation vertex for
19 "rule" is unconnected in the maximal matching.
20

E Mathematical consequences of the fast attribute on Reaction

(Appendix contributed by James C. Schaff, University of Connecticut Health Center, Connecticut, U.S.A.)

Section 4.13.1 described the `fast` flag available on `Reaction`. In this appendix, we discuss the principles involved in interpreting this attribute in the context of a simple biochemical reaction model. The derivation presented here is not fully rigorous and this section is not considered normative; achieving a higher level of rigor would require considerably more background exposition and a much longer appendix. Nevertheless, we hope this section is sufficient to answer unambiguously the question “How should a system of reactions be treated if some of the reactions have `fast=true`?”

Identification of “fast” reactions

First, it is worth noting that the identification of so-called *fast* reactions is actually a modeling issue, not an SBML representation issue. The notion of fast reactions is the following. A system may be decomposable into two sets of reactions, where one set may have characteristic times that are much faster than the other time scales in the system. An approximation that is sometimes useful is to assume that the fast reactions have kinetics that settle infinitely fast compared to the other reactions in the system. In other words, the fast reactions are assumed to be always in equilibrium. This is called a pseudo-steady state approximation (PSSA), and is also known as a quasi-steady state approximation (QSSA). Given a case where the time-scale separation between fast and other reactions in the system is large, an accurate and efficient approach for computing the time-course of the system behavior is to treat the fast reactions as being always in equilibrium.

The key to successful application of a PSSA is that there should be a significant separation of time scales between these fast reactions and other reactions in the system. The determination of which reactions qualify as fast is up to the creator of the model, because there is currently no known general algorithm for doing so.

Simple one-compartment biochemical system model

To explain how to solve a system containing fast reactions, we use a simple model of a biochemical reaction network located in a single compartment. Let \mathbf{x}^* represent a vector of all the species in the system, \mathbf{v}^* a vector of the reaction rates, and \mathbf{A}^* the stoichiometry matrix, with the vector dimension being \mathbf{n}^* . Then the system can be described using the following matrix equation:

$$\frac{d\mathbf{x}^*}{dt} = \mathbf{A}^* \mathbf{v}^*(\mathbf{x}^*)$$

This system can be optionally reduced by noting that mass conservation usually implies there are linear combinations of species quantities in the system and the value of these combinations do not change over time. Identifying these combinations is the topic of structural analysis and is described in the literature (Reder, 1988; Sauro and Ingalls, 2003). Briefly, let \mathbf{N} be defined as the left null space of \mathbf{A}^* :

$$\mathbf{N}\mathbf{A}^* = \mathbf{0}$$

Now, premultiply the previous equation by \mathbf{N} to get

$$\mathbf{N} \frac{d\mathbf{x}^*}{dt} = \mathbf{N}\mathbf{A}^* \mathbf{v}^*(\mathbf{x}^*) = \mathbf{0}$$

Thus, \mathbf{N} captures the space of solutions to the equation

$$\mathbf{m}^T \left(\frac{d\mathbf{x}^*}{dt} \right) = 0$$

where \mathbf{m} is a vector representing the coefficients in a mass conservation relationship, that is, combinations of species that are time-invariant. Now, let

$$r = \text{rank}(\mathbf{A}^*)$$

$$n = \text{dim}(\mathbf{x}^*)$$

Then the system has $n - r$ mass conservation relationships, each of which is a linear equation. We can use these $n - r$ linear equations to solve for $n - r$ dependent variables in terms of r independent variables and the initial masses of all species. Doing that allows us to decompose \mathbf{x}^* into $n - r$ dependent variables \mathbf{x}_d and r independent variables \mathbf{x}_i where \mathbf{L} is an $(n - r) \times r$ matrix that is derived from \mathbf{N} and represents \mathbf{x}_d in terms of \mathbf{x}_i , \mathbf{I} is the $r \times r$ identity matrix, and \mathbf{T} is an $n \times r$ matrix:

$$\mathbf{x}^* \equiv \begin{bmatrix} \mathbf{x}_i \\ \mathbf{x}_d \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{L} \end{bmatrix} \mathbf{x}_i = \mathbf{T} \mathbf{x}_i$$

Using this equation, we can define a new vector of reaction velocities \mathbf{v} in terms of \mathbf{x}_i only:

$$\mathbf{v}(\mathbf{x}_i) \equiv \mathbf{v}^*(\mathbf{T} \mathbf{x}_i)$$

With this \mathbf{v} , we can now write a reduced system by substituting terms. First we define \mathbf{A} as the r independent rows of \mathbf{A}^* corresponding to \mathbf{x}_i . Then:

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{A} \mathbf{v}(\mathbf{x}_i)$$

This is a set of r independent differential equations in r unknowns (i.e., an r -dimensional system). To simplify the notation slightly, let

$$\mathbf{x} \equiv \mathbf{x}_i$$

and, thus,

$$\frac{d\mathbf{x}}{dt} = \mathbf{A} \mathbf{v}(\mathbf{x})$$

Application of a PSSA to biochemical systems

Assume that we have eliminated redundant variables and equations using the mass conservation analysis above. Further assume that we have some external means of classifying some reactions in a given system as being *fast* as discussed earlier. We now need to apply this to the system under study. We begin by decomposing the vector of reaction velocities \mathbf{v} according to fast and slow reactions:

$$\frac{d\mathbf{x}}{dt} = \mathbf{A}_1 \mathbf{v}_f(\mathbf{x}) + \mathbf{A}_2 \mathbf{v}_s(\mathbf{x})$$

In the expression above, \mathbf{A}_1 represents the stoichiometry of the set of reactions operating on the fast time scale, and \mathbf{A}_2 the stoichiometry of the set of reactions operating on a slower time scale. We find the left null space of \mathbf{A}_1 (i.e., the space of solutions to $\mathbf{m}^T [d\mathbf{x}/dt] = 0$ on a fast time scale), and call this matrix \mathbf{B} :

$$\mathbf{B} \mathbf{A}_1 = \mathbf{0}$$

The matrix \mathbf{B} represents the linear combination of species that do not change on a fast time scale, i.e., the slow species in the system. Now, we premultiply the equation for $d\mathbf{x}/dt$ by \mathbf{B} :

$$\begin{aligned} \mathbf{B} \frac{d\mathbf{x}}{dt} &= \mathbf{B} \mathbf{A}_1 \mathbf{v}_f(\mathbf{x}) + \mathbf{B} \mathbf{A}_2 \mathbf{v}_s(\mathbf{x}) \\ &= \mathbf{B} \mathbf{A}_2 \mathbf{v}_s(\mathbf{x}) \end{aligned}$$

where the second line follows from the fact that $\mathbf{B} \mathbf{A}_1 = \mathbf{0}$. The above is an ordinary differential equation in terms of only the slow dynamics. The remaining fast dynamics are handled by applying the pseudo-steady state approximation, with fast transients assumed to have settled with respect to the slow time scale. This produces a system of nonlinear algebraic equations:

$$\mathbf{A}_1 \mathbf{v}_f = \mathbf{0}$$

The last two equations form the system of equations resulting from the application of the PSSA. If $r_1 = \text{rank}(\mathbf{A}_1)$ and $r = \text{rank}(\mathbf{A})$, then there will be r_1 degrees of freedom that will be determined by solving an algebraic system (the equation $\mathbf{A}_1 \mathbf{v}_f = \mathbf{0}$ above), and there will be $r - r_1$ degrees of freedom that will be determined by ordinary differential equations (the equation for $\mathbf{B} d\mathbf{x}/dt$).

F Processing and validating SBase notes and Constraint message content

In Section 3.2.3 and Section 4.12.2, we discussed the **notes** element on **Sbase** and the **message** element on **Constraint**, respectively. These elements can contain a number of possible forms of XHTML content. In this appendix, we describe a general procedure for how application software can process such content. We concentrate on the common case of an SBML-reading application that needs to take the contents and pass it to an XHTML display and/or editing function obtained from a third-party API library. The content of the **notes** and **message** may not be a complete XHTML document, so the application will have to perform some processing before handing it to the XHTML editor or validator. How should this be done?

Based on the three forms of **Sbase notes** content described in Section 3.2.3 and the identical forms for **Constraint message** described in Section 4.12.2, there are only three cases possible. Here we give an example approach for handling them, although the actual implementation details will differ depending on various factors such as the requirements of the software libraries being used. This example approach would be performed for each **notes** and **message** to be viewed or edited:

Step 1. If the XHTML viewing/editing function requires a fully compliant XML document, the SBML application could create a temporary data object containing an appropriate XML declaration and a DOCTYPE declaration; otherwise, the XML data object can be initially blank.

Step 2. The application should look at the first element inside the **notes** or **message** (or rather, the first element that is not an XML comment), and take action based on the following three possibilities:

- If the first element begins with `<html xhtml="http://www.w3.org/1999/xhtml">`, the application could assume that the content is a complete XHTML document and insert this into the temporary data object.
- Else, if the first element is `<body>`, the application should insert the following into the temporary data object,

```
<html xhtml="http://www.w3.org/1999/xhtml">
  <head><title></title></head>
```

then insert the content of the **notes** or **message**, and finally insert a closing `</html>`.

- Else, if the content begins with neither of the above elements, the application should insert the following into the temporary data object,

```
<html xhtml="http://www.w3.org/1999/xhtml">
  <head><title></title></head>
  <body>
```

then insert the content of the **notes** or **message**, and finally insert `</body></html>` to close the XHTML document.

The result of the above would be a temporary XML data object that the application could then pass to the XHTML viewing/editing API function.

G Major changes between versions of SBML Level 2 and implications for backward compatibility

In Section 1.3, we summarized the language features removed and deprecated in SBML Level 2 Version 3. Here we describe the changes in more detail. We describe the cumulative changes introduced since SBML Level 2 Version 1.

G.1 Between Version 2 and Version 1

The following features were removed between SBML Level 2 Version 1 and Version 2:

- The **offset** attribute on **UnitDefinition**. (See Section 4.4.2.) The definition of offsets in SBML Level 2 Version 1 was in fact incorrect; moreover, a proper implementation would have required a complete change in the SBML unit scheme. Few models appeared to use offsets on unit definitions, so the impact of this change on models is expected to be small.
- The “Celsius” predefined unit. (See Section 4.4.2.) The removal of offsets on unit definitions meant an inconsistency existed if the Celsius predefined unit was left in the system. Removing Celsius removes the inconsistency. Alternative ways of using Celsius are discussed in Section 4.4.2.
- The **substanceUnit** and **timeUnits** attributes on **KineticLaw**. (See Section 4.13.5.) The ability to redefine the substance units on each reaction separately, coupled with other features in SBML Level 2 Version 1, created the opportunity for defining a valid system of reactions which potentially could not be combined into a consistent system of equations without external knowledge.

The following features were deprecated between SBML Level 2 Version 1 and Version 2:

- The **charge** attribute on **Species**. (See Section 4.8.7.) This attribute does not appear to be supported by any existing software, and moreover, since its value cannot be accessed from mathematical formulas in SBML, the impact of this change is expected to be small.

The following additional changes were made between SBML Level 2 Version 1 and Version 3:

- SBML Level 2 Version 1 did not clearly specify the value space of integer and floating-point numbers permitted in the MathML expressions in SBML; moreover, it used the XML Schema type “**integer**” instead of SBML Level 2 Version 2’s “**int**”. Although extremely unlikely, some previously valid SBML Level 2 Version 1 documents *may* not be valid in Version 2 and Version 3 as a result of these changes. See Sections 3.1.3, 3.1.5 and 3.4.2 for more information.
- SBML Level 2 Version 1 did not define a default value for the attribute **fast** on **Reaction**. SBML Level 2 Version 2 introduced a default value, and the value is “**false**”. Further, SBML now requires that software tools *must* respect the value or indicate to the user that they do not have the capacity to do so. See Section 4.13.1.
- As of SBML Level 2 Version 2, SBML is somewhat stricter about how the content of **annotation** elements must be organized and written.. Previously valid SBML Level 2 Version 1 documents *may* need changes to their **annotation** elements to comply with the specification beginning with Version 2 and Version 3. See Section 3.2.4 for more details.
- As of SBML Level 2 Version 2, SBML is slightly stricter about how the content of **notes** elements must be organized. Previously valid SBML Level 2 Version 1 documents *may* need changes to their **notes** elements to comply with the specification beginning with Version 2 and Version 3. See Section 3.2.3 for more details.
- SBML Level 2 Version 2 corrected numerous errata and ambiguities discovered in SBML Level 2 Version 1. These errata are listed on the project web site at <http://sbml.org>. As a result of changes

1 to SBML Level 2 implied by these errata, some existing SBML Level 2 Version 1 models, even when
2 modified as explained above, may still not be compliant with Version 2 or Version 3. The ultimate
3 impact of the changes depends on the specific features used by a given model and the assumptions
4 under which the model was created.

5 G.2 Between Version 3 and Version 2

6 The following features were removed between SBML Level 2 Version 2 and Version 3:

- 7 • The **spatialSizeUnits** attribute on **Species**. (See Section 4.8.5.) This attribute introduced an implicit
8 unit conversion between the spatial units used in defining the quantity of a species and the size of
9 the compartment in which the species was located. Moreover, the semantics of **spatialSizeUnits**
10 were confusing and required complicated unit conversions to be written explicitly into reaction rate
11 expressions by either the modeler or the software. Although the conversions could be worked out
12 unambiguously, the potential for error was judged to exceed by far the utility of this feature.
- 13 • The **timeUnits** attribute **Event**. (See Section 4.14.) This attribute was judged to add needless com-
14 plexity and inconsistency. For instance, the ability to change the time units of the delay of an **Event**
15 to be different from the units of time for the whole model meant that computing an **Event**'s time of
16 triggering and its delay might have to be done using two different sets of units. The ability to redefine
17 the units of time for the delay of an **Event** was also inconsistent with the lack of such an attribute on
18 other SBML components involving an element of time; for example, **RateRule**, and now **KineticLaw**,
19 have no such attributes.

20 The following additional changes were made between SBML Level 2 Version 2 and Version 3:

- 21 • The definition of the XML type **ID** was incorrectly given in the SBML Level 2 Version 2 specification.
22 This type is used as the type of the attribute **metaid** on **Sbase**. The error in the definition of **ID**
23 was such that the type did not include the colon (:) character and all Unicode characters actually
24 permitted in XML **ID**. This change is therefore entirely backward compatible: all models with valid
25 **metaid** values valid prior to SBML Level 2 Version 3 are still valid under the new definition.
- 26 • The SBML specifications prior to SBML Level 2 Version 3 did not indicate what units are as-
27 sumed for literal numbers appearing in MathML expressions (i.e., numbers inside MathML's **cn** el-
28 ements). SBML Level 2 Version 3 stipulates that there are no units associated with numbers (not even
29 "dimensionless"), and provides suggestions for how to associate units with numbers (Section 3.4.2).
- 30 • The SBML specifications prior to SBML Level 2 Version 3 did not make clear what units are required
31 by the arguments to various MathML operators and other constructs. SBML Level 2 Version 3 clarifies
32 this (Section 3.4.10).
- 33 • The **UnitKind** enumeration previously defined in the context of **Unit** and **UnitDefinition** has been elim-
34 inated in favor of simply defining the symbols as reserved words in the value space of **UnitSid**. This
35 has no effect on written models and is completely backward compatible. It was done to resolve the
36 problem that, previously, the values in **UnitKind** were technically inaccessible from attributes whose
37 data type was **UnitSid**.
- 38 • The SBML specification did not point out that the value space of the data type **boolean** is differ-
39 ent in MathML 2.0 and XML Schema 1.0. This means that the permitted values of attributes on
40 SBML objects is different from the values permitted in MathML formulas. (Section 3.1.2 explains the
41 difference.)
- 42 • The SBML specifications were inconsistent about the permitted number of items inside **listOf_____**
43 lists: the text portion of the specifications claimed the lists could have zero length, but the XML Schema
44 definition for SBML required one or more items. As of SBML Level 2 Version 3, the specification is
45 consistent on requiring one or more items inside these lists.

- 1 ● The SBO term hierarchy (Section 5) has grown in the time intervening between SBML Level 2 Version 2
2 and Version 3, and the mapping of terms between SBO and SBML components was revised as a result
3 of community discussions during the 2006 SBML Forum meeting.
- 4 ● The **sboTerm** attribute, introduced on many components in SBML Level 2 Version 2, has been moved
5 to **Sbase** as an optional attribute and removed from the individual components. The result is that
6 all model components may now have SBO terms associated with them. This change is completely
7 backward compatible.
- 8 ● A number of validation rules in Appendix C have been introduced; some were missing from previous
9 specifications, and some were added to cover changes introduced in SBML Level 2 Version 3 (for
10 example, for validation of SBO terms assigned to various SBML model components).
- 11 ● The SBML specifications prior to SBML Level 2 Version 3 did not adequately explain the assumptions
12 regarding XML namespace declarations within the **annotation** element on SBML components. SBML
13 Level 2 Version 3 makes these assumptions more explicit, including the assumption that applications
14 may not preserve another applications annotations unless those annotations are self-contained with the
15 XML Namespace declaration. See Section 3.2.4 for more details.

References

- Abbott, A. (1999). Alliance of US labs plans to build map of cell signalling pathways. *Nature*, 402:219–200.
- Abramowitz, M. and Stegun, I. A., editors (1977). *Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*. Dover Publications Inc.
- Ausbrooks, R., Buswell, S., Carlisle, D., Dalmas, S., Devitt, S., Diaz, A., Froumentin, M., Hunter, R., Ion, P., Kohlhase, M., Miner, R., Poppelier, N., Smith, B., Soiffer, N., Sutor, R., and Watt, S. (2003). Mathematical Markup Language (MathML) Version 2.0 (second edition): W3C Recommendation 21 October 2003. Available via the World Wide Web at <http://www.w3.org/TR/2003/REC-MathML2-20031021/>.
- Biron, P. V. and Malhotra, A. (2000). XML Schema part 2: Datatypes (W3C candidate recommendation 24 October 2000). Available via the World Wide Web at <http://www.w3.org/TR/xmlschema-2/>.
- Bosak, J. and Bray, T. (1999). XML and the second-generation Web. *Scientific American*, 280(5):89–93.
- Bray, T., D. Hollander, D., and Layman, A. (1999). Namespaces in XML. World Wide Web Consortium 14-January-1999. Available via the World Wide Web at <http://www.w3.org/TR/1999/REC-xml-names-19990114/>.
- Bray, T., Paoli, J., Sperberg-McQueen, C. M., and Maler, E. (2000). Extensible markup language (XML) 1.0 (second edition), W3C recommendation 6-October-2000. Available via the World Wide Web at <http://www.w3.org/TR/1998/REC-xml-19980210/>.
- Bray, T., Paoli, J., Sperberg-McQueen, C. M., Maler, E., and Yergeau, F. (2004). Extensible markup language (XML) 1.0 (third edition), W3C recommendation 4-February-2004. Available via the World Wide Web at <http://www.w3.org/TR/2004/REC-xml-20040204/>.
- Bureau International des Poids et Mesures (2000). The International System of Units (SI) supplement 2000: addenda and corrigenda to the 7th edition (1998). Available via the World Wide Web at <http://www.bipm.fr/pdf/si-supplement2000.pdf>.
- Chartrand, G. (1977). *Introductory Graph Theory*. Dover Publishing, Inc., New York.
- DCMI Usage Board (2005). DCMI Metadata Terms. Available via the World Wide Web at <http://www.dublincore.org/documents/dcmi-terms/>.
- Dublin Core Metadata Initiative (2005). Dublin core metadata initiative. Available via the World Wide Web at <http://dublincore.org/>.
- Eriksson, H.-E. and Penker, M. (1998). *UML Toolkit*. John Wiley & Sons, New York.
- Fallside, D. C. (2000). XML Schema part 0: Primer (W3C candidate recommendation 24 October 2000). Available via the World Wide Web at <http://www.w3.org/TR/xmlschema-0/>.
- Gillespie, D. (1977). Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81:2340–2361.
- Gilman, A. (2000). A letter to the signaling community. Alliance for Cellular Signaling, The University of Texas Southwestern Medical Center. Available via the World Wide Web at http://afcs.swmed.edu/afcs/Letter_to_community.htm.
- Harold, E. R. and Means, E. S. (2001). *XML in a Nutshell*. O'Reilly & Associates.
- Hedley, W. J., Nelson, M. R., Bullivant, D., Cuellar, A., Ge, Y., Grehlinger, M., Jim, K., Lett, S., Nickerson, D., Nielsen, P., and Yu, H. (2001). CellML specification. Available via the World Wide Web at http://www.cellml.org/public/specification/20010810/cellml_specification.html.
- Hopcroft, J. E. and Karp, R. M. (1973). An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2(4):225–231.

- 1 Hucka, M., Finney, A., Sauro, H. M., and Bolouri, H. (2001). Systems Biology Markup Language (SBML)
2 Level 1: Structures and facilities for basic model definitions. Available via the World Wide Web at
3 <http://www.sbml.org>.
- 4 Hucka, M., Finney, A., Sauro, H. M., Bolouri, H., Doyle, J. C., Kitano, H., Arkin, A. P., Bornstein, B. J.,
5 Bray, D., Cornish-Bowden, A., Cuellar, A. A., Dronov, S., Gilles, E. D., Ginkel, M., Gor, V., Goryanin, I. I.,
6 Hedley, W. J., Hodgman, T. C., Hofmeyr, J.-H., Hunter, P. J., Juty, N. S., Kasberger, J. L., Kremling, A.,
7 Kummer, U., Le Novère, N., Loew, L. M., Lucio, D., Mendes, P., Minch, E., Mjolsness, E. D., Nakayama,
8 Y., Nelson, M. R., Nielsen, P. F., Sakurada, T., Schaff, J. C., Shapiro, B. E., Shimizu, T. S., Spence,
9 H. D., Stelling, J., Takahashi, K., Tomita, M., Wagner, J., and Wang, J. (2003). The Systems Biology
10 Markup Language (SBML): A medium for representation and exchange of biochemical network models.
11 *Bioinformatics*, 19(4):524–531.
- 12 Iannella, R. (2001). Representing vCard objects in RDF/XML. Available via the World Wide Web at
13 <http://www.w3.org/TR/vcard-rdf>.
- 14 Jacobs, I. (2004). World wide web consortium process document. Available via the World Wide Web at
15 <http://www.w3.org/2004/02/Process-20040205/>.
- 16 Kokkeliink, S. and Schwänzl, R. (2002). Expressing qualified Dublin Core in RDF/XML. Available via the
17 World Wide Web at <http://dublincore.org/documents/dcq-rdf-xml/index.shtml>.
- 18 Lassila, O. and Swick, R. (1999). Resource description framework (RDF) model and syntax specification.
19 Available via the World Wide Web at <http://www.w3.org/TR/REC-rdf-syntax/>.
- 20 Le Novère, N., Finney, A., Hucka, M., Bhalla, U., Campagne, F., Collado-Vides, J., Crampin, E. J., Halstead,
21 M., Klipp, E., Mendes, P., Nielsen, P., Sauro, H., Shapiro, B., Snoep, J. L., Spence, H. D., and Wanner,
22 B. L. (2005). Minimum information requested in the annotation of biochemical models (MIRIAM). *Nature*
23 *Biotechnology*, 23:1509–1515.
- 24 Oestereich, B. (1999). *Developing Software with UML: Object-Oriented Analysis and Design in Practice*.
25 Addison-Wesley Publishing Company.
- 26 Pemberton, S., Austin, D., Axelsson, J., Celik, T., Dominiak, D., Elenbaas, H., Epperson, B., Ishikawa, M.,
27 Matsui, S., McCarron, S., Navarro, Peruvemba, S., Relyea, R., Schnitzenbaumer, S., and Stark, P. (2002).
28 XHTMLTM 1.0 the Extensible HyperText Markup Language (second edition): W3C Recommendation 26
29 January 2000, revised 1 August 2002. Available via the World Wide Web at [http://www.w3.org/TR/](http://www.w3.org/TR/xhtml1/)
30 [xhtml1/](http://www.w3.org/TR/xhtml1/).
- 31 Popel, A. and Winslow, R. L. (1998). A letter from the directors... Center for Computational Medicine
32 & Biology, Johns Hopkins School of Medicine, Johns Hopkins University. Available via the World Wide
33 Web at <http://www.bme.jhu.edu/ccmb/ccmbletter.html>.
- 34 Powell, A. and Johnston, P. (2003). Guidelines for implementing Dublin Core in XML. Available via the
35 World Wide Web at <http://dublincore.org/documents/dc-xml-guidelines/index.shtml>.
- 36 Reder, C. (1988). Metabolic Control Theory: a structural approach. *Journal of Theoretical Biology*, 135:175–
37 201.
- 38 Sauro, H. M. and Ingalls, B. (2003). Conservation analysis in biochemical networks: Computational issues
39 for software writers. Available via the World Wide Web at [http://www.math.uwaterloo.ca/~bingalls/](http://www.math.uwaterloo.ca/~bingalls/Pubs/conservation.pdf)
40 [Pubs/conservation.pdf](http://www.math.uwaterloo.ca/~bingalls/Pubs/conservation.pdf).
- 41 Smaglik, P. (2000). For my next trick... *Nature*, 407:828–829.
- 42 Thompson, H. S., Beech, D., Maloney, M., and Mendelsohn, N. (2000). XML Schema part 1: Structures
43 (W3C candidate recommendation 24 October 2000). Available via the World Wide Web at [http://www.](http://www.w3.org/TR/xmlschema-1/)
44 [w3.org/TR/xmlschema-1/](http://www.w3.org/TR/xmlschema-1/).

- 1 Unicode Consortium (1996). *The Unicode Standard, Version 2.0*. Addison-Wesley Developers Press, Reading,
2 Massachusetts.
- 3 W3C (2000a). Naming and addressing: URIs, URLs, Available via the World Wide Web at [http:
4 //www.w3.org/Addressing/](http://www.w3.org/Addressing/).
- 5 W3C (2000b). W3C's math home page. Available via the World Wide Web at <http://www.w3.org/Math/>.
- 6 W3C (2004a). Rdf/xml syntax specification (revised). Available via the World Wide Web at [http://www.
7 w3.org/TR/rdf-syntax-grammar/](http://www.w3.org/TR/rdf-syntax-grammar/).
- 8 W3C (2004b). Resource description framework (RDF). Available via the World Wide Web at [http://www.
9 w3.org/RDF/](http://www.w3.org/RDF/).
- 10 Wilkinson, D. J. (2006). *Stochastic Modelling for Systems Biology*. Chapman & Hall/CRC.
- 11 Wolf, M. and Wicksteed, C. (1998). Date and time formats. Available via the World Wide Web at [http:
12 //www.w3.org/TR/NOTE-datetime](http://www.w3.org/TR/NOTE-datetime).
- 13 Zwillinger, D., editor (1996). *Standard Mathematical Tables and Formulae*. CRC Press LLC, 30th edition.