JML Reference Manual

DRAFT, $Revision: 1.65 $
11 December 2003

Gary T. Leavens, Erik Poll, Curtis Clifton, Yoonsik Cheon, Clyde Ruby, David Cok, Joseph Kiniry

Department of Computer Science
226 Atanasoff Hall
Iowa State University
Ames, Iowa 50011-1041, USA
e-mail: jml@cs.iastate.edu

Version Information:
@(#) $Id: jmlrefman.texinfo,v 1.65 2003/12/09 20:49:42 cruby Exp $

1 Introduction

JML is a notation for formally specifying the behavior and interfaces of Java [Arnold-Gosling-Holmes00] [Gosling-etal00] classes and methods.

The goal of this reference manual is to precisely record the design of JML. We include both informal semantics (intentions) and where possible formal semantics (usually in the form of verification conditions). We also discuss the implications for various tools (such as the run-time assertion checker, static checkers such as ESC/Java, and documentation generators such as jmldoc [Burdy-etal03]).

We try to give examples and explanations, and we hope that these will be helpful to readers trying to learn about formal specification using JML. However, this manual is not designed to give all the background needed to write JML specifications, nor to give the prospective user an overview of a useful subset of the language. For this background, we recommend starting with the paper "JML: A notation for detailed design" [Leavens-Baker-Ruby99], and continuing with the paper "Preliminary Design of JML" [Leavens-Baker-Ruby02]. Both of these are available from the JML web site '`http://www.jmlspecs.org/`', where further readings and examples may also be found.

Readers with the necessary background, and users wanting more details may, we hope, profit from reading this manual. We suggest reading this manual starting with chapters 1-3, skimming chapter 4 quickly, skimming chapter 5 to get the idea of what declarations mean in JML (paying a bit more attention in section [[[5.4??]]] to the sorts of declarations), and then reading the chapters on class specifications (chapter 6) and method specifications (chapter 7), paying particular attention to the examples. After that, one can use the rest of this manual as a reference.

The rest of this chapter describes some of the fundamental ideas and background behind JML.

1.1 Behavioral Interface Specifications

JML is a *behavioral interface specification language* (BISL) that builds on the Larch approach [Guttag-Horning93] [Guttag-Horning-Wing85b] and that found in APP [Rosenblum95] Eiffel [Meyer92b] [Meyer97]. In this style of specification, which might be called model-oriented [Wing90a], one specifies both the interface of a method or abstract data type and its behavior [Lamport89]. In particular JML builds on the work done by Leavens and others in Larch/C++ [Leavens-Baker99] [Leavens96b] [Leavens97c]. (Indeed, large parts of this manual are adapted wholesale from the Larch/C++ reference manual [Leavens97c].)

The *interface* of the method or type is the information needed to use it from other programs. In the case of JML, this is the Java syntax and type information needed to call a method or use a type. For a method the interface includes such things as the name of the method, its number of arguments, its return type, what exceptions it may throw, and so on. JML specifies interface information using Java's syntax.

The *behavior* of a method or type describes the state transformations it can perform. The behavior of a method is specified by describing: the set of states in which calling the method is defined, the set of locations the method is allowed to assign to (and hence change), and the relation between the states before and after the method is invoked. The states for

which the method is defined can be formally described by a logical assertion, called the method's *precondition*. The allowed relationships between these states and the states that may result from a call are formally described by another logical assertion called the method's *postcondition*. The set of locations the method is allowed to assign to is described by listing the names of these locations in what is called the method's *frame axiom* [Borgida-etal95].

The behavior of an abstract data type (ADT), which is implemented by a class in Java, is specified by describing a set of abstract fields for its objects and by specifying the behavior of its methods (as described above). The abstract fields for an object can be specified either by using JML's model and ghost fields [Cheon-etal03], which are specification-only fields, or by specifying some of the fields used in the implementation as `spec_public` or `spec_protected`. Thus, like other specification languages, such as Z [Hayes93] [Spivey92], or Fresco [Wills92b] JML requires the user to model an instance as a collection of instance variables.

In the rest of this section we give a first example of an ADT specification, and then describe more of the related work behind such specifications and JML's design in this area.

1.1.1 A First Example

For example, consider the following JML specification of a simple Java abstract class `IntHeap`. (An explanation of the notation follows the specification.)

```
package org.jmlspecs.samples.jmlrefman;                // line 1
                                                       // line 2
public abstract class IntHeap {                        // line 3
                                                       // line 4
    //@ public model non_null int [] elements;         // line 5
                                                       // line 6
    /*@ public normal_behavior                         // line 7
      @   requires elements.length >= 1;               // line 8
      @   assignable \nothing;                         // line 9
      @   ensures \result                              // line 10
      @        == (\max int j;                         // line 11
      @               0 <= j && j < elements.length;   // line 12
      @               elements[j]);                    // line 13
      @*/                                              // line 14
    public abstract /*@ pure @*/ int largest();        // line 15
                                                       // line 16
    //@ ensures \result == elements.length;           // line 17
    public abstract /*@ pure @*/ int size();           // line 18
};                                                     // line 19
```

The interface of this class consists of lines 1, 3, 15, and 18. Line 3 specifies class name, and the fact that it is both public and abstract. Lines 15 and 18, apart from their comments, give the interface information for the methods of this class.

The behavior of this class is specified in the JML annotations found in the special comments that have an at-sign (`@`) as their first character following the usual comment beginning. Such lines look like comments to Java, but are interpreted by JML and its tools. For example, line 5 starts with an annotation comment marker of the form `//@`, and

this annotation continues until the `//` towards the end of the line, which starts a comment within the annotation which even JML ignores. The other form of such annotations can be seen on lines 7 through 14 line 17, and on lines 15 and 18. These annotations start with the characters `/*@` and end with either `@*/` or `*/`; within such annotations, at-signs (`@`) at the beginnings of lines are ignored by JML. (See Chapter 4 [Lexical Conventions], page 13, for more details about annotations.)

The first annotation, on line 5 of the figure above, gives the specification of a field, named `elements`, which is part of this class's behavioral specification. Ignoring, for the moment the extra JML modifiers, one should think of this field, in essence, as being declared like:

```
public int[] elements;
```

That is, it is a public field with an integer array type. And within specifications it is treated as such. However, because it is declared in an annotation, this field cannot be manipulated by Java code. Therefore, for example, the fact that the field is declared public is not a problem, because it cannot be directly changed by Java code.

Such declarations of fields in annotations should be marked as specification-only fields, using the JML modifier `model`.[1] A model field should be thought of as an abstraction of the concrete fields that are used to implement this type and its subtypes. (See Section 10.3 [Represents Clauses], page 36, for a discussion of how to specify the connection between the concrete fields and such model fields. See also the paper by Cheon et al. [Cheon-etal03]) That is, we imagine that objects that are instances of the type `IntHeap` have such a field, whose value is determined by the concrete fields that are known to Java in the actual object. Of course at runtime, objects of type `IntHeap` have no such field, the model fields are purely imaginary. Model fields are thus a convenient fiction that is useful for describing the behavior of an ADT. One does not have to worry about their cost (in space or time), and should only be concerned with how they clarify the behavior of an ADT.

The other annotation used on line 5 is `non_null`. This just says that in any publicly-visible state, the value of `elements` must not be `null`. It is thus a simple kind of invariant (see Section 10.1 [Invariants], page 28).

In the above specification of `IntHeap`, the specification of each method precedes its interface declaration. This follows the usual convention of Java tools, such as JavaDoc, which put such descriptive information in front of the method. In JML, it is also possible to put the specification just before the semicolon (`;`) following the method's interface information (see Chapter 11 [Method Specifications], page 38), but we will usually not to do that in this document.

The specification of the method `largest` is given on lines 7 through 15. Line 7 says that this is a public, normal behavior specification. JML permits several different specifications for a given method, which can be of different privacy levels [Ruby-Leavens00]. The modifier `public` says that the specification is intended for the use of clients. (If the privacy modifier had been `protected`, for example, then the specification would have been intended for subclasses.)

The keyword `normal_behavior` tells JML several things. First, it says that the specification is a heavyweight method specification, as opposed to a lightweight method specification like that given on line 17. A *heavyweight* specification uses one of JML's behavior keywords,

[1] Another way to declare a specification-only field is to use the `ghost` modifier ([[[Add pxref]]]).

like `normal_behavior`, which tells JML that the method specification is intended to be complete. By contrast, a *lightweight* specification does not use one of JML's behavior keywords, which tells JML that the specification is incomplete in the sense of the only contains what the specifier wanted to write down, and not everything possible.[2] Second, the keyword `normal_behavior` tells JML that when the precondition of this method is met, then the method must return normally, without throwing an exception. (See Chapter 11 [Method Specifications], page 38, for more details.)

The heart of the method specification of `largest` is found on lines 7 through 13. This part of the specification gives the method's precondition, on line 8, frame axiom, on line 9, and postcondition, on lines 10 through 13. The precondition is contained in the `requires` clause. The frame axiom is contained in the `assignable` clause. The postcondition is contained in the `ensures` clause.[3]

The precondition in the requires clause says that the length of `elements` must be at least 1 before this method can be called. If that is not true, then the method is under no obligation to fulfill the rest of the specified behavior.

The frame axiom in the assignable clause says that the method may not assign to any locations (i.e. fields of objects) that are visible outside the method and which existed before the method started execution. (The method may still modify its local variables.) This form of the frame axiom is quite common.[4] Note that in assignable clauses and in assertions, JML uses keywords that start with a backslash (\), to avoid interfering with identifiers in the user's program. Examples of this are `\nothing` on line 9 and `\result` on line 10.

The postcondition in the ensures clause, on lines 10 through 13, says that the result of the method (`\result`) must be equal to the maximum integer found in the array `elements`. This postcondition uses JML's `\max` quantifier (lines 11 through 13). Such a quantifier is always parenthesized, and can consist of three parts. The first is a declaration of some quantified variables, in this case the integer `j` on line 11. The second is a *range predicate*, on line 12, which constrains the quantified variables. The third is the *body* of the quantifier, on line 13, which in this case describes the elements of the array from which the maximum value is taken.

The methods `largest` and `size` are both specified using the JML modifier `pure`. This modifier says that the method has no side effects, and allows the method to be used in assertions if desired.

Line 17 contains the single annotation that makes up the lightweight specification of the method `size`. As a lightweight specification with only an ensures clause, this says nothing about the precondition or frame axiom of `size`, although the use of `pure` on line 18 gives an implicit frame axiom. Such a form of specification is useful when one only cares to state (the important) part of a method's specification. It is also useful when first learning JML,

---

[2] Lightweight specifications come from ESC/Java, which does not understand JML's heavyweight specification syntax.

[3] JML also has various synonyms for these keywords; one can use `pre` for `requires`, `modifies` or `modifiable` for `assignable`, and `post` for `ensures` if desired. See Chapter 11 [Method Specifications], page 38, for more details.

[4] However, unlike Larch BISLs and earlier versions of JML, this is not the default for an omitted `assignable` clause (see Section 11.9.8 [Assignable Clauses], page 52). Thus line 9 cannot be omitted without changing the meaning of the specification.

and when one is using tools, such as ESC/Java, that do not need or understand heavyweight specifications.

The specifications of the methods `largest` and `size` above are very precise: they give a complete specification of what the methods do. We can also give JML specifications that are far less detailed. For example, we could just specify that the result of `size` is non-negative, with a postcondition

```
//@  \result >= 0;
```

instead of the postcondition given earlier.

1.1.2 Historical Precedents

JML combines ideas from Eiffel [Meyer92a] [Meyer92b] [Meyer97] with ideas from model-based specification languages such as VDM [Jones90] and the Larch family [Guttag-Horning93] [LeavensLarchFAQ] [Wing87] [Wing90a]. It also adds some ideas from the refinement calculus [Back88] [Back-vonWright89a] [Back-vonWright98] [Morgan-Vickers94] [Morgan94] (see Chapter 19 [Refinement], page 78). In this section we describe the advantages and disadvantages of these approaches.

Formal, model-based languages such as those typified by the Larch family build on ideas found originally in Hoare's work. Hoare used pre- and postconditions to describe the semantics of computer programs in his famous article [Hoare69]. Later Hoare adapted these axiomatic techniques to the specification and correctness proofs of abstract data types [Hoare72a]. To specify an ADT, Hoare described a mathematical set of abstract values for the type, and then specified pre- and postconditions for each of the operations of the type in terms of how the abstract values of objects were affected. For example, one might specify a class `IntHeap` using abstract values of the form `empty` and `add(i,h)`, where `i` is an `int` and `h` is an `IntHeap`. These notations form a mathematical vocabulary used in the rest of the specification.

There are two advantages to writing specifications with abstract values instead of directly using Java variables and data structures. The first is that by using abstract values, the specification does not have to be changed when the particular data structure used in the program is changed. This permits different implementations of the same specification to use different data structures. Therefore the specification forms a contract between the rest of the program in the implementation, which ensures that the rest of the program is also independent of the particular data structures used [Liskov-Guttag86] [Meyer97] [Meyer92a] [Parnas72]. Second, it allows the specification to be written even when there are no implementation data structures, as is the case for `IntHeap`.

This idea of model-oriented specification has been followed in VDM [Jones90], VDM-SL [Fitzgerald-Larsen98] [ISO96], Z [Hayes93] [Spivey92], and the Larch family [Guttag-Horning93]. In the Larch approach, the essential elaboration of Hoare's original idea is that the abstract values also come with a set of operations. The operations on abstract values are used to precisely describe the set of abstract values and to make it possible to abbreviate interface specifications (pre- and postconditions for methods). In Z one builds abstract values using tuples, sets, relations, functions, sequences, and bags; these all come with pre-defined operations that can be used in assertions. In VDM one has a similar collection of mathematical tools to describe abstract values, and another set of pre-defined

operations. In the Larch approach, there are some pre-defined kinds of abstract values (found in Guttag and Horning's LSL Handbook, Appendix A of [Guttag-Horning93]), but these are expected to be extended as needed. (The advantage of being able to extend the mathematical vocabulary is similar to one advantage of object-oriented programming: one can use a vocabulary that is close to the way one thinks about a problem.)

However, there is a problem with using mathematical notations for describing abstract values and their operations. The problem is that such mathematical notations are an extra burden on a programmer who is using a specification language to document programs. In the case of JML, this is a problem for Java programmers. This is the essential insight that JML takes from the Eiffel language [Meyer92a] [Meyer92b] [Meyer97]. Eiffel is a programming language with built-in specification constructs. It features pre- and postconditions, although it has no direct support for frame axioms. Programmers like Eiffel because they can easily read the assertions, which are written in Eiffel's own expression syntax. However, Eiffel does not provide support for specification-only variables, and it does not provide much explicit support for describing abstract values. Because of this, it is difficult to write specifications that are as mathematically complete in Eiffel as one can write in a language like VDM or Larch/C++.

JML attempts to combine the good features of these approaches. From Eiffel we have taken the idea that assertions can be written in a language that is based on Java expressions. We also adopt the "`old`" notation from Eiffel, which appears in JML as `\old`, instead of the Larch-style annotation of names with state functions. To make it easy to write more complete specifications, however, we use various semantic ideas from model-based specification languages. In particular we use a variant of abstract value specifications, where one describes the abstract value of an object implicitly using several model fields. These specification-only fields allow one to implicitly partition the abstract value of an object into smaller chunks, which helps in stating frame axioms. More importantly, we hide the mathematical notation behind a facade of Java classes. This makes it so the operations on abstract values appear in familiar (although perhaps verbose) Java notation, and also insulates JML from the details of the particular mathematical logic used to do reasoning.

1.2 What is JML Good For?

JML is a formal specification language tailored to Java. Its basic use is thus the formal specification of the behavior of Java program modules. As it is a behavioral interface specification language, JML specifies how to use such Java program modules from *within* a Java program; hence JML is *not* designed for specifying the behavior of an entire program. So the question "what is JML good for?" really boils down to the following question: what good is formal specification for Java program modules?

The two main benefits in using JML are

- the precise, unambiguous specification of the behavior of Java program modules (i.e., classes and interfaces), and documentation of Java code,
- the possibility of tool support [Burdy-etal03].

A JML specification can be a completely formal contract about an interface and its behavior. Because it is an interface specification, one can record all the Java details about the interface, such as the parameter mechanisms, whether the method is `final`, `protected`,

etc.; if one used a specification language such as VDM-SL or Z, which is not tailored to Java, then one could not record such details of the interface, which could cause problems in code integration. For example, in JML one can specify the precise conditions under which certain exceptions may be thrown, something which is difficult in a specification language that is not tailored to Java and that doesn't have the notion of an exception.

In addition to precisely specifying the behavior of a Java class that is visible to its clients, JML can also be used to document design decisions taken in the implementation of a class, notably by recording the class invariants that the implementation is designed to maintain.

One can use JML either before coding, or as documentation of the code. The notation is indifferent to the methodological questions; designing before coding is recommended, but documentation after the fact is better than none.

Reasons for formal documentation of interfaces and their behavior, using JML, include the following.

- One can ship the object code for a class library to customers, without sending the source code. Customers would have documentation that is precise, unambiguous, but not overly specific. Customers would not have the code, protecting proprietary rights. In addition, customers would not rely on details of the implementation of the library that they might otherwise glean from the code, easing the process of improving the code in future releases.
- One can use a formal specification to analyze certain properties of a design carefully or formally (see [Hall90] and Chapter 7 of [Guttag-Horning93]). In general, the act of formally specifying a program module has salutary effects on the quality of the design.
- One can use the JML specification as an aid to careful reasoning about the correctness of code, or even for formal verification [Huisman01] [Poll-Jacobs00].

There is one additional benefit from using JML. It is that JML allows one to record not just public interfaces and behavior, but also some detailed design decisions. That is, in JML, one can specify not just the public interface of a Java class, but also behavior of a class's protected and private interfaces. Formally documenting a base class's protected interface and "subclassing contract" allows programmers to implement derived classes of such a base class without looking at its code [Ruby-Leavens00].

Recording the private interface of a class may be helpful in program development or maintenance. Usually one would expect that the public interface of a class would be specified, and then separate, more refined specifications would be given for use by derived classes and for detailed implementation (and friend classes). (See Chapter 19 [Refinement], page 78, for how to record each level in JML.)

The reader may also wish to consult the "Preliminary Design of JML" [Leavens-Baker-Ruby02] for a discussion of the goals that are behind JML's design. Apart from the improved precision in the specifications and documentation of code, the main advantage of using a formal specification language, as opposed to informal natural language, is the possibility of tool support. One specific goal that has emerged over time is that JML should be able to unify several different tool-building efforts in the area of formal methods.

The most basic tool support for JML – simply parsing and typechecking – is already useful. Whereas informal comments in code are typically not kept up to date as the code is changed, the simple act of running the typechecker will catch any JML assertions referring

to parameter or field names that no longer exist, and all other typos of course. Enforcing the visibility rules can also provide useful feedback; for example, a precondition of a `public` method which refers to a `private` field of an object is suspect.

Of course, there are more exciting forms of tool support than just parsing and type-checking. In particular JML is designed to support static analysis (as in ESC/Java [Leino-etal00]), formal verification (as in the LOOP tool [Huisman01] [Jacobs-etal98]), recording of dynamically obtained invariants (as in Daikon [Ernst-etal01]), runtime assertion checking (as in JML's runtime assertion checker, jmlc [Cheon-Leavens02b] [Cheon03]), unit testing [Cheon-Leavens02], and documentation (as in JML's jmldoc tool). The paper by Burdy et al. [Burdy-etal03] is a recent survey of tools for JML. The utility of these tools is the ultimate answer to the question of what JML is good for.

1.3 Status and Plans for JML

JML is still in development. As you can see, this reference manual is still a draft, and there are some holes in it. [[[And some notes for the authors by the authors that look like this.]]]

Influences on JML that may lead to changes in its design include our desire to specify programs written using the unique features of MultiJava [Clifton-etal00], an eventual integration with Bandera [[[refs]]] or other tools for specification of concurrency, aspect-oriented programming, and the evolution of Java itself [[[refs]]]. Another influence is the ongoing effort to use JML on examples, in designing the JML tools, and efforts to give a formal semantics to JML.

1.4 Acknowledgments

The work of Leavens and Ruby was supported in part by a grant from Rockwell International Corporation and by NSF grant CCR-9503168. Work on JML by Leavens, and Ruby was also supported in part by NSF grant CCR-9803843. Work on JML by Cheon, Clifton, Leavens, Ruby, and others is supported in part by NSF grants CCR-0097907 and CCR-0113181. The work of Poll is partly supported by the Information Society Technologies (IST) Programme of the European Union, as part of the VerifiCard project, IST-2000-26328.

Thanks to Bart Jacobs, Rustan Leino, Peter Müller, Arnd Poetzsch-Heffter, and Joachim van den Berg, for many discussions about the semantics of JML specifications. Thanks for Raymie Stata for spear-heading an effort at Compaq SRC to unify JML and ESC/Java, and to Rustan and Raymie for many interesting ideas and discussions that have profoundly influenced JML. Thanks to David Cok and Todd Millstein for their work on the prototype JML checker used to check the specifications in this document. Thanks to David and Joe Kiniry for various corrections to this document.

See the "Preliminary Design of JML" [Leavens-Baker-Ruby02] for more acknowledgments relating to the earlier history, design, and implementation of JML.

2 Fundamental Concepts

[[[Visibility restrictions in JML vs. Java.]]]

[[[type vs. class and interface.]]]

[[[static vs. instance]]]

[[[Aliasing and such]]]

[[[Model of expression evaluation in assertions and what happens when exceptions occur. See [Jones95e] [Gries-Schneider95] [Leavens-Wing97a].]]]

2.1 Privacy Modifiers and Visibility

Java code that is not within an annotation uses the usual access control rules for determining visibility (or accessibility) of Java [Arnold-Gosling-Holmes00] [Gosling-etal00]. That is, a name declared in package P and type $P.T$ may be referenced from outside P only if it is declared as `public`, or if it is declared as `protected` and the reference occurs within a subclass of $P.T$. This name may be referenced from within P but outside of $P.T$ only if it is declared as `public`, default access, or `protected`. Such a name may always be referenced from within $P.T$, even if it is declared as `private`. See the Java language specification [Gosling-etal00] for details on visibility rules applied to nested and inner classes.

Within annotations, JML imposes some extra rules in addition to the usual Java visibility rules [Leavens-Baker-Ruby02]. These rules depend not just on the declaration of the name but also on the visibility level of the context that is referring to the name in question. This context can for instance be a method specification or an invariant. The visibility level of such an annotation context can be `public`, `protected`, `private`, or default (package) visibility. [[[In essence, the visibility of the referring context must intersect that of the declaration.]]]

Suppose x is a name declared in package P and type $P.T$.

- An expression in a public annotation context (e.g., in a public method specification) can refer to x only if x is declared as `public`.
- An expression in a protected annotation context (e.g., in a protected method specification) can refer to x only if x is declared as `public` or `protected`, and x must also be visible according to Java's rules (so if x is `protected`, then the reference must either be from within P or, if it is from outside P, then the reference must occur in a subclass of $P.T$).
- An expression in a default (package) visibility annotation context (e.g., in a default visibility method specification) can refer to x only if x is declared as `public`, `protected`, or with default visibility, and x must also be visible according to Java's rules (so if x has default visibility, then the reference must be from within P).
- An expression in a `private` visibility annotation context (e.g., in a private method specification) can refer to x only if x is visible according to Java's rules (so if x has private visibility, then the reference must be from within $P.T$).

In the following example, the comments on the right show which uses of the various privacy level names are legal and illegal. Similar examples could be given for method specifications, history constraints, and so on.

```
public class PrivacyDemoLegalAndIllegal {
   public int pub;
   protected int prot;
   int def;
   private int priv;

   //@ public invariant pub > 0;      // legal
   //@ public invariant prot > 0;     // illegal!
   //@ public invariant def > 0;      // illegal!
   //@ public invariant priv < 0;     // illegal!

   //@ protected invariant pub > 1;   // legal
   //@ protected invariant prot > 1;  // legal
   //@ protected invariant def > 1;   // illegal!
   //@ protected invariant priv < 1;  // illegal!

   //@ invariant pub > 1;              // legal
   //@ invariant prot > 1;             // legal
   //@ invariant def > 1;              // legal
   //@ invariant priv < 1;             // illegal!

   //@ private invariant pub > 1;      // legal
   //@ private invariant prot > 1;     // legal
   //@ private invariant def > 1;      // legal
   //@ private invariant priv < 1;     // legal
}
```

Note that in a lightweight method specification, the privacy level is assumed to be the same privacy level as the method itself. That is, for example, a protected method with a lightweight method specification uses a protected annotation context for purposes of checking proper visibility usage [Leavens-Baker-Ruby02] [Mueller02].

[[[The following is not true for the revised escjava - David]]] Note that the old version of ESC/Java [Leino-Nelson-Saxe00] does not permit one to give privacy modifiers on invariants. This conflicted with the JML rules, because an invariant with no privacy modifier in its declaration is assumed to be a default access invariant, and thus tools following the JML rule will complain if such invariants use names declared to be private. One can fix this by using a lexical trick in JML, as shown in the following (see Section 4.4 [Annotation Markers], page 14).

```
//+@ public                       // seen by JML
//@    invariant pub > 1;         // seen by ESC/Java and JML
```

The JML keyword `spec_public` provides a way to make a declaration that has different visibilities for Java and JML. For example, the following declaration declares an integer field that Java regards as private but JML regards as public.

```
private /*@ spec_public @*/ int length;
```

Thus for example, `length` in the above declaration could be used in a public method specification or invariant.

However, `spec_public` is more than just a way to change the visibility of a name for specification purposes. When applied to fields it can be considered to be shorthand for the declaration of a model field with the same name. That is, the declaration of `length` above can be thought of as equivalent to the following declarations, together with a rewrite of the Java code that uses `length` to use `_length` instead (where we assume `_x` is fresh, i.e., not used elsewhere).

```
//@ public model int length;
private int _length; //@ in length;
//@ private represents x <- _length;
```

The above desugaring allows one to change the underlying field without affecting the readers of the specification.

2.2 Model and Ghost

In JML one can declare various names with the modifier `model`; for example one can declare model types, methods, and fields. One can also declare some fields as `ghost` fields. JML also has a `model import` directive (see Chapter 5 [Compilation Units], page 20), which imports names.

The meaning of a feature declared with `model` is that it is only present for purposes of specification. For example a model type is a type that is used for specification purposes, a model method is a method that is used for specification purposes, and a model field is a field that is used for specification purposes. A model import directive imports names that are used only for specification purposes. A model field should be thought of as the abstraction of various non-model (Java) fields [Cheon-etal03]. Its value is determined by a represents clause (see Section 10.3 [Represents Clauses], page 36) and it also can act as a data group [Leino98] for specifying what non-model fields may change their state when it is used in frame axioms (see Chapter 12 [Data Groups], page 56). A model method or a model type, however, is not an abstraction of a non-model method or type.

A `ghost` field is also only present for purposes of specification. However, unlike a model field, it does not have a value determined by a represents clause, instead its value is directly determined by a *set-statement* (see Chapter 15 [Statements and Annotation Statements], page 68).

Although these model and ghost names are used only for specifications, JML uses the same namespace for such names as for normal Java names. Thus, one cannot declare a field to be both a model (or ghost) field and a normal Java field in the same class (or in a refinement, see Chapter 19 [Refinement], page 78). Similarly, a method is either a model method or not. In part, this is done because JML has no syntactic distinction between Java and JML field access or method calls. This decision makes it an error for someone to use the same name as a model or ghost feature in an implmentation. In such a case if the Java code is considered to be the goal, one can either change the name of the JML feature, or declare the JML feature to be `spec_public` instead of a model or ghost feature.

3 Syntax Notation

We use an extended BNF grammar to describe the syntax of JML. The extensions are as follows [Ledgard80].

- Nonterminal symbols are written as follows: *nonterminal*. That is, nonterminal symbols appear in an *italic* font (in the printed manual).
- Terminal symbols are written as follows: `terminal`. In a few cases it is also necessary to quote terminal symbols, such as when using '`|`' as a terminal symbol instead of a meta-symbol.
- Square brackets ([and]) surround optional text. Note that `[` and `]` are terminals.
- The notation . . . means that the preceding nonterminal or group of optional text can be repeated zero (0) or more times.

For example, the following gives a production for a non-empty list of *init-declarator*s, separated by commas.

init-declarator-list ::= *init-declarator* [`,` *init-declarator*] . . .

To remind the reader that the notation '. . .' means zero or more repetitions, we try to use '. . .' only following optional text, although in cases such as the following the brackets and the enclosed nonterminal could have been omitted.

spec-case-seq ::= *spec-case* [`also` *spec-case*] . . .

As in the above examples, we follow the C++ standard's conventions [ANSI95] in using nonterminal names of the form *X-list* to mean a comma-separated list, and nonterminal names of the form *X-seq* to mean a sequence not separated by commas.

We use "//" to start a comment (to you, the reader) in the grammar.

4 Lexical Conventions

This chapter presents the lexical conventions of JML; that is, the microsyntax of JML. At the end of the chapter, support for international character sets is described.

Throughout this chapter, grammatical productions are to be understood lexically. That is, no *white-space* (see Section 4.1 [White Space], page 13) may intervene between the characters of a token.

The microsyntax of JML is described by the production *microsyntax* below; it describes what a program looks like from the point of view of a lexical analyzer [Watt91].

microsyntax ::= *lexeme* [*lexeme*] . . .
lexeme ::= *white-space* | *lexical-pragma* | *comment*
| *annotation-marker* | *doc-comment* | *token*
token ::= *ident* | *keyword* | *special-symbol* | *java-literal*
| *informal-description*

In the rest of this section we provide more details on each of the major nonterminals used in the above grammar.

4.1 White Space

Blanks, horizontal and vertical tabs, carriage returns, formfeeds, and newlines, collectively called *white space*, are ignored except as they serve to separate tokens. Newlines and carriage returns are special in that they cannot appear in some contexts where other whitespace can appear, and are also used to end Java-style comments (see Section 4.3 [Comments], page 14).

white-space ::= *non-nl-white-space* | *end-of-line*
non-nl-white-space ::= a blank, tab, or formfeed character
end-of-line ::= *newline* | *carriage-return* | *carriage-return newline*
newline ::= a newline character
carriage-return ::= a carriage return character

4.2 Lexical Pragmas

ESC/Java [Leino-etal00] has a single kind of "lexical pragma", `nowarn`, whose syntax is described below in general terms. The JML checker currently ignores these lexical pragmas, but `nowarn` is only recognized within an annotation. Note that, unlike ESC/Java, the semicolon is mandatory. This restriction seems to be necessary to prevent lexical ambiguity.

lexical-pragma ::= *nowarn-pragma*
nowarn-pragma ::= `nowarn` [*spaces* [*nowarn-label-list*]] `;`
spaces ::= *non-nl-white-space* [*non-nl-white-space*] . . .
nowarn-label-list ::= *nowarn-label* [*spaces*] [`,` [*spaces*] *nowarn-label* [*spaces*]] . . .
nowarn-label ::= *letter* [*letter*] . . .

4.3 Comments

Both kinds of Java comments are allowed in JML: old C-style comments and new C++-style comments. However, if what looks like a comment starts with the at-sign (`@`) character, or with a plus sign and an at-sign (`+@`), then it is considered to be the start of an annotation by JML, and not a comment. Furthermore, if what looks like a comment starts with an asterisk (`*`), then it is a documentation comment, which is parsed by JML.

comment ::= *C-style-comment* | *C++-style-comment*
C-style-comment ::= `/*` [*C-style-body*] *C-style-end*
C-style-body ::= *non-at-plus-star* [*non-star-slash*] . . .
| `+` *non-at* [*non-star-slash*] . . .
| stars-non-slash [non-star-slash] . . .
non-star-slash ::= *non-star*
| *stars-non-slash*
stars-non-slash ::= `*` [`*`] . . . *non-slash*
non-at-plus-star ::= any character except `@`, `+`, or `*`
non-at ::= any character except `@`
non-star ::= any character except `*`
non-slash ::= any character except `/`
C-style-end ::= [`*`] . . . `*/`
C++-style-comment ::= `//` [`+`] *end-of-line*
| `//` *non-at-plus-end-of-line* [*non-end-of-line*] . . . *end-of-line*
| `//+` *non-at-end-of-line* [*non-end-of-line*] . . . *end-of-line*
non-end-of-line ::= any character except a newline or carriage return
non-at-plus-end-of-line ::= any character except `@`, `+`, newline, or carriage return
non-at-end-of-line ::= any character except `@`, newline, or carriage return

4.4 Annotation Markers

If what looks to Java like a comment starts with an at-sign (`@`) as its first character, then it is not considered a comment by JML. We refer to the tokens between `//@` and the following *end-of-line*, and between pairs of annotation start (`/*@` or `/*+@`) and end (`*/` or `@*/` or `@+*/`) markers as *annotations*.

Annotations must hold entire grammatical units of JML specifications, in the sense that an annotation must hold all the text that is parsed in any given production of the grammar contained in this reference manual. [[[Clarification needed.]]] For example the following is illegal, because the *postcondition* is split over two annotations, and thus each contains a fragment instead of a complete grammatical unit.

```
//@ ensures 0 <= x              // illegal!
//@       && x < a.length;
```

Implementations are not required to check for such errors. However, note that ESC/Java [Leino-Nelson-Saxe00] assumes that such units are not split into separate annotations, and so effectively does check for them.

Annotations look like comments to Java, and are thus ignored by it, but they are significant to JML. One way that this can be achieved is by having JML drop (ie., ignore)

the character sequences that are *annotation-markers*: `//@`, `//+@`, `/*@`, `/*+@`, and `@+*/`, `@*/`. The at-sign (`@`) in `@*/` is optional, and more than one at-sign may appear in the other annotation markers. However, JML will recognize *jml-keywords* only within annotations.

Within annotations, on each line, initial white-space and any immediately following at-signs (`@`) are ignored. The definition of an annotation marker is given below.

annotation-marker ::= `//@ | //+@`
`| /*@ | /*+@ | @+*/ | @*/ | */`
ignored-at-in-annotation ::= `@`

4.5 Documentation Comments

If what looks like a C-style comment starts with an asterisk (`*`) then it is a *documentation comment*. The syntax is given below. The syntax *doc-comment-ignored* is used for documentation comments that are ignored by JML.

doc-comment ::= `/**` [`*`] . . . *doc-comment-body* `*/`
doc-comment-ignored ::= *doc-comment*

At the level of the rest of the JML grammar, a documentation comment that does not contain an embedded JML method specification is essentially described by the above, and the fact that a *doc-comment-body* cannot contain the two-character sequence `*/`.

However, JML and `javadoc` both pay attention to the syntax inside of these documentation comments. This syntax is really best described by a context-free syntax that builds on a lexical syntax. However, because much of the documentation is free-form, the context-free syntax has a lexical flavor to it, and is quite line-oriented. Thus it should come as no surprise that the first non-whitespace, non-asterisk (ie., not `*`) character on a line determines its interpretation.

doc-comment-body ::= [*description*] . . .
[*tagged-paragraph*] . . .
[*jml-specs*]
description ::= *doc-non-empty-textline*
tagged-paragraph ::= *paragraph-tag* [*doc-non-nl-ws*] . . .
[*doc-atsign*] . . . [*description*] . . .
jml-specs ::= *jml-tag* [*method-specification*] *end-jml-tag*
[*jml-tag* [*method-specification*] *end-jml-tag*] . . .

The microsyntax or lexical grammar used within documentation comments is as follows. Note that the token *doc-nl-ws* can only occur at the end of a line, and is always ignored within documentation comments. Ignoring *doc-nl-ws* means that any asterisks at the beginning of the next line, even in the part that would be a JML *method-specification*, is also ignored. Otherwise the lexical syntax within a *method-specification* is as in the rest of JML. This method specification is attached to the following method or constructor declaration. (Currently there is no useful way to use such specifications in the documentation comments for other declarations.) Note the exception to the grammar of *doc-non-empty-textline*.

paragraph-tag ::= `@author` | `@deprecated` | `@exception`
| `@param` | `@return` | `@see`
| `@serial` | `@serialdata` | `@serialfield`
| `@since` | `@throws` | `@version`
| `@` *letter* [*letter*] . . .
doc-atsign ::= `@`
doc-nl-ws ::= *end-of-line* [*doc-non-nl-ws*] . . . [`*` [`*`] . . . [*doc-non-nl-ws*] . . .]
doc-non-nl-ws ::= *non-nl-white-space*
doc-non-empty-textline ::= *non-at-end-of-line* [*non-end-of-line*] . . .
jml-tag ::= `<jml>` | `<JML>` | `<esc>` | `<ESC>`
end-jml-tag ::= `</jml>` | `</JML>` | `</esc>` | `</ESC>`
A *jml-tag* marks the (temporary) end of a documentation comment and the beginning of text contributing to a method specification. The corresponding *end-jml-tag* marks the reverse transition.
The *end-jml-tag* must match the corresponding *jml-tag*.

4.6 Tokens

Character strings that are Java reserved words are made into the token for that reserved word, instead of being made into an *ident* token. Within an *annotation* this also applies to *jml-keyword*s. The details are given below.

ident ::= *letter* [*letter-or-digit*] . . .
letter ::= `_`, `$`, `a` through `z`, or `A` through `Z`
digit ::= `0` | `1` | `2` | `3` | `4` | `5` | `6` | `7` | `8` | `9`
letter-or-digit ::= *letter* | *digit*

Several strings of characters are recognized as keywords or reserved words in JML. These fall into three separate categories: Java keywords, JML predicate keywords (which start with a backslash), and JML keywords. Java keywords are truly reserved words, and are recognized in all contexts. The nonterminal *java-keywords* represents the reserved words in Java (as in the JDK version 1.4). JML keywords are only recognized as such if they occur outside of a *spec-expression* but within an annotation. JML predicate keywords are, as their name implies, used within *spec-expression*s; they are also used in *store-ref-list*s and *constrained-list*s. The details are given below.

keyword ::= *java-keyword* | *jml-predicate-keyword* | *jml-keyword*
jml-predicate-keyword ::= `\duration` | `\elemtype`
| `\everything` | `\exists`
| `\forall` | `\fresh` | `\invariant_for`
| `\is_initialized` | `\lblneg` | `\lblpos`
| `\lockset` | `\max` | `\min`
| `\nonnullelements` | `\nothing` | `\not_modified`
| `\not_specified` | `\num_of` |`\old`
| `\other` | `\private_data` | `\product`
| `\reach` | `\result` | `\space`
| `\such_that` | `\sum` | `\type`
| `\typeof` | `\TYPE` | `\working_space`

jml-keyword ::= `abrupt_behavior` | `accessible_redundantly` | `accessible`
| `also` | `assert_redundantly`
| `assignable_redundantly` | `assignable`
| `assume_redundantly` | `assume` | `axiom`
| `behavior` | `breaks_redundantly` | `breaks`
| `callable_redundantly` | `callable` | `choose_if`
| `choose` | `code_contract`
| `constraint_redundantly` | `constraint`
| `constructor` | `continues_redundantly` | `continues`
| `decreases_redundantly` | `decreases`
| `decreasing_redundantly` | `decreasing`
| `diverges_redundantly` | `diverges` | `duration_redundantly`
| `duration` | `ensures_redundantly`
| `ensures` | `example` | `exceptional_behavior`
| `exceptional_example` | `exsures_redundantly` | `exsures`
| `field` | `forall`
| `for_example` | `ghost`
| `implies_that` | `helper` | `hence_by_redundantly`
| `hence_by` | `initializer` | `initially`
| `instance` | `invariant_redundantly` | `invariant`
| `loop_invariant_redundantly` | `loop_invariant`
| `maintaining_redundantly` | `maintaining`
| `measured_by_redundantly` | `measured_by` | `method`
| `model_program` | `model` | `modifiable_redundantly`
| `modifiable` | `modifies_redundantly` | `modifies`
| `monitors_for` | `monitored` | `non_null`
| `normal_behavior` | `normal_example` | `nowarn`
| `old` | `or` | `post_redundantly` | `post`
| `pre_redundantly` | `pre`
| `pure` | `readable` | `refine`
| `represents_redundantly` | `represents` | `requires_redundantly`
| `requires` | `returns_redundantly` | `returns`
| `set` | `signals_redundantly` | `signals`
| `spec_protected` | `spec_public` | `static_initializer`
| `uninitialized` | `unreachable`
| `weakly` | `when_redundantly` | `when`
| `working_space_redundantly` | `working_space`

The following describes the special symbols used in JML. The nonterminal *java-special-symbol* is the special symbols of Java, taken without change from Java [Gosling-Joy-Steele96].

special-symbol ::= *java-special-symbol* | *jml-special-symbol*
java-special-symbol ::= *java-separator* | *java-operator*
java-separator ::= `(` | `)` | `{` | `}` | `'['` | `']'` | `;` | `,` | `.`
java-operator ::= `=` | `<` | `>` | `!` | `~` | `?` | `:`
| `==` | `<=` | `>=` | `!=` | `&&` | `'||'` | `++` | `--`
| `+` | `-` | `*` | `/` | `&` | `'|'` | `^` | `%` | `<<` | `>>` | `>>>`

| += | -= | *= | /= | &= | '|=' | ^= | %= | <<= | >>= | >>>=
jml-special-symbol ::= ==> | <== | <==> | <=!=> | -> | <- | <: | .. | '{|' | '|}'

The nonterminal *java-literal* represents Java literals which are taken without change from Java [Gosling-Joy-Steele96].

java-literal ::= *integer-literal* | *floating-point-literal* | *boolean-literal*
| *character-literal* | *string-literal* | *null-literal*

integer-literal ::= *decimal-integer-literal* | *hex-integer-literal* | *octal-integer-literal*
decimal-integer-literal ::= *decimal-numeral* [*integer-type-suffix*]
decimal-numeral ::= 0 | *non-zero-digit* [*digits*]
digits ::= *digit* [*digit*] . . .
digit ::= 0 | *non-zero-digit*
non-zero-digit ::= 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
integer-type-suffix ::= l | L
hex-integer-literal ::= *hex-numeral* [*integer-type-suffix*]
hex-numeral ::= 0x *hex-digit* [*hex-digit*] . . . | 0X *hex-digit* [*hex-digit*] . . .
hex-digit ::= *digit* | a | b | c | d | e | f
| A | B | C | D | E | F
octal-integer-literal ::= *octal-numeral* [*integer-type-suffix*]
octal-numeral ::= 0 *octal-digit* [*octal-digit*] . . .
octal-digit ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7

floating-point-literal ::= *digits* . [*digits*] [*exponent-part*] [*float-type-suffix*]
| . *digits* [*exponent-part*] [*float-type-suffix*]
| *digits exponent-part* [*float-type-suffix*]
| *digits* [*exponent-part*] *float-type-suffix*
exponent-part ::= *exponent-indicator signed-integer*
exponent-indicator ::= e | E
signed-integer ::= [*sign*] *digits*
sign ::= + | -
float-type-suffix ::= f | F | d | D

boolean-literal ::= true | false

character-literal ::= ' *single-character* ' | ' *escape-sequence* '
single-character ::= any character except ', \, carriage return, or newline
escape-sequence ::= \b // backspace
| \t // tab
| \n // newline
| \r // carriage return
| \' // single quote
| \" // double quote
| \\ // backslash
| *octal-escape*
| *unicode-escape*
octal-escape ::= \ *octal-digit* [*octal-digit*]

| \ *zero-to-three octal-digit octal-digit*
zero-to-three ::= `0` | `1` | `2` | `3`
unicode-escape ::= `\u` *hex-digit hex-digit hex-digit hex-digit*

string-literal ::= `"` [*string-character*] . . . `"`
string-character ::= *escape-sequence*
| any character except `"`, `\`, carriage return, or newline

null-literal ::= `null`

An *informal-description* looks like `(* some text *)`. It is used in predicates and store-ref expressions. The exact syntax is given below.

informal-description ::= `(*` *non-star-close* [*non-star-close*] . . . `*)`
non-star-close ::= *non-star*
| *stars-non-close*
stars-non-close ::= `*` [`*`] . . . *non-close*
non-close ::= any character except `)`

5 Compilation Units

A compilation unit in JML is similar to that in Java, with some additions. It consists of, in order,

- an optional package definition,
- an optional JML refines declaration,
- zero or more import and model import definitions,
- and one or more type definitions.

The following is the syntax of compilation units in JML. The *compilation-unit* rule is the start rule for the JML grammar.

compilation-unit ::= [*package-definition*]
[*refine-prefix*]
[*import-definition*] . . .
[*type-definition*] . . .
package-definition ::= `package` *name* `;`
import-definition ::= [`model`] `import` *name-star* `;`
name ::= *ident* [`.` *ident*] . . .
name-star ::= *ident* [`.` *ident*] . . . [`.` `*`]

5.1 Package definition

The meaning of a *package-definition* is the same as in Java [Gosling-etal00].

5.2 Refines declaration

See Chapter 19 [Refinement], page 78, for a discussion of the *refine-prefix* and its uses.

5.3 Import and model import declaration

The meaning of an *import-definition* without the `model` keyword is the same as in Java [Gosling-etal00]. When the `model` keyword is used, the import is only effective within JML annotations. An *import-definition* may have a `model` keyword modifier if and only if it is within an annotation. The syntax

`/*@ model @*/ import` *name-star* `;`

is legal, but the `model` modifier has no additional effect. [[[The above seems contradictory, though the last statement is the current behavior - DRCok]]]

5.4 Class and interface definitions

A compilation unit may have one or more type definitions. At most one of them may be declared `public`. If there is a `public` type definition, the name of that type must match the file name, as is the case for the Java implementation on the given platform. Class and interface definitions are described in Chapter 6 [Type Definitions], page 21.

6 Type Definitions

The following is the syntax of type definitions. The order of the *modifier* productions suggests the relative order for writing the modifiers in code and specifications, with `public` before all other modifiers, and `non_null` last. Note that although the *modifiers* grammar non-terminal is used in many places throughout the grammar, not all modifiers are semantically permitted in conjunction with every grammar construct. See the discussion regarding each grammar construct and in [[[Appendix ?]]] for the semantic limitations.

type-definition ::= [*doc-comment*] *modifiers class-or-interface-def*
| `;`
class-or-interface-def ::= *class-definition* | *interface-definition*
type-spec ::= *type* [*dims*] | `\TYPE` [*dims*]
type ::= *reference-type* | *builtInType*
reference-type ::= *name*
modifiers ::= [*modifier*] . . .
modifier ::= `public` | `private` | `protected`
| `spec_public` | `spec_protected`
| `abstract` | `static` |
| `model` | `ghost` | `pure`
| `final` | `synchronized`
| `instance` | `helper`
| `transient` | `volatile`
| `native` | `strictfp`
| `const`[1]
| `monitored` | `uninitialized`
| `non_null`
class-definition ::= `class` *ident* [`extends` *name* [`weakly`]]
[*implements-clause*] *class-block*
interface-definition ::= `interface` *ident* [*interface-extends*] *class-block*
interface-extends ::= `extends` *name-weakly-list*
implements-clause ::= `implements` *name-weakly-list*
name-weakly-list ::= *name* [`weakly`] [`,` *name* [`weakly`]] . . .
class-block ::= `{` [*field*] . . . `}`

The modifiers `spec_public`, `spec_protected`, `model`, `ghost`, `pure`, `instance`, `helper`, `monitored`, `uninitialized`, `non_null` are JML and not Java modifiers and are only allowed in annotation comments.

Type declarations in Java may contain

- field declarations,
- method declarations,
- constructor declarations,
- class and interface declarations, and
- static and non-static initializer blocks.

They may also contain these JML declarations:

[1] `const` is reserved but not used in Java

- ghost fields,
- model fields,
- model methods and constructors,
- model types (classes and interfaces),
- initializer and static initializer declarations, and
- several kinds of specification clauses (invariant, constraint, represents, axiom).

Each of these may be placed wherever a Java type declaration element may be placed. The specification clauses are discussed in the next chapter. The other elements are discussed here.

Type definitions that appear as elements of compilation units are called top-level types. Type definitions may also appear as elements within a type definition (nested and inner type definitions). Class definitions may also appear as statements within a statement block.

6.1 Type modifiers

The following keywords may be used as modifiers for a type definition. They are placed just before the `class` or `interface` keyword.

6.1.1 abstract

[[[discussion needed - classes only]]]

6.1.2 pure

A type definition may be modified with the JML modifier keyword `pure`. This may occur where a Java access or `static` keyword may be placed. The effect of declaring a type `pure` is that all constructor and method declarations within the type are automatically pure. It does not imply that any derived type is necessarily `pure`.

6.1.3 model

A type definition modified with the `model` keyword indicates that the type defined is only used in annotations. The entire type definition must be contained within an annotation comment, and consequently annotations within the type definition do not need to be separately enclosed in annotation comments, as is demonstrated in the example below. The scope rules for a model type definition are the same as for Java type definitions, except that a model type definition is not in scope for any Java code, only for annotations.

[[[May a model type definition appear in more than one specification file of a refinement sequence, with any member declarations being combined together? I'd prefer that it only be allowed to appear once and be required to be completely defined in one spec file - easier for tools. – DRCok]]]

Example:

```
public class Example {

private int i;

/*@
model public class ModelExample {

ghost int j;

requires true;
ensures i > 0;
public void m();

requires i > 0;
ensures i > 1;
public void inc() {
set j = 0;
++i;
}
}
@*/
}
```

[[[discussion needed]]]

6.1.4 weakly

[[[discussion needed]]]

7 Field declarations

7.1 Java Field Declarations

[[[discussion needed]]]

[[[This section ought to be moved to a section describing all the elements of a type declaration.]]] The following gives the syntax of fields.

field ::= [*doc-comment*] ... *modifiers member-decl*
| *modifiers jml-declaration*
| [*method-specification*] [`static`] *compound-statement*
| *method-specification* `static_initializer`
| *method-specification* `initializer`
| `axiom` *predicate* `;`
| `;`

member-decl ::= *variable-decls* | *method-decl*
| *class-definition* | *interface-definition*

variable-decls ::= [`field`] *type-spec variable-declarators* `;` [*jml-data-group-clause*] ...

variable-declarators ::= *variable-declarator* [`,` *variable-declarator*] ...

variable-declarator ::= *ident* [*dims*] [`=` *initializer*]

initializer ::= *expression* | *array-initializer*

array-initializer ::= `{` [*initializer-list*] `}`

initializer-list ::= *initializer* [`,` *initializer*] ... [`,`]

method-decl ::= *method-specification*
modifiers method-or-constructor-keyword
[*type-spec*] *method-head method-body*
| *method-or-constructor-keyword*
[*type-spec*] *method-head*
[*method-specification*]
method-body

method-or-constructor-keyword ::= `method` | `constructor`

method-head ::= *ident* `(` [*param-declaration-list*] `)`
[*dims*] [*throws-clause*]

method-body ::= *compound-statement* | `;`

throws-clause ::= `throws` *name* [`,` *name*] ...

param-declaration-list ::= *param-declaration* [`,` *param-declaration*] ...

param-declaration ::= [*param-modifier*] ... *type-spec ident* [*dims*]

param-modifier ::= `final` | `non_null`

In a non-Java file, such as a file with suffix '`.refines-java`' (see Chapter 19 [Refinement], page 78), one can omit the initializer of a *variable-declarator*, even one declared to be `final`. In such a file, one can also omit the body of a *method-decl*. Of course, in a '`.java`' file, one must obey all the rules of Java for declarations that are not in annotations.

[[[JML tools allow methods in Java files to omit method bodies.]]] [[[jml-var-assertions is not defined, and has changed]]]

7.2 Ghost Fields

[[[discussion needed]]]

7.3 Model Fields

[[[discussion needed]]]

7.4 JML Modifiers for Fields

7.4.1 non_null

[[[needs discussion]]]

7.4.2 monitored

[[[needs discussion]]]

7.4.3 instance

[[[needs discussion]]]

8 Methods and constructors

[[[discussion needed]]]

8.1 Java methods and constructor declarations

[[[discussion needed]]]

8.2 Model Methods and Constructors

[[[discussion needed]]]

8.3 Modifiers for Routines

[[[needs discussion]]]

8.3.1 pure

[[[needs discussion]]]

8.3.2 non_null

[[[needs discussion]]]

8.3.3 helper

[[[needs discussion]]]

8.4 Modifiers for Formal Parameters

8.4.1 non_null

[[[needs discussion]]] [[[comments on inheritance!]]]

9 Other elements of type declarations

[[[discussion needed]]]

9.1 Nested type definitions

[[[discussion needed]]]

9.2 Model Types

[[[discussion needed]]]

9.3 Initializer blocks

[[[discussion needed]]]

9.4 initializer and static_initializer declarations

[[[discussion needed - semantics unclear]]]

9.5 Type specifications

Type specifications within a type definition are JML annotations; these are discussed in Chapter 10 [Type Specifications], page 28.

10 Type Specifications

This chapter describes the way JML allows one to specify ADTs. [[[It should also start with some introductory examples...]]]

The following gives the syntax of behavioral specifications for types.

jml-declaration ::= *invariant* | *history-constraint*
| *represents-decl* | *initially-clause*
| *monitors-for-clause* | *readable-if-clause*

The various parts of type specifications and the assertions that can be added to variable declarations are described below.

[[[Discussion of modifiers allowed on type specifications.]]]

10.1 Invariants

[[[Most of the following discussion ought to be moved to a subsection of Fundamental Concepts, and then this section can be shortened to reflect the 'reference manual' content regarding invariants - namely the rules regarding their syntax and semantics. Besides a lot of the discussion applies to other forms of specs as well. – DRCok]]]

The syntax of an invariant declaration is as follows.

invariant ::= *invariant-keyword predicate* `;`
invariant-keyword ::= `invariant` | `invariant_redundantly`

An example of an invariant is given below. The invariant in the example has default (package) visibility, and says that in every state that is a visible state for an object of type `Invariant`, the object's field `b` is not null and the array it refers to has exactly 6 elements. In this example, no postcondition is necessary for the constructor since the invariant is an implicit postcondition for it.

```
package org.jmlspecs.samples.jmlrefman;

public abstract class Invariant {

    boolean[] b;

    //@ invariant b != null && b.length == 6;

    //@ assignable b;
    Invariant() {
        b = new boolean[6];
    }
}
```

Invariants are properties that have to hold in all visible states. The notion of visible state is of crucial importance in the explanation of the semantics of both invariants and constraints. A state is a *visible state* for an object o if it occurs in one of the following:

- at end of a non-helper constructor invocation that is initializing o,
- at the beginning of a non-helper destructor invocation that is finalizing o,

- at the beginning or end of a non-helper non-static method invocation with o as the receiver,
- at the beginning or end of a non-helper static method invocation for a method in o's class or some superclass of o's class, or
- when no constructor, destructor, non-static method invocation with o as receiver, or static method invocation for a method in o's class or some superclass of o's class is in progress.

Note that visible states for an object o do not include states at the beginning and end of invocations of constructors, destructors, and methods declared with the `helper` modifier.

A state is a *visible state* for a type T if it occurs after static initialization for T is complete and it is a visible state for some object that has type T.

JML distinguishes *static* and *instance* invariants. These are mutually exclusive and any invariant is either a static or instance invariant. An invariant may be explicitly declared to be static or instance by using one of the modifiers `static` or `instance` in the declaration of the invariant. An invariant declared in a class declaration is, by default, an instance invariant. An invariant declared in an interface declaration is, by default, a static invariant.

For example, the invariant declared in the class `Invariant` above is an instance invariant, because it occurs inside a class declaration. If `Invariant` had been an interface instead of a class, then this invariant would have been a static invariant.

A static invariant may only refer to static fields of an object. An instance invariant, on the other hand, may refer to both static and non-static fields.

The distinction between static and instance invariants also affects when the invariants are supposed to hold. A static invariant declared in a type T must hold in every state that is a visible state for type T. An instance invariant declared in a type T must hold for every object o of type T, for every state that is a visible state for o.

For reasoning about invariants we make a distinction between assuming, establishing, and preserving an invariant. A method or constructor *assumes* an invariant if the invariant must hold in its pre-state. A method or constructor *establishes* an invariant if the invariant must hold in its post-state. A method or constructor *preserves* an invariant if the invariant is both assumed and established.

JML's verification logic enforces invariants by making sure that each non-helper method, constructor, or destructor:

- assumes the static invariants of all types, T, for which its pre-state is a visible state for T,
- establishes the static invariants of all types, T, for which its post-state is a visible state for T,
- assumes the instance invariants of all objects, o, for which its pre-state is a visible state for o, and
- establishes the instance invariants of all objects, o, for which its post-state is a visible state for o.

This means that each non-helper constructor found in a class C preserves the static invariants of all types, including C, that have finished their static initialization, establishes the instance invariant of the object under construction, and, modulo creation and deletion

of objects, preserves the instance invariants of all other objects. (Objects that are created by a constructor must have their instance invariant established; and objects that are deleted by the action of the constructor can be assumed to satisfy their instance invariant in the constructor's pre-state.) Note in particular that, at the beginning of a constructor invocation, the instance invariant of the object being initialized does not have to hold yet.

Furthermore, each non-helper non-static method found in a type T preserves the static invariants of all types that have finish their static initialization, including T, and, modulo creation and deletion of objects, preserves the instance invariants of all objects, in particular the receiver object.

The semantics given above is highly non-modular, but is in general necessary for the enforcement of invariance when no mechanisms are available to prevent aliasing problems, or when constructs like (concrete) public fields are used [Poetzsch-Heffter97]. Of course, one would like to enforce invariants in a more modular way. By a modular enforcement of invariants, we mean that one could verify each type independently of the types that it does not use, and that a well-formed program put together from such verified types would still satisfy the semantics for invariants given above. That is, each type would be responsible for the enforcement of the invariants it declares and would be able to assume, without checking, the invariants of other types it uses.

To accomplish this ideal, it seems that some mechanism for object ownership and alias control [Noble-Vitek-Potter98] [Mueller-Poetzsch-Heffter00] [Mueller-Poetzsch-Heffter00a] [Mueller-Poetzsch-Heffter01a] [Mueller02] seems necessary. However, this mechanism is still not a part of JML, although some design work in this direction has taken place [Mueller-Poetzsch-Heffter-Leavens02].

On the other hand, people generally assume that there are no object ownership alias problems; this is perhaps a reasonable strategy for some tools, like run-time assertion checkers, to take. The alternative, tracking which types and objects are in visible states, and checking every applicable invariant for every type and object in a visible state, is obviously impractical.

Therefore, assuming or ignoring the problems with object ownership and alias control, one obtains a simple and more modular way to check invariants. This is as follows.

- Each non-helper constructor declared in a class C, must preserve the static invariant of C, if C is finished with its static initialization, and must establish the instance invariant of the object being constructed.
- Each non-helper non-static method declared in a type T, must preserve the static invariant of T, if T is finished with its static initialization, and must preserve the instance invariant of the receiver object.
- Each non-helper static method declared in a type T, must preserve the static invariant of T, if T is finished with its static initialization.

When doing such proofs, one may assume the static invariant of any type (that is finished with its static initialization), and one may also assume the instance invariant of any other object.

In this more modular style of checking invariants, one can think of all the static invariants in a class as being implicitly conjoined to the pre- and postconditions of all non-helper constructors and methods, and the instance invariants in a class as being implicitly conjoined

to the postcondition of all non-helper constructors, and to the pre- and postconditions of all non-helper methods.

As noted above, `helper` methods and constructors are exempt from the normal rules for checking invariants. That is because the beginning and end of invocations of these `helper` methods and constructors are not visible states, and therefore they do not have to preserve or establish invariants. Note that only `private` methods and constructors can be declared as `helper`.

The following subsections discuss other points about the semantics of invariants:

- Invariants can be declared `static`; see Section 10.1.1 [Static vs. instance invariants], page 31.
- Invariants can be declared with the access modifiers `public`, `protected`, and `private`, or be left with default access; see Section 10.1.3 [Access Modifiers for Invariants], page 32.
- Invariants should also hold in case a constructor or method terminates abruptly, by throwing an exception; see Section 10.1.2 [Invariants and Exceptions], page 32.
- A class inherits all visible invariants specified in its superclasses and superinterfaces; see Section 10.1.4 [Invariants and Inheritance], page 33.
- Although some aspects of invariants are discussed in isolation here, the full explanation of their semantics can only be given considered together with that of method specifications. After all, a method only has to preserve invariants when one of the preconditions (i.e., `requires` clauses) specified for that method holds. So invariants are an integral part of the explanation of method specifications in Chapter 11 [Method Specifications], page 38.
- When considering an individual method body, remember that invariants should not just hold in the beginning and the end of it, but also at any program point halfway where another (non-`helper`) method or constructor is invoked. After all, these program points are also visible states, and, as stated above, invariants should hold at all visible states.
- A method invocation on an object should not just preserve the instance invariants of that object and the static invariants of the class, but it should preserve the invariants of all other (reachable) objects as well [Poetzsch-Heffter97].

It should be noted that the last two points above are not specific to Java or JML, but these are tricky issues that have to be considered for any notion of invariant in an object-oriented languages. Indeed, these two issues make the familiar notion of invariant a lot more complicated than one might guess at first sight!

10.1.1 Static vs. instance invariants

As discussed above (see Section 10.1 [Invariants], page 28), invariants can be declared `static` or `instance`. Just like a static method, a static invariant cannot refer to the current object `this` and thus cannot refer to instance fields of `this` or non-static methods of the type.

Instance invariants must be established by the constructors of an object, and must be preserved by all non-helper instance methods. Static methods can neither assume nor establish instance invariants, because such invariants are meaningless to them.

Static invariants must be established by the static initialization of a class, and must be preserved by all non-helper constructors and methods, i.e., by both static and instance methods.

The table below summarizes this:

	static initialization	non-helper static method	non-helper constructor	non-helper instance method
static invariant	establish	preserve	preserve	preserve
instance invariant	(irrelevant)	(irrelevant)	establish	preserve

A word of warning about terminology. In standard Java terminology static members are also called class members. However, static invariants should *never* be called class invariants! This would conflict with the standard use of the term "class invariant" in the literature, where "class invariant" always means instance invariant.

10.1.2 Invariants and Exceptions

Methods and constructors should preserve and establish invariants both in the case of normal termination and in the case of abrupt termination (i.e., when an exception is thrown). In other words, invariants are implicitly included in both normal postconditions, i.e., `ensures` clauses, and in exceptional postconditions, i.e., `signals` clauses, of methods and constructors.

The requirement that invariants hold after abrupt termination of a method or constructor may seen excessively strong. However, it is the only sound option in the long run. After all, once an object's invariant is broken, no guarantees whatsoever can be made about subsequent method invocations on that object. When faced with a method or constructor that may violate an invariant in case it throws an exception, one will typically try to strengthen the precondition of the method to rule out this exceptional behavior or try to weaken the invariant. Note that a method that does not have any side effects when it throws an exception automatically preserves all invariants.

10.1.3 Access Modifiers for Invariants

Invariants can be declared with any one of the Java access modifiers `private`, `protected`, and `public`. Like class members, invariants declared in a class have `package` visibility if they do not have one of these keywords as modifier. Similarly, invariants declared in an interface implicitly have `public` visibility if they do not have one of these keywords as modifier.

The access modifier of an invariant affects which members, i.e. which fields and which (pure) methods, may be used in it, according to JML's usual visibility rules. See Section 2.1 [Privacy Modifiers and Visibility], page 9, for the details and an example using invariants.

The access modifier of invariants do *not* affect the obligations of methods and constructors to maintain and establish them. That is, *all* non-`helper` methods are expected to

preserve invariants irrespective of the access modifiers of the invariants and the methods. For example, a public method must preserve private invariants as well as public ones.

[[[JML's visibility restrictions still allow some highly dubious invariants. E.g., a private invariant can refer to a public field, which, if this public field is not final, means the invariant is not really enforceable. Tools should warn about (or forbid??) invariants which refer to non-final non-model fields that have a looser access control than the invariant itself has.]]]

10.1.4 Invariants and Inheritance

Each type inherits all the instance invariants specified in its superclasses and superinterfaces. [[[Erik wrote: "Static invariants are not inherited", but there seems to be some kind of static field inheritance in Java...]]]

The fact that (instance) invariants are inherited is one of the reasons why the use of the keyword `super` is not allowed in invariants. [[[Is this true? - I don't understand this. DRCok]]]

10.2 Constraints

History constraints [Liskov-Wing93b] [Liskov-Wing94], which we call *constraints* for short, are related to invariants. But whereas invariants are predicates that should hold in all visible states, history constraints are relationships that should hold for the combination of each visible state and any visible state that occurs later in the program's execution. Constraints can therefore be used to constrain the way that values change over time.

The syntax of history constraints in JML is as follows.

history-constraint ::= *constraint-keyword predicate*
[`for` *constrained-list*] `;`
constraint-keyword ::= `constraint` | `constraint_redundantly`
constrained-list ::= *method-name-list* | `\everything`
method-name-list ::= *method-name* [`,` *method-name*] . . .
method-name ::= *method-ref* [`(` [*param-disambig-list*] `)`]
method-ref ::= *method-ref-start* [`.` *method-ref-rest*] . . .
| `new` *reference-type*
method-ref-start ::= `super` | `this` | *ident* | `\other`
method-ref-rest ::= `this` | *ident* | `\other`
param-disambig-list ::= *param-disambig* [`,` *param-disambig*] . . .
param-disambig ::= *type-spec* [*ident* [*dims*]]

Because methods will not necessarily change the values referred to in a constraint, a constraint must always describe reflexive and transitive relations.

For example, the constraints in the example below say that the value of field `a` and the length of the array `b` will never change, and that the length of the array `c` will only ever increase.

```
package org.jmlspecs.samples.jmlrefman;

public abstract class Constraint {
```

```
    int a;
    //@ constraint a == \old(a);

    boolean[] b;

    //@ invariant b != null;
    //@ constraint b.length == \old(b.length) ;

    boolean[] c;

    //@ invariant c != null;
    //@ constraint c.length >= \old(c.length) ;

    //@ requires bLength >= 0 && cLength >= 0;
    Constraint(int bLength, int cLength) {
      b = new boolean[bLength];
      c = new boolean[cLength];
    }
}
```

Note that, unlike invariants, constraints can – and typically do – use the JML keyword `\old`.

A constraint declaration may optionally explicitly list one or more methods. It is the listed methods that must *respect* the constraint. If no methods are listed, then all non-helper methods of the class (and any subclasses) must respect the constraint. A method respects a history constraints iff the pre-state and the post-state of a non-static method invocation are in the relation specified by the history constraint (modulo static initialization). So one can think of history constraints as being implicitly included in the postcondition of relevant methods. However, history constraints do not apply to constructors and destructors, since constructors do not have a pre-state and destructors do not have a post-state.

Private methods declared as `helper` methods do not have to respect history constraints, just like these do not have to preserve invariants.

A few points to note about history constraints:

- Constraints can be declared `static`; see Section 10.2.1 [Static vs. instance constraints], page 35.
- Constraints can be declared with the access modifiers `public`, `protected`, and `private`; see Section 10.2.2 [Access Modifiers for Constraints], page 35.
- Constraints should also hold if a method terminates abruptly by throwing an exception.
- A class inherits all constraints specified in its superclasses and superinterfaces; see Section 10.2.3 [Constraints and Inheritance], page 36.
- Although some aspects of constraints are discussed in isolation here, the full explanation of their semantics can only be given considered together with that of method specifications. After all, a method only has to respect constraints when one of the preconditions (ie. `requires` clauses) specified for that method holds. So constraints

are an integral part of the explanation of method specifications in Chapter 11 [Method Specifications], page 38.

- When considering an individual method body, remember that constraints not only have to hold between the pre-state and the post-state, but between all visible state that arise during execution of the method. So, given that any program points in the method where (non-`helper`) methods or constructors are invoked are also visible states, constraints should also hold between the pre-state and any such program points, between these program points themselves, and between any such program points and the post-state.
- A method invocation on an object `o` should not just respect the constraints of `o`, but should respect the constraints of all other (reachable) objects as well.

These aspects of constraints are discussed in more detail below.

10.2.1 Static vs. instance constraints

History constraints can be declared `static`. Non-`static` constraints are also called *instance* constraints. Like a static invariant, a static history constraint cannot refer to the current object `this` or to its fields.

Static constraints should be respected by all constructors and all methods, i.e., both static and instance methods.

Instance constraints must be respected by all instance methods.

The table below summarizes this:

```
            | static          non-helper     non-helper    non-helper
            | initialization  static method  constructor   instance method
  ------------------------------------------------------------------------
  static    | (irrelevant)    respect        respect       respect
  invariant |
            |
  instance  | (irrelevant)    (irrelevant)   (irrelevant)  respect
  invariant |
```

Instance constraints are irrelevant for constructors, in that here there is no pre-state for a constructor that can be related (or not) to the post-state. However, if a visible state arises during the execution of a constructor, then any instance constraints have to be respected.

In the same way, and for the same reason, static constraints are irrelevant for static initialization.

10.2.2 Access Modifiers for Constraints

The access modifiers `public`, `private`, and `protected` pose exactly the same restrictions on constraints as they do on invariants, see Section 10.1.3 [Access Modifiers for Invariants], page 32.

10.2.3 Constraints and Inheritance

Any class inherits all the instance constraints specified in its superclasses and superinterfaces. [[[Static constraints are not inherited.]]] [[[But they still apply to subclasses, no ? - David]]]

The fact that (instance) constraints are inherited is one of the reasons why the use of the keyword `super` is not allowed in constraints. [[[Needs explanation - David]]]

10.3 Represents Clauses

The following is the syntax for `represents` clauses.

represents-decl ::= *represents-keyword store-ref-expression l-arrow-or-eq spec-expression* `;`
| *represents-keyword store-ref-expression* `\such_that` *predicate* `;`
represents-keyword ::= `represents` | `represents_redundantly`
l-arrow-or-eq ::= `<-` | `=`

The first form of `represents` clauses is called a *functional abstraction* and the second form is called a *relational abstraction*.

- The left-hand side of a `represents` clause must be a reference to a model field (See Section 7.1 [Fields], page 24, for details of model fields).
- In a functional abstraction form, the type of right-hand side of a `represents clause` must be assignment-compatible to the type of left-hand side.

[[[I would expect that a represents clause would be static iff the corresponding model variable is static. - David]]]

A `represents` clause can be declared as `static` (See Chapter 6 [Type Definitions], page 21, for `static` declarations). In a `static represents` clause, only static elements can be referenced both in the left-hand side and the right-hand side. In addition, the following restriction is enforced.

- A `static represents` clause must be declared in the type where the model field on the left-hand side is declared.

Unless explicitly declared as `static`, a `represents` clause is non-`static` (for exceptions see see Chapter 6 [Type Definitions], page 21). A non-`static represents` clause can refer to both `static` and non-`static` elements on the right-hand side.

- A non-`static represents` clause must not have a static model field in its left-hand side.
- A non-`static represents` clause must be declared in a type descended from (or nested within in) the type where the model field on the left-hand side is declared.

Note that represents clauses can be recursive. That is, a represents clause may name a field on its right hand side that is the same as the field being represented (named on the left hand side). It is the specifier's responsibilty to make sure such definitions are well-defined. But such recursive represents clauses can be useful when dealing with recursive datatypes [Mueller-Poetzsch-Heffter-Leavens02].

10.4 Initially Clauses

initially-clause ::= `initially` *predicate* `;`

10.5 Axioms

[[[description needed]]]

10.6 Readable If Clauses

readable-if-clause ::= `readable` *ident* `if` *predicate* `;`

10.7 Monitors For Clause

The *monitors-for-clause* is adapted from ESC/Java. It specifies an object and a set of objects, one of which must be locked for the first object to be accessed.

monitors-for-clause ::= `monitors_for` *ident l-arrow-or-eq spec-expression-list* `;`

11 Method Specifications

Although the use of pre- and postconditions for specification of the behavior of methods is very standard, JML offers some features that are not so standard. A good example is the distinction between normal and exceptional postconditions (in `ensures` and `signals` clauses, respectively), and the specification of frame conditions using `assignable` clauses. Another example is the use of privacy modifiers to specify for different readers, and the use of redundancy [Tan94] [Leavens-Baker99]. [[[Are we using this document as a reference manual or to motivate the design of JML as well? In any case the last sentence above needs improvement- David]]]

JML provides two constructs for specifying methods and constructors:

- pre- and postconditions, and
- model programs.

This chapter only discusses the first of these, which is by far the most common. Model programs are discussed in Chapter 17 [Model Programs], page 73.

11.1 Basic Concepts in Method Specification

[[[Discuss the "client viewpoint" here and give some basic examples here.]]]

[[[Perhaps discuss other common things to avoid repeating ourselves below...]]]

11.2 Organization of Method Specifications

The following gives the syntax of behavioral specifications for methods. We start with the top-level syntax that organizes these specifications.

method-specification ::= *specification* | *extending-specification*
extending-specification ::= `also` *specification*
specification ::= *spec-case-seq* [*redundant-spec*]
| *redundant-spec*
spec-case-seq ::= *spec-case* [`also` *spec-case*] . . .

Redundant specifications (*redundant-spec*) are discussed in Chapter 16 [Redundancy], page 72.

A *method-specification* of a method in a class or interface *must* start with the keyword `also` if this method is already declared in the parent type that the current type extends, in one of the interfaces the class implements, or in a previous file of the refinement sequence for this type. Starting a *method-specification* with the keyword `also` is intended to tell the reader that this specification is in addition to some specifications of the method that are given in the superclass of the class, one of the interfaces it implements, or in another file in the refinement sequence.

A *method-specification* can include any number of *spec-case*s, joined by the keyword `also`, as well as a *redundant-spec*. Aside from the *redundant-spec*, each of these components specifies a property that must be satisfied by the implementation of the method or constructor in question. A *method-specification* must satisfy all the specified properties

together. (So, speaking loosely, its meaning is the "conjunction" of the semantics of these individual components.)

The *spec-cases* in a *method-specification* can have several forms:

spec-case ::= *lightweight-spec-case* | *heavyweight-spec-case*
| *model-program* | *code-contract-spec*

Model programs are discussed in Chapter 17 [Model Programs], page 73. The remainder of this chapter concentrates on lightweight and heavyweight behavior specification cases. JML distinguishes between

- *heavyweight specification cases*, which start with one of the keywords `behavior`, `normal_behavior` or `exceptional_behavior` (these are also called behavior, normal behavior, and exceptional behavior specification cases, respectively), and
- *lightweight specification cases*, which do not contain one of these behavior keywords.

A lightweight specification case is similar to a behavior specification case, but with different defaults [Leavens-Baker-Ruby02]. It also is possible to desugar all such specification cases into behavior specification cases [Raghavan-Leavens00].

11.3 Access Control in Specification Cases

Heavyweight specification cases may be declared with an explicit access modifier, according to the following syntax.

privacy ::= `public` | `protected` | `private`

The access modifier of a heavyweight specification case cannot allow more access than the method being specified. So a `public` method may have a `private` behavior specification, but a `private` method may not have a `public` public specification. A heavyweight specification case without an explicit access modifier is considered to have default (package) access.

Lightweight specification cases have no way to explicitly specify an access modifier, so their access modifier is implicitly the same as the method being specified. For example, a lightweight specification of a `public` method has `public` access, implicitly, but a lightweight specification of a `private` method has `private` access, implicitly. Note that this is a different default than that for heavyweight specifications, where an omitted access modifier always means package access.

The access modifier of a specification case affects only which annotations are visible in the specification and does *not* affect the semantics of a specification case in any other way.

JML's usual visibility rules apply to specification cases. So, for example, a public specification case may only refer to public members, a protected specification case may refer to both public and protected members, as long as the protected members are otherwise accessible according to Java's rules, etc. See Section 2.1 [Privacy Modifiers and Visibility], page 9, for more details and examples.

11.4 Lightweight Specification Cases

Syntax

The following is the syntax of lightweight specification cases. These are the most concise specification cases.

lightweight-spec-case ::= *generic-spec-case*
generic-spec-case ::= [*spec-var-decls*] [*spec-header*] *simple-spec-body*
| [*spec-var-decls*] [*spec-header*] `{|` *generic-spec-case-seq* `|}`
generic-spec-case-seq ::= *generic-spec-case* [`also` *generic-spec-case*] . . .
spec-header ::= *requires-clause*
simple-spec-body ::= *simple-spec-body-clause* [*simple-spec-body-clause*] . . .
simple-spec-body-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause*
| *duration-clause* | *ensures-clause* | *signals-clause*

As far as the syntax is concerned, the only difference between a lightweight specification cases and a *behavior-specification-case* (see Section 11.6 [Behavior Specification Cases], page 41) is that the latter has the keyword `behavior` and possibly an access control modifier.

A lightweight specification case always has the same access modifier as the method being specified, see Section 11.3 [Access Control in Specification Cases], page 39. To specify a different access control modifier, one must use a heavyweight specification.

Semantics

A lightweight specification case can be understood as syntactic sugar for a behavior specification case, except that the defaults for omitted specification clauses are different for lightweight specification cases than for behavior specification cases. So, for example, apart from the class names, method `m` in class `Lightweight` below

```
package org.jmlspecs.samples.jmlrefman;

public abstract class Lightweight {

    protected boolean P, Q, R;
    protected int X;

    /*@ requires P;
      @  assignable X;
      @  ensures Q;
      @  signals (Exception) R;
      @*/
    protected abstract int m() throws Exception;
}
```

has a specification that is equivalent to that of method `m` in class `Heavyweight` below.

```
package org.jmlspecs.samples.jmlrefman;

public abstract class Heavyweight {
```

```
    protected boolean P, Q, R;
    protected int X;

    /*@ protected behavior
      @   requires P;
      @   diverges \not_specified;
      @   assignable X;
      @   when \not_specified;
      @   working_space \not_specified;
      @   duration \not_specified;
      @   ensures Q;
      @   signals (Exception) R;
      @*/
    protected abstract int m() throws Exception;
}
```

As this example illustrates, the default for an omitted clause in a lightweight specification is `\not_specified` for every clause [Leavens-Baker-Ruby02]. It is intended that the meaning of `\not_specified` may vary between different uses of a JML specification. For example, a static checker might treat a `requires` clause that is `\not_specified` as if it were `true`, while a verification logic would need to treat it as if it were `false`.

In JML, a completely omitted specification is taken to be a lightweight specification.

11.5 Heavyweight Specification Cases

There are three kinds of heavyweight specification cases, called behavior, normal behavior, and exceptional behavior specification cases, beginning (after an optional privacy modifier) with the one of the keywords `behavior`, `normal_behavior`, or `exceptional_behavior`, respectively.

heavyweight-spec-case ::= *behavior-spec-case*
| *exceptional-behavior-spec-case*
| *normal-behavior-spec-case*

Like lightweight specification cases, normal behavior and exceptional behavior specification cases can be understood as syntactic sugar for special kinds of `behavior` specification cases [Raghavan-Leavens00].

11.6 Behavior Specification Cases

The behavior specification case is the most general form of specification case. All other forms of specification cases simply provide some syntactic sugar for special kinds of `behavior` specification cases.

Syntax

As far as the syntax is concerned, the only difference between a `behavior` specification case and a lightweight one is the keyword `behavior`, and the optional access control modifier.

behavior-spec-case ::= [*privacy*] `behavior` *generic-spec-case*

Semantics

To explain the semantics of a behavior specification case we make a distinction between *flat* and *nested* specification cases:

- *Flat* specification cases are of the form

  `behavior`
  [*spec-var-decls*] [*spec-header*] *simple-spec-body*

  i.e., of the form

  `behavior`
  [*spec-var-decls*] [*spec-header*]
  simple-spec-body-clause [*simple-spec-body-clause*] . . .

  A flat specification case is just made up of a sequence of method specification clauses, ie. `require`, `ensures`, etc. clauses, and its semantics is explained directly in Section 11.6.1 [Semantics of flat behavior specification cases], page 42.

- *Nested* specification cases are all other specification cases. They use the special brackets `{|` and `|}` to nest specification clauses and possibly also `also` inside these brackets to join several specification cases.

  A nested specification case can be syntactically desugared into a list of one or more simple specification cases, joined by the `also` keyword [Raghavan-Leavens00]. This is explained in Section 11.6.5 [Semantics of nested behavior specification cases], page 44.

Invariants and constraints

The semantics of a behavior specification case for a method or constructor in a class depends on the invariants and constraints that have been specified. This is discussed in Section 10.1 [Invariants], page 28 and Section 10.2 [Constraints], page 33. In a nutshell, methods must to preserve invariants and respect constraints, and constructors must establish invariants.

11.6.1 Semantics of flat behavior specification cases

Below we explain the semantics of a simple behavior specification case with precisely one `requires` clause, one `ensures` clause, one `when` clause one `assignable` clause one `accessible` clause, and at least one `diverges` clause.

A `behavior` specification case can contain any number of these clauses, but, as explained in Section 11.9 [Method Specification Clauses], page 47, any `behavior` specification case is equivalent with a `behavior` specification case of this form.

11.6.2 Non-helper methods

The semantics of a specification

```
behavior
  requires P;
  diverges D;
  assignable A;
  when W;
  ensures Q;
  signals (E1 e1) R1;
  ...
  signals (En en) Rn;
also
code_contract
  accessible C;
  callable p();
```

for a non-`helper` instance method `m` is as follows. [[[What about duration and working_space clauses? - David]]]

If the method is invoked in a pre-state where

- the precondition P holds, and
- all applicable invariants hold

then either:

- the Java virtual machine throws an error that inherits from `java.lang.Error`, or
- if the execution of the method does not terminate (i.e., it loops forever or exits without returning or throwing an exception), the predicate D holds in the pre-state, or
- the method terminates by returning or throwing an exception, and:

  during execution of the method (which includes all called methods and constructors), only locations that either did not exist in the pre-state, that are local to the method (including the method's formal parameters), or that are either named by the assignable clause's list A, or are dependees (see Chapter 12 [Data Groups], page 56) of such locations, are assigned to by the method, and

  in all visible states, all applicable invariants and history constraints hold, as explained in detail in Section 10.1 [Invariants], page 28 and see Section 10.2 [Constraints], page 33, and

  if the execution of the method terminates normally, then the normal postcondition Q holds, as do all applicable invariants and history constraints, and

  if the execution of the method terminates throwing an exception of type Ei, then the exceptional postcondition Ri holds, with the exception object thrown substituted for ei, as do all invariants and constraints, and

  in the body of the method, only the locations mentioned in the accessible clause list C are directly accessed, and

  in the body of the method, only the methods mentioned in the callable clause list p are directly called.

Note that if there is more than one `signals` clause, and the types *Ei* are related in the subclass hierarchy, then more than one of the exceptional postconditions *Ri* may apply if an exception is thrown. This is explained in detail in See Section 11.9.4 [Signals Clauses], page 49.

If the formal parameters of the method are used in a (normal or exceptional) postcondition, then these always take the value the parameters had in the pre-state, and not the value they have in the post-state, as explained in See Section 11.9.5 [Parameters in Postconditions], page 50.

11.6.3 Non-helper constructors

The semantics of a flat specification case for a (non-`helper`) constructor is the same as that for a (non-`helper`) method given above, except that:

any instance invariants of the object being initialized by the constructor are not assumed to hold in the precondition,

any instance constraints do not have to be established as implicit part of the postcondition of the constructor.

These two differences are also discussed in Section 10.1 [Invariants], page 28 and Section 10.2 [Constraints], page 33.

11.6.4 Helper methods and constructors

The semantics of a flat specification case for a helper method (or constructor) is the same as that for a non-helper method (or constructor) given above, except that:

- the instance invariants for the current object and the static invariants for the current class are not assumed to hold in the pre-state, and do not have to be established in the post-state.
- the instance constraints for current object and the static constraints for the current class do not have to be established in the post-state

These differences are also discussed in Section 10.1 [Invariants], page 28 and Section 10.2 [Constraints], page 33.

11.6.5 Semantics of nested behavior specifications

We now explain how all behavior specification cases can be desugared into a list of one or more flat specification cases joined by the `also` keyword [Raghavan-Leavens00]. The semantics of a behavior specification case is then simply the semantics of this desugared version. The meaning of a such a list of specification cases is explained in Section 11.2 [Organization of Method Specifications], page 38, the meaning of a single simple specification case is explained in Section 11.6.1 [Semantics of flat behavior specification cases], page 42.

```
[ spec-var-decls ] [ spec-header ]
{|  GenSpecCase1
    also
     ...
```

```
            also
             GenSpecCase1
        |}
```

can be desugared into

```
        [ spec-var-decls ] [ spec-header ]
          GenSpecCase1
      also
          ...
      also
        [ spec-var-decls ] [ spec-header ]
          GenSpecCasen
```

[[[EXAMPLE]]]

11.7 Normal Behavior Specifications Cases

A `normal_behavior` specification case is just syntactic sugar for a `behavior` specification case with an implicit `signals` clause

```
signals (java.lang.Exception) false;
```

ruling out abrupt termination, ie. the throwing of any exception.

The following gives the syntax of the body of a normal behavior specification case.

normal-behavior-spec-case ::= [*privacy*] `normal_behavior` *normal-spec-case*
normal-spec-case ::= [*spec-var-decls*] *normal-spec-body*
| [*spec-var-decls*] [*spec-header*] `{|` *normal-spec-case-seq* `|}`
normal-spec-case-seq ::= *normal-spec-case* [`also` *normal-spec-case*] . . .
normal-spec-body ::= *normal-spec-clause* [*normal-spec-clause*] . . .
normal-spec-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause*
| *duration-clause* | *ensures-clause*

As far as syntax is concerned, the only difference with the base behavior specification case is that normal behavior specification cases use a different behavior keyword and cannot include *signals-clause*s.

The semantics of a normal behavior specification case is the same as the `behavior` specification case obtained by adding the following *signals-clause*

```
signals (java.lang.Exception) false;
```

So a normal behavior specification case specifies a precondition which guarantees normal terminates; i.e., it prohibits the method from throwing an exception.

11.8 Exceptional Behavior Specifications Cases

The following gives the syntax of the body of an exceptional behavior specification case.

exceptional-behavior-spec-case ::= [*privacy*] `exceptional_behavior` *exceptional-spec-case*

exceptional-spec-case ::= [*spec-var-decls*] *exceptional-spec-body*
| [*spec-var-decls*] [*spec-header*] `{|` *exceptional-spec-case-seq* `|}`
exceptional-spec-case-seq ::= *exceptional-spec-case* [`also` *exceptional-spec-case*] . . .
exceptional-spec-body ::= *exceptional-spec-clause* [*exceptional-spec-clause*] . . .
exceptional-spec-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause*
| *duration-clause* | *signals-clause*

As far as syntax is concerned, the only difference from a standard behavior specification case is that an exceptional behavior specification case uses a different behavior keyword and cannot include an `ensures` clause.

The semantics of a exceptional behavior specification case is the same as the behavior specification case obtained by adding the following `ensures` clause.

```
ensures false;
```

So an exceptional behavior specification case specifies a precondition which guarantees that the method throws an exception, i.e., a precondition which prohibits the method from terminating normally.

11.8.1 Pragmatics of Exceptional Behavior Specifications Cases

Note that an exceptional behavior specification case says that an exception *must* be thrown if its precondition is met (assuming the diverges clause predicate is `false`, as is the default.) Beware of the difference between specifying that an exception *must* be thrown and specifying that an exception *may* be thrown. To specify that an exception *may* be thrown you should *not* use an exceptional behavior, but should instead use a behavior specification case [Leavens-Baker-Ruby02].

For example, the following method specification

```
package org.jmlspecs.samples.jmlrefman;

public abstract class InconsistentMethodSpec {

    /** A specification that can't be satisfied. */
    /*@  public normal_behavior
     @    requires z <= 99;
     @    assignable \nothing;
     @    ensures \result > z;
     @ also
     @  public exceptional_behavior
     @    requires z < 0;
     @    assignable \nothing;
     @    signals (IllegalArgumentException) true;
     @*/
    public abstract int cantBeSatisfied(int z)
        throws IllegalArgumentException;
}
```

is *inconsistent* because the preconditions `z <= 99` and `z < 0` overlap, for example when `z` is `-1`. When both preconditions hold then the exceptional behavior case specifies that an exception *must* be thrown and the normal behavior case specifies that an exception *may not* be thrown, but the implementation cannot both throw and not throw an exception.

Similarly, multiple exceptional specification cases with overlapping preconditions may give rise to an inconsistent specification. [[[For example,]]]

This specification is inconsistent (ie. it is impossible to come up with an implementation that meets this specification), because if [[[..]]] and [[[...]]] then the specification requires that two different exception are thrown, which is clearly impossible.

11.9 Method Specification Clauses

The different kinds of clauses that can be used in method specifications are discussed in this section. See Section 11.4 [Lightweight Specification Cases], page 39, for the overall syntax that ties these clauses together.

11.9.1 Specification Variable Declarations

The syntax of *spec-var-decls* is as follows.

spec-var-decls ::= *forall-var-decls* [*old-var-decls*]
| *old-var-decls*
forall-var-decls ::= *forall-var-decl* [*forall-var-decl*] . . .
forall-var-decl ::= `forall` *quantified-var-decl* `;`
old-var-decls ::= `old-var-decl` [*old-var-decl*] . . .
old-var-decl ::= `old` *type-spec spec-variable-declarators* `;`

11.9.2 Requires Clauses

A requires clause specifies a precondition of method or constructor. Its syntax is as follows.

requires-clause ::= *requires-keyword pred-or-not* `;`
requires-keyword ::= `requires` | `pre`
| `requires_redundantly` | `pre_redundantly`
pred-or-not ::= *predicate* | `\not_specified`

The *predicate* in a `requires` clause can refer to any visible fields and to the parameters of the method. See Section 2.1 [Privacy Modifiers and Visibility], page 9, for more details on visibility in JML.

Any number of requires clauses can be included a single specification case. Multiple requires clauses in a specification case mean the same as a single requires clause whose precondition predicate is the *conjunction* of these precondition predicates in the given requires clauses. For example,

```
requires P;
requires Q;
```

means the same thing as:

```
requires P && Q;
```

[[[How do we deal with `\not_specified` for this conjunction semantics?]]] [[[Also explain the meaning if P or Q throws an exception. Is the above equivalence true if Q is defined only if P is true (i.e. making use of the short-circuit nature of &&)?]]]

When a requires clause is omitted in a specification case, a default requires clause is used. For a lightweight specification case, the default precondition is `\not_specified`. The default precondition for a heavyweight specification case is `true`.

11.9.3 Ensures Clauses

An ensures clause specifies a normal postcondition, i.e., a property that is guaranteed to hold at the end of the method (or constructor) invocation in the case that this method (or constructor) invocation returns without throwing an exception. The syntax is as follows See Section 11.9.2 [Requires Clauses], page 47, for the syntax of *pred-or-not*.

ensures-clause ::= *ensures-keyword pred-or-not* `;`
ensures-keyword ::= `ensures` | `post`
| `ensures_redundantly` | `post_redundantly`

A *predicate* in an `ensures` clause can refer to any visible fields, the parameters of the method, `\result` if the method is non-void, and may contain expressions of the from `\old(E)`. See Section 2.1 [Privacy Modifiers and Visibility], page 9, for more details on visibility in JML.

Informally,

```
ensures Q;
```

means

> if the method invocation terminates normally (ie. without throwing an exception), then predicate Q holds in the post-state.

In an ensures clause, `\result` stands for the result that is returned by the method. The postcondition Q may contain expressions of the form `\old(e)`. Such expressions are evaluated in the pre-state, and not in the post-state, and allow Q to express a relation between the pre- and the post-state. If parameters of the method occur in the postcondition Q, these are always evaluated in the pre-state, not the post-state. In other words, if a method parameter x occurs in Q, it is treated as `\old(`x`)`. For a detailed explanation of this see Section 11.9.5 [Parameters in Postconditions], page 50.

Any number of ensures clauses can be given in a single specification case. Multiple ensures clauses in a specification case mean the same as a single ensures clause whose postcondition predicate is the *conjunction* of the postcondition predicates in the given ensures clauses. So

```
ensures P;
ensures Q;
```

means the same as

```
ensures P && Q;
```

[[[Also explain the meaning if P or Q throws an exception. Is the above equivalence true if Q is defined only if P is true (i.e. making use of the short-circuit nature of &&)?]]]

When an ensures clause is omitted in a specification case, a default ensures clause is used. For a lightweight specification case, the default precondition is `\not_specified`. The default precondition for a heavyweight specification case is `true`.

11.9.4 Signals Clauses

In a specification case a `signals` clause specifies the exceptional or abnormal postcondition, i.e., the property that is guaranteed to hold at the end of a method (or constructor) invocation when this method (or constructor) invocation terminates abruptly by throwing a given exception.

The syntax is as follows See Section 11.9.2 [Requires Clauses], page 47, for the syntax of *pred-or-not*.

signals-clause ::= *signals-keyword*
`(` *reference-type* [*ident*] `)` [*pred-or-not*] `;`
signals-keyword ::= `signals | signals_redundantly`
`| exsures | exsures_redundantly`

In a *signals-clause* of the form

```
signals (E e) P;
```

E has to be a subclass of `java.lang.Exception`, and the variable *e* is bound in *P*. If *E* is a checked exception (i.e., if it does not inherit from `java.lang.RuntimeException` [Arnold-Gosling-Holmes00] [Gosling-etal00]), it must either be one of the exceptions listed in the method or constructor's `throws` clause, or a subclass or a superclass of such a declared exception.

Informally,

```
signals (E e) P;
```

means

> If the method (or constructor) invocation terminates abruptly by throwing an exception of type *E*, then predicate *P* holds in the final state for this exception object *E*.

A signals clause of the form

```
signals (E e) R;
```

is equivalent to the signals clause

```
signals (java.lang.Exception e) (e instanceof E) ==> R;
```

Several signals clauses can be given in a single lightweight, behavior or exceptional behavior specification case. Multiple signals clauses in a specification case mean the same as a single signals clause whose exceptional postcondition predicate is the *conjunction* of the exceptional postcondition predicates in the given signals clauses. This should be understood to take place after the desugaring given above, which makes all the signals clauses refer to exceptions of type `java.lang.Exception`. Also, the names in the given signals clauses have to be standardized [Raghavan-Leavens00]. So for example,

```
signals (E1 e) R1;
signals (E2 e) R2;
```

means the same as

```
signals (Exception e)   ((e instanceof E1) ==> R1)
                      && ((e instanceof E2) ==> R2);
```

Note that this means that if an exception is thrown that is both of type *E1* and of type *E2*, then both *R1* and *R2* must hold.

[[[EXAMPLE]]]

Beware that a `signals` clause specifies when a certain exception *may* be thrown, not when a certain exception *must* be thrown. To say that an exception must be thrown in some situation, one has to exclude that situation from other signals clauses and from ensures clause (and any diverges clauses).

[[[EXAMPLE?]]]

11.9.5 Parameters in Postconditions

Parameters of methods are passed by value in Java, meaning that parameters are local variables in a method body, which are initialized when the method is called with the values of the parameters for the invocation.

This leads us to the following two rules:

- The parameters of a method or constructor can never be listed in the its assignable clause.
- If parameters of a method (or constructor) are used in a normal or exceptional postcondition for that method (or constructor), i.e., in an ensures or signals clause, then these always have their value in the pre-state of the method (or constructor), not the post-state. In other word, there is an implicit `\old()` placed around any occurrence of a formal parameter in a postcondition.

The justification for the first convention is that clients cannot observe assignments to the parameters anyway, as these are local variables that can only be used by the implementation of the method. Given that clients can never observe these assignments, there is no point in making them part of the contract between a class and its clients.

The justification for the second convention is that clients only know the initial values of the parameter that they supply, and do not have any knowledge of the final values that these variables may have in the post-state.

The reason for this is best illustrated by an example. Consider the following class and its method specifications. Without the convention described above the implementations given for methods `notCorrect1` and `notCorrect2` would satisfy their specifications. However, clearly neither of these satisfies the specification when read from the caller's point of view.

```
package org.jmlspecs.samples.jmlrefman;

public abstract class ImplicitOld {

    /*@ ensures 0 <= \result && \result <= x;
      @ signals (Exception) x < 0;
      @*/
    public static int notCorrect1(int x) throws Exception {
        x = 5;
```

```
        return 4;
    }

    /*@ ensures 0 <= \result && \result <= x;
      @ signals (Exception) x < 0;
      @*/
    public static int notCorrect2(int x) throws Exception {
        x = -1;
        throw new Exception();
    }

    /*@ ensures 0 <= \result && \result <= x;
      @ signals (Exception) x < 0;
      @*/
    public static int correct(int x) throws Exception {
        if (x < 0) {
            throw new Exception();
        } else {
            return 0;
        }
    }
}
```

The convention above rules out such pathological implementations as `notCorrect1` above; because mention of a formal parameter name, such as `x` above, in postconditions always means the pre-state value of that name, e.g., `\old(x)` in the example above.

11.9.6 Diverges Clauses

The diverges clause is a seldom-used feature of JML. It says when a method may loop forever or otherwise not return to its caller, by either throwing an exception or returning normally. The syntax is as follows See Section 11.9.2 [Requires Clauses], page 47, for the syntax of *pred-or-not*.

diverges-clause ::= *diverges-keyword pred-or-not* `;`
diverges-keyword ::= `diverges` | `diverges_redundantly`

When a diverges clause is omitted in a specification case, a default diverges clause is used. For a lightweight specification case, the default diverges condition is `\not_specified`. The default diverges condition for a heavyweight specification case is `false`. Thus by default, heavyweight method specification cases are total correctness specifications [Dijkstra76]. Explicitly writing a diverges clause allows one to obtain a partial correctness specification [Hoare69]. Being able to specify both total and partial correctness specification cases for a method leads to additional power [Hesselink92] [Nelson89].

As an example of the use of `diverges`, consider the `exit` method in the following class. (This example is simplified from the specification of Java's `System.exit` method. This specification says that the method can always be called (the implicit precondition is `true`), may always not return to the caller (i.e., diverge), and may never return normally, and

may never throw an exception. Thus the only thing the method can legally do, aside from causing a JVM error, is to not return to its caller.

```
package org.jmlspecs.samples.jmlrefman;

public abstract class Diverges {

    /*@ public behavior
      @    diverges true;
      @    assignable \nothing;
      @    ensures false;
      @    signals (Exception) false;
      @*/
    public static void abort();

}
```

The diverges clause is also useful to specify things like methods that are supposed to abort the program when certain conditions occur, although that isn't really good practice in Java. In general, it is most useful for examples like the one given above, when you want to say when a method cannot return to its caller.

11.9.7 When Clauses

The `when` clause allows concurrency aspects of a method or constructor to be specified [Lerner91] [Sivaprasad95]. A caller of a method will be delayed until the condition given in the `when` clause holds. (Note that support for concurrency in JML is in its infancy.)

The syntax is as follows See Section 11.9.2 [Requires Clauses], page 47, for the syntax of *pred-or-not*.

when-clause ::= *when-keyword pred-or-not* `;`
when-keyword ::= `when` | `when_redundantly`

When a when clause is omitted in a specification case, a default when clause is used. For a lightweight specification case, the default when condition is `\not_specified`. The default when condition for a heavyweight specification case is `true`.

[[[Need an example of a when clause and how it is used.]]]

11.9.8 Assignable Clauses

An assignable clause gives a frame axiom for a specification. It says that, from the client's point of view, only the locations named (and their dependees) can be assigned to during the execution of the method. However, locations that are local to the method (or methods it calls) and locations that are created during the method's execution are not subject to this restriction.

The syntax is as follows. See Section 13.3 [Store Refs], page 66, for the syntax of *store-ref*.

assignable-clause ::= *assignable-keyword conditional-store-ref-list* `;`
assignable-keyword ::= `assignable` | `assignable_redundantly`
| `modifiable` | `modifiable_redundantly`

| `modifies` | `modifies_redundantly`
conditional-store-ref-list ::= *conditional-store-ref*
[`,` *conditional-store-ref*] . . .
conditional-store-ref ::= *store-ref* [`if` *predicate*]
| *other-ref* [`if` *predicate*]
other-ref ::= `\other` [*store-ref-name-suffix*] . . .

When an assignable clause is omitted in a specification case, a default assignable clause is used. This default has a default *conditional-store-ref-list*. For a lightweight specification case, the default *conditional-store-ref-list* is `\not_specified`. The default *conditional-store-ref-list* for a heavyweight specification case is `\everything`.

If one wants the opposite of the default for a heavyweight specification case, one can specify that a method cannot assign to any locations by writing:

```
assignable \nothing;
```

Using the modifier `pure` on a method achieves the same effect as specifying `assignable \nothing`, but does so for the method's entire specification as opposed to a single *specification-case*.

11.9.9 Working Space Clauses

A *working-space-clause* can be used to specify the maximum amount of heap space used by a method, over an above that used by its callers. The clause applies only to the particular specification case it is in, of course This is a adapted from the work of Krone, Ogden, and Sitaraman on RESOLVE [Krone-Ogden-Sitaraman03].

working-space-clause ::= *working-space-keyword* `\not_specified ;`
| *working-space-keyword spec-expression* [`if` *predicate*] `;`
working-space-keyword ::= `working_space` | `working_space_redundantly`

The *spec-expression* in a working space clause must have type `int`. It is to be understood in units of bytes.

The *spec-expression* in a working space clause may use `\old` and other JML operators appropriate for postconditions. This is because it is considered to be evaluated in the post-state, and provides a guarantee of the maximum amount of additional space used by the call. In some cases this space may depend on the `\result`, exceptions thrown, or other post-state values.

In both lightweight and heavyweight specification cases, an omitted working space clause means the same as a working space clause of the following form.

```
working_space \not_specified;
```

See Section 13.1.7 [Backslash working space], page 59, for information about the `\working_space` expression that can be used to describe the working space needed by a method call. See Section 13.1.6 [Backslash space], page 59, for information about the `\space` expression that can be used to describe the heap space occupied by an object.

11.9.10 Duration Clauses

A duration clause can be used to specify the maximum (i.e., worst case) time needed to process a method call in a particular specification case. [[[Tools are simpler if the

argument can simply be an arbitrary expression rather than a method call. – DRCok]]] This is a adapted from the work of Krone, Ogden, and Sitaraman on RESOLVE [Krone-Ogden-Sitaraman03].

duration-clause ::= *duration-keyword* `\not_specified ;`
| *duration-keyword spec-expression* [`if` *predicate*] `;`
duration-keyword ::= `duration | duration_redundantly`

The *spec-expression* in a duration clause must have type `long`. It is to be understood in units of [[[the JVM instruction that takes the least time to execute, which may be thought of as the JVM's cycle time.]]] The time it takes the JVM to execute such an instruction can be multiplied by the number of such cycles to arrive at the clock time needed to execute the method in the given specification case. [[[This time should also be understood as not counting garbage collection time.]]]

The *spec-expression* in a duration clause may use `\old` and other JML operators appropriate for postconditions. This is because it is considered to be evaluated in the post-state, and provides a guarantee of the maximum amount of additional space used by the call. In some cases this space may depend on the `\result`, exceptions thrown, or other post-state values. [[[But if the method throws an exception, then \result is not valid. - DRCok]]]

In both lightweight and heavyweight specification cases, an omitted duration clause means the same as a duration clause of the following form.

```
duration \not_specified;
```

See Section 13.1.5 [Backslash duration], page 59, for information about the `\duration` expression that can be used in the duration clause to specify the duration of other methods.

11.9.11 Measured By Clauses

[[[No longer a clause in this section.]]]

A measured by clause can be used in a termination argument for a recursive specification.

measured-clause ::= *measured-by-keyword* `\not_specified ;`
| *measured-by-keyword spec-expression* [`if` *predicate*] `;`
measured-by-keyword ::= `measured_by | measured_by_redundantly`

The *spec-expression* in a measured by clause must have type `int`.

In both lightweight and heavyweight specification cases, an omitted measured by clause means the same as a measured by clause of the following form.

```
measured_by \not_specified;
```

11.9.12 Accessible Clauses

accessible-clause ::= *accessible-keyword conditional-store-ref-list* `;`
accessible-keyword ::= `accessible | accessible_redundantly`

When an accessible clause is omitted in a code contract specification case, a default accessible clause is used. This default has a default *object-ref-list* which is `\everything`.

[[[Need some discussion of the meaning of an accessible clause.]]]

11.9.13 Callable Clauses

callable-clause ::= *callable-keyword* *callable-methods-list* `;` *callable-keyword* ::= `callable` | `callable_redundantly` *callable-methods-list* ::= *method-name-list* | *store-ref-keyword*

When a callable clause is omitted in a code contract specification case, a default callable clause is used. This default has a default *callable-methods-list* which is `\everything`.

[[[Need some discussion of the meaning of an callable clause.]]]

12 Frame Conditions and Data Groups

[[[discussion needed]]]

12.1 Data Groups

[[[needs discussion - general usage, default data groups, use in modifies statements]]]

The following is the syntax for `in` and `maps` data group clauses.

jml-data-group-clause ::= *in-group-clause* | *maps-into-clause*
in-group-clause ::= *in-keyword group-list* ;
in-keyword ::= `in` | `in_redundantly`
group-list ::= *group-name* [, *group-name*] . . .
group-name ::= [*group-name-prefix*] *ident*
group-name-prefix ::= `super .` | `this .`
maps-into-clause ::= *maps-keyword member-field-ref* `\into` *group-list* ;
maps-keyword ::= `maps` | `maps_redundantly`
member-field-ref ::= *ident* . *maps-member-ref-expr*
| *maps-array-ref-expr* [. *maps-member-ref-expr*]
maps-member-ref-expr ::= *ident* | `*`
maps-array-ref-expr ::= *ident maps-spec-array-dim* [*maps-spec-array-dim*] . . .
maps-spec-array-dim ::= '`[`' *spec-array-ref-expr* '`]`'

See [Leino98].

12.2 Static Data Group Inclusions

[[[`in`, `in_redundantly`, needs discussion]]]

12.3 Dynamic Data Group Mappings

[[[`maps`, `maps_redundantly`, `\into` needs discussion]]]

13 Predicates and Specification Expressions

13.1 Predicates

The following gives the syntax of predicates and specification expressions. Within a *spec-expression*, one cannot use any of the operators (such as `++`, `--`, and the assignment operators) that would necessarily cause side effects. See Section 13.2 [Specification Expressions], page 64, for the syntax of expressions.

predicate ::= *spec-expression*
spec-expression-list ::= *spec-expression* [`,` *spec-expression*] . . .
spec-expression ::= *expression*

jml-primary ::= `\result`
| `\old` `(` *spec-expression* `)`
| `\not_modified` `(` *store-ref-list* `)`
| `\fresh` `(` *spec-expression-list* `)`
| `\reach` `(` *spec-expression* `)`
| `\duration` `(` *expression* `)`
| `\space` `(` *spec-expression* `)`
| `\max` `(` *spec-expression* `)`
| `\working_space` `(` *expression* `)`
| *informal-description*
| `\nonnullelements` `(` *spec-expression* `)`
| `\typeof` `(` *spec-expression* `)`
| `\elemtype` `(` *spec-expression* `)`
| `\type` `(` *type* `)`
| `\lockset`
| `\is_initialized` `(` *reference-type* `)`
| `\invariant_for` `(` *spec-expression* `)`
| `(` `\lblneg` *ident spec-expression* `)`
| `(` `\lblpos` *ident spec-expression* `)`
| *spec-quantified-expr*

set-comprehension ::= `{` *type-spec quantified-var-declarator*
'`|`' *set-comprehension-pred* `}`
set-comprehension-pred ::= *postfix-expr* `.` `has` `(` *ident* `)`
`&&` *predicate*

spec-quantified-expr ::= `(` *quantifier quantified-var-decls* `;` [[*predicate*] `;`]
spec-expression `)`
quantifier ::= `\forall` | `\exists` | `\max` | `\min` | `\num_of` | `\product` | `\sum`
quantified-var-decls ::= *type-spec quantified-var-declarator*
[`,` *quantified-var-declarator*] . . .
quantified-var-declarator ::= *ident* [*dims*]

spec-variable-declarators ::= *spec-variable-declarator*
[`,` *spec-variable-declarator*] . . .
spec-variable-declarator ::= *ident* [*dims*] [`=` *spec-initializer*]
spec-array-initializer ::= `{` [*spec-initializer*
[`,` *spec-initializer*] . . . [`,`]] `}`
spec-initializer ::= *spec-expression*
| *spec-array-initializer*

All of the JML keywords that can be used in expressions which would otherwise start with an alphabetic character start with a backslash (\), so that they cannot clash with the program's variable names.

The new expressions that JML introduces are described below. Several of the descriptions below quote, without attribution, descriptions from [Leavens-Baker-Ruby02].

13.1.1 `\result`

The JML keyword `\result` can only be used in `ensures` clauses of a non-void method. Its value is the value returned by the method. Its type is the return type of the method; hence it is a type error to use `\result` in a void method or in a constructor. The keyword `\result` can only be used in an *ensures-clause*; it cannot be used, for example, in preconditions or in signals clauses. [[[ALso can be used in duration and workingspace specifications – DRCok]]]

13.1.2 `\old`

The JML keyword `\old` can be used in both normal and exceptional postconditions (ie. in `ensures` and `signals` clauses), and in history constraints. An expression of the form `\old(`*Expr*`)` refers to the value that the expression *Expr* had in the pre-state of a method. The type of `\old(`*Expr*`)` is simply the type of *Expr*. [[[ALso can be used in duration and workingspace specifications – DRCok]]] [[[Would be nice to use \old in assert, assume, set statements – DRCok]]]

13.1.3 `\not_modified`

The JML keyword `\not_modified` can be used in both normal and exceptional preconditions (ie. in `ensures` and `signals` clauses), and in history constraints. It asserts that the values of objects [[[why not simply expressions ? – DRCok]]] are the same in the post-state as in the pre-state; for example, `\not_modified(xval,yval)` says that `xval` and `yval` have the same value in the pre- and post-states (in the sense of an `equals` method).

The type of a `\not_modified` expression is `boolean`.

13.1.4 `\fresh`

The operator `\fresh` asserts that objects were freshly allocated; for example, `\fresh(x,y)` asserts that `x` and `y` are not null and that the objects bound to these identifiers were not allocated in the pre-state. The arguments to `\fresh` can have any reference type, and the type of the overall expression is `boolean`

Note that it is wrong to use `\fresh(this)` in the specification of a constructor, because Java's `new` operator allocates storage for the object; the constructor's job is just to initialize that storage.

13.1.5 `\duration`

`\duration`, which describes the specified maximum number of virtual machine cycle times needed to execute the method call or explicit constructor invocation expression that is its argument; e.g., `\duration(myStack.push(o))` is the maximum number of virtual machine cycles needed to execute the call `myStack.push(o)`, according to the contract of the static type of `myStack`'s type's `push` method, when passed argument `o`. Note that the expression used as an argument to `\duration` should be thought of as quoted, in the sense that it is not to be executed; thus the method or constructor called need not be free of side effects. Note that the argument to `\duration` is an *expression* instead of just the name of a method, because because different method calls, i.e., those with different parameters, can use take different amounts of space [Krone-Ogden-Sitaraman03].

The argument expression passed to `\duration` must be a method call or explicit constructor invocation expression; the type of a `\duration` expression is `long`. [[[Why not make this simply an arbitrary expression? - DRCok]]]

For a given Java Virtual Machine, a *virtual machine cycle* is defined to be the minimum of the maximum over all Java Virtual Machine instructions, i, of the length of time needed to execute instruction i.

13.1.6 `\space`

`\space`, which describes the amount of heap space, in bytes, allocated to the object refered to by its argument [Krone-Ogden-Sitaraman03]; e.g., `\space(myStack)` is number of bytes in the heap used by `myStack`, not including the objects it contains. The type of the *spec-expression* that is the argument must be a reference type, and the result type of a `\space` expression is `long`.

13.1.7 `\working_space`

`\working_space`, which describes the maximum specified amount of heap space, in bytes, used by the method call or explicit constructor invocation expression that is its argument; e.g., `\working_space(myStack.push(o))` is the maximum number of bytes needed on the heap to execute the call `myStack.push(o)`, according to the contract of the static type of `myStack`'s type's `push` method, when passed argument `o`. [[[Why not allow the argument to be an expression ? - DRCok]]] Note that the expression used as an argument to `\working_space` should be thought of as quoted, in the sense that it is not to be executed; thus the method or constructor called need not be free of side effects. The detailed arguments are needed in the specification of the call because different method calls, i.e., those with different parameters, can use take different amounts of space [Krone-Ogden-Sitaraman03]. The argument expression must be a method call or explicit constructor invocation expression; the result type of a `\working_space` expression is `long`.

13.1.8 \reach

The `\reach` expression allows one to refer to the set of objects reachable from some particular object. The syntax `\reach(`*x*`)` denotes the smallest `JMLObjectSet` containing the object denoted by *x*, if any, and all objects accessible through all fields of objects in this set. That is, if *x* is `null`, then this set is empty otherwise it contains *x*, all objects accessible through all fields of *x*, all objects accessible through all fields of these objects, and so on, recursively. If *x* denotes a model field (or data group), then `\reach(`*x*`)` denotes the smallest `JMLObjectSet` containing the objects reachable from *x* or reachable from the objects referenced by fields in that data group.

13.1.9 \nonnullelements

The operator `\nonnullelements` can be used to assert that an array and its elements are all non-null. For example, `\nonnullelements(myArray)`, is equivalent to [Leino-Nelson-Saxe00]

```
myArray != null &&
(\forall int i; 0 <= i && i < myArray.length;
                myArray[i] != null)
```

13.1.10 Subtype operator

The relational operator `<:` compares two reference types and returns true when the type on the left is a subtype of the type on the right [Leino-Nelson-Saxe00]. Although the notation might suggest otherwise, this operator is also reflexive; a type will compare as `<:` with itself. In an expression of the form *E1* `<:` *E2*, both *E1* and *E2* must have type `\TYPE`; since in JML `\TYPE` is the same as `java.lang.Class` the expression *E1* `<:` *E2* means the same thing as the expression *E2*`.isAssignableFrom(`*E1*`)`.

13.1.11 \typeof

The operator `\typeof` returns the most-specific dynamic type of an expression's value [Leino-Nelson-Saxe00]. The meaning of `\typeof(`*E*`)` is unspecified if *E* is null. If *E* has a static type that is a reference type, then `\typeof(`*E*`)` means the same thing as *E*`.getClass()`. For example, if `c` is a variable of static type `Collection` that holds an object of class `HashSet`, then `\typeof(c)` is `HashSet.class`, which is the same thing as `\type(HashSet)`. If *E* has a static type that is not a reference type, then `\typeof(`*E*`)` means the instance of `java.lang.Class` that represents its static type. For example, `\typeof(true)` is `Boolean.TYPE`, which is the same as `\type(boolean)`. Thus an expression of the form `\typeof(`*E*`)` has type `\TYPE`, which JML considers to be the same as `java.lang.Class`.

13.1.12 \elemtype

The `\elemtype` operator returns the most-specific static type shared by all elements of its array argument [Leino-Nelson-Saxe00]. For example, `\elemtype(\type(int[]))` is

`\type(int)`. The argument to `\elemtype` must be an expression of type `\TYPE`, which JML considers to be the same as `java.lang.Class`, and its result also has type `\TYPE`.

[[[It might be helpful for consistency among tools if we defined \elemtype(t) where t is not an array type - perhaps to be one common value such as \typeof(null).]]]

13.1.13 `\type`

The operator `\type` can be used to introduce literals of type `\TYPE` in expressions. An expression of the form `\type(T)`, where `T` is a type name, has the type `\TYPE`. Since in JML `\TYPE` is the same as `java.lang.Class`, an expression of the form `\type(`T`)` means the same thing as T`.class`. For example, in

```
\typeof(myObj) <: \type(PlusAccount)
```

the use of `\type(PlusAccount)` is used to introduce the type `PlusAccount` into this expression context.

13.1.14 `\lockset`

The `\lockset` primitive denotes the set of locks held by the current thread. It is of type `JMLObjectSet`. (This is an adaptation from ESC/Java [Leino-etal00] [Leino-Nelson-Saxe00] for dealing with threads.)

13.1.15 `\max`

The `\max` operator returns the "largest" (as defined by `<`) of a set of lock objects, given a lock set as an argument. (This is an adaptation from ESC/Java [Leino-etal00] [Leino-Nelson-Saxe00] for dealing with threads.)

13.1.16 `\is_initialized`

The `\is_initialized` operator returns true just when its *reference-type* argument is a class that has finished its static initialization. It is of type `boolean`.

13.1.17 `\invariant_for`

The `\invariant_for` operator returns true just when its argument satisfies the invariant of its static type; for example, `\invariant_for((MyClass)o)` is true when `o` satisfies the invariant of `MyClass`. The entire `\invariant_for` expression is of type `boolean`.

13.1.18 `\lblneg` and `\lblpos`

Parenthesized expressions that start with `\lblneg` and `\lblpos` can be used to attach labels to expressions [Leino-Nelson-Saxe00]; these labels might be printed in various messages by support tools, for example, to identify an assertion that failed. Such an expression has a *label* and a *body*; for example, in

```
(\lblneg indexInBounds 0 <= index && index < length)
```

the label is `indexInBounds` and the body is the expression `0 <= index && index < length`. The value of a labeled expression is the value of its body, hence its type is the type of its body. The idea is that if this expression is used in an assertion and its value is `false` (e.g., when doing run-time checking of assertions), then a warning will be printed that includes the label `indexInBounds`. The form using `\lblpos` has a similar syntax, but should be used for warnings when the value of the enclosed expression is `true`.

13.1.19 Universal and Existential Quantifiers

The quantifiers `\forall` and `\exists`, are universal and existential quantifiers (respectively). For example,

```
(\forall int i,j; 0 <= i && i < j && j < 10; a[i] < a[j])
```

says that `a` is sorted at indexes between 0 and 9. The quantifiers range over all potential values of the variables declared which satisfy the *range* predicate, given between the semicolons (`;`). If the range predicate is omitted, it defaults to `true`. Since a quantifier quantifies over all potential values of the variables, when the variables declared are reference types, they may be null, or may refer to objects not constructed by the program; one should use a range predicate to eliminate such cases if they are not desired. The type of a universal and existential quantifier is `boolean`. [[[May the bound formal identifier in a quantifier expression be any identifier, or must it not redeclare a local variable? - DRCok]]]

13.1.20 Generalized Quantifiers

The quantifiers `\max`, `\min`, `\product`, and `\sum`, are generalized quantifiers that return the maximum, minimum, product, or sum of the values of the expressions given, where the variables satisfy the given range. The range predicate must be of type `boolean`. The expression in the body must be a built-in numeric type, such as `int` or `double`; the type of the quantified expression as a whole is the type of its body. The *body* of a quantified expression is the last top-level expression it contains; it is the expression following the range predicate, if there is one. As with the universal and existential quantifiers, if the range predicate is omitted, it defaults to `true`. For example, the following equations are all true (see chapter 3 of [Cohen90]):

```
(\sum int i; 0 <= i && i < 5; i) == 0 + 1 + 2 + 3 + 4
(\product int i; 0 < i && i < 5; i) == 1 * 2 * 3 * 4
(\max int i; 0 <= i && i < 5; i) == 4
(\min int i; 0 <= i && i < 5; i-1) == -1
```

For computing the value of a sum or product, Java's arithmetic is used. The meaning thus depends on the type of the expression. For example, in Java, floating point numbers use the IEEE 754 standard, and thus when an overflow occurs, the appropriate positive or negative infinity is returned. However, Java integers wrap on overflow. Consider the following examples.

```
(\product float f; 1.0e30f < f && f < 1.0e38f; f)
  == Float.POSITIVE_INFINITY
```

```
(\sum int i; i == Integer.MAX_VALUE || i == 1; i)
  == Integer.MAX_VALUE + 1
  == Integer.MIN_VALUE
```

When the range predicate is not satisfiable, the sum is 0 and the product is 1; for example:

```
(\sum int i; false; i) == 0
(\product double d; false; d*d) == 1.0
```

When the range predicate is not satisfiable for `\max` the result is the smallest number with the type of the expression in the body; for floating point numbers, negative infinity is used. Similarly, when the range predicate is not satisfiable for `\min`, the result is the largest number with the type of the expression in the body.

13.1.21 Numerical Quantifier

The numerical quantifier, `\num_of`, returns the number of values for its variables for which the range and the expression in its body are true. Both the range predicate and the body must have type `boolean`, and the entire quantified expression has type `long`. The meaning of this quantifier is defined by the following equation (see p. 57 of [Cohen90]).

```
(\num_of T x; R(x); P(x)) == (\sum T x; R(x) && P(x); 1L)
```

13.1.22 Set Comprehension

The set comprehension notation can be used to succinctly define sets. For example, the following is the `JMLObjectSet` that is the subset of non-`null` `Integer` objects found in the set `myIntSet` whose values are between 0 and 10, inclusive.

```
new JMLObjectSet {Integer i | myIntSet.has(i) &&
    i != null && 0 <= i.getInteger() && i.getInteger() <= 10 }
```

The syntax of JML limits set comprehensions so that following the vertical bar (`|`) is always an invocation of the has method of some set on the variable declared. (This restriction is used to avoid Russell's paradox [Whitehead-Russell25].) In practice, one either starts from some relevant set at hand, or one can start from the sets containing the objects of primitive types found in `org.jmlspecs.models.JMLObjectSet` and (in the same Java package) `JMLValueSet`. The type of such an expression is the type named following new, which must be `JMLObjectSet` or `JMLValueSet`.

[[[May the bound variable hide other variables that are in scope? – DRCok]]]

13.1.23 <==> and <=!=>

[[[discussion needed]]]

13.1.24 ==> and <==

[[[discussion needed]]]

13.1.25 informal predicates

[[[discussion needed]]]

13.1.26 primitives for safe arithmetic

[[[discussion needed]]]

13.1.27 lockset membership

[[[discussion needed]]]

13.2 Specification Expressions

The JML syntax for expressions extends the Java syntax with several operators and primitives.

The precedence of operators in JML expressions is similar to that in Java The precedence levels are given in the following table, where the parentheses, quantified expressions, `[]`, `.`, and method calls on the first three lines all have the highest precedence, and for the rest, only the operators on the same line have the same precedence.

highest `new () \forall \exists \max \min`
`\num_of \product \sum` *informal-description*
`[] .` and method calls
unary `+` and `- ~ ! (typecast)`
`* / %`
`+` (binary) `-` (binary)
`<< >> >>>`
`< <= > >= <: instanceof`
`== !=`
`&`
`^`
`|`
`&&`
`||`
`==> <==`
`<==> <=!=>`
`?:`
lowest `= *= /= %= += -= <<= >>= >>>= &= ^= |=`

The following is the syntax of Java expressions, with JML additions. The additions are the operators `==>`, `<==`, `<==>`, `<=!=>`, and `<:`, and the syntax found under the nonterminals *jml-primary*, *set-comprehension*, and *spec-quantified-expr* (see Chapter 13 [Predicates and Specification Expressions], page 57). The JML additions to the Java syntax can only be used in assertions and other annotations. Furthermore, within assertions, one cannot use any of the operators (such as `++`, `--`, and the assignment operators) that would necessarily cause side effects.

expression-list ::= *expression* [`,` *expression*] . . .
expression ::= *assignment-expr*
assignment-expr ::= *conditional-expr* [*assignment-op assignment-expr*]
assignment-op ::= `=` | `+=` | `-=` | `*=` | `/=` | `%=` | `>>=`
| `>>>=` | `<<=` | `&=` | `'|='` | `^=`
conditional-expr ::= *equivalence-expr*
[`?` *conditional-expr* `:` *conditional-expr*]
equivalence-expr ::= *implies-expr* [*equivalence-op implies-expr*] . . .
equivalence-op ::= `<==>` | `<=!=>`
implies-expr ::= *logical-or-expr*
[`==>` *implies-non-backward-expr*]
| *logical-or-expr* `<==` *logical-or-expr*
[`<==` *logical-or-expr*] . . .
implies-non-backward-expr ::= *logical-or-expr*
[`==>` *implies-non-backward-expr*]
logical-or-expr ::= *logical-and-expr* [`'||'` *logical-and-expr*] . . .
logical-and-expr ::= *inclusive-or-expr* [`&&` *inclusive-or-expr*] . . .
inclusive-or-expr ::= *exclusive-or-expr* [`'|'` *exclusive-or-expr*] . . .
exclusive-or-expr ::= *and-expr* [`^` *and-expr*] . . .
and-expr ::= *equality-expr* [`&` *equality-expr*] . . .
equality-expr ::= *relational-expr* [`==` *relational-expr*] . . .
| *relational-expr* [`!=` *relational-expr*] . . .
relational-expr ::= *shift-expr* `<` *shift-expr*
| *shift-expr* `>` *shift-expr*
| *shift-expr* `<=` *shift-expr*
| *shift-expr* `>=` *shift-expr*
| *shift-expr* `<:` *shift-expr*
| *shift-expr* [`instanceof` *type-spec*]
shift-expr ::= *additive-expr* [*shift-op additive-expr*] . . .
shift-op ::= `<<` | `>>` | `>>>`
additive-expr ::= *mult-expr* [*additive-op mult-expr*] . . .
additive-op ::= `+` | `-`
mult-expr ::= *unary-expr* [*mult-op unary-expr*] . . .
mult-op ::= `*` | `/` | `%`
unary-expr ::= `(` *type-spec* `)` *unary-expr*
| `++` *unary-expr*
| `--` *unary-expr*
| `+` *unary-expr*
| `-` *unary-expr*
| *unary-expr-not-plus-minus*
unary-expr-not-plus-minus ::= `~` *unary-expr*
| `!` *unary-expr*
| `(` *builtinType* `)` *unary-expr*
| `(` *reference-type* `)` *unary-expr-not-plus-minus*
| *postfix-expr*
postfix-expr ::= *primary-expr* [*primary-suffix*] . . . [`++`]
| *primary-expr* [*primary-suffix*] . . . [`--`]

| *builtinType* ['[' ']'] ... `.` `class`
primary-suffix ::= `.` *ident*
| `.` `this`
| `.` `class`
| `.` *new-expr*
| `.` `super` `(` [*expression-list*] `)`
| `(` [*expression-list*] `)`
| '[' *expression* ']'
| ['[' ']'] ... `.` `class`
primary-expr ::= *ident* | *new-expr*
| *constant* | `super` | `true`
| `false` | `this` | `null`
| `(` *expression* `)`
| *jml-primary*
| *informal-description*
builtInType ::= `void` | `boolean` | `byte`
| `char` | `short` | `int`
| `long` | `float` | `double`
constant ::= *java-literal*
new-expr ::= `new` *type new-suffix*
new-suffix ::= `(` [*expression-list*] `)` [*class-block*]
| *array-decl* [*array-initializer*]
| *set-comprehension*
array-decl ::= *dim-exprs* [*dims*]
dim-exprs ::= '[' *expression* ']' ['[' *expression* ']'] ...
dims ::= '[' ']' ['[' ']'] ...
array-initializer ::= `{` [*initializer* [`,` *initializer*] ... [`,`]] `}`
initializer ::= *expression*
| *array-initializer*

13.3 Store Refs

[[[Needs discussion, where used, what do they represent?]]]

The syntax related to the *store-ref* production is used in several places.

store-ref-list ::= *store-ref* [`,` *store-ref*] ...
store-ref ::= *store-ref-expression*
| *informal-description*
| *store-ref-keyword*
store-ref-expression ::= *store-ref-name* [*store-ref-name-suffix*] ...
store-ref-name ::= *ident* | `super` | `this`
store-ref-name-suffix ::= `.` *ident* | `.` `this` | '[' *spec-array-ref-expr* ']' | `.` `*`
spec-array-ref-expr ::= *spec-expression*
| *spec-expression* `..` *spec-expression*
| `*`
store-ref-keyword ::= `\nothing` | `\everything` | `\not_specified` | `\private_data`

14 JML primitive types

14.1 \TYPE

[[[needs discussion]]]

14.2 \real

[[[needs discussion]]]

14.3 \bigint

[[[needs discussion]]]

15 Statements and Annotation Statements

JML also defines a number of annotation statements that may be interspersed with Java statements in the body of a method, constructor, or initialization block.

The following gives the syntax of statements. These are the standard Java statements, with the addition of annotations, the *hence-by-statement*, *assert-redundantly-statement*, *assume-statement*, *set-statement*, and *unreachable-statement*, and the various forms of *model-prog-statement*. See Chapter 17 [Model Programs], page 73, for the syntax of *model-prog-statement*, which is only allowed in model programs. [[[Does this include local class declarations?]]]

compound-statement ::= `{` *statement* [*statement*] . . . `}`
statement ::= *compound-statement*
| *local-declaration* `;`
| *ident* `:` *statement*
| *expression* `;`
| `if (` *expression* `)` *statement* [`else` *statement*]
| [*loop-invariant*] . . . [*variant-function*] . . . [*ident*] `:` *loop-stmt*
| `break` [*ident*] `;`
| `continue` [*ident*] `;`
| `return` [*expression*] `;`
| *switch-statement*
| *try-block*
| `throw` *expression* `;`
| `synchronized (` *expression* `)` *statement*
| `;`
| *assert-statement*
| *hence-by-statement*
| *assert-redundantly-statement*
| *assume-statement*
| *set-statement*
| *unreachable-statement*
| *model-prog-statement* // only allowed in model programs
loop-stmt ::= `while (` *expression* `)` *statement*
| `do` *statement* `while (` *expression* `) ;`
| `for (` [*for-init*] `;` [*expression*] `;` [*expression-list*] `)`
statement
for-init ::= *local-declaration* | *expression-list*
local-declaration ::= *local-modifiers variable-decls*
local-modifiers ::= [*local-modifier*] . . .
local-modifier ::= `model | ghost | final | non_null`
[[[I don't think model can be a local-modifier - David]]]
switch-statement ::= `switch (` *expression* `) {` [*switch-body*] . . . `}`
switch-body ::= *switch-label-seq* [*statement*] . . .
switch-label-seq ::= *switch-label* [*switch-label*] . . .
switch-label ::= `case` *expression* `:` | `default :`
try-block ::= `try` *compound-statement* [*handler*] . . .

[`finally` *compound-statement*]
handler ::= `catch (` *param-declaration* `)` *compound-statement*
assert-statement ::= `assert` *expression* [`:` *expression*] `;`

Note that Java (as of J2SDK 1.4) also has its own `assert` statement. JML distinguishes between assert statements that occur inside and outside annotations. Inside an annotation, such a statement is a JML assert statement, and the first *expression* thus can't have side effects. (The second *expression* is a `String` that, as in Java, is printed if the assertion fails.) Outside an annotation, an assert statement is a Java assert statement, and so the first *expression* can have side effects (potentially, although it shouldn't).

The following gives the syntax of JML annotations that can be used on statements. See Chapter 17 [Model Programs], page 73, for the syntax of statements that can only be used in model programs.

hence-by-statement ::= *hence-by-keyword predicate* `;`
hence-by-keyword ::= `hence_by | hence_by_redundantly`
assert-redundantly-statement ::= `assert_redundantly` *predicate* [`:` *expression*] `;`
assume-statement ::= *assume-keyword predicate* [`:` *expression*] `;`
assume-keyword ::= `assume | assume_redundantly`
set-statement ::= `set` *assignment-expr* `;`
unreachable-statement ::= `unreachable ;`
loop-invariant ::= *maintaining-keyword predicate* `;`
maintaining-keyword ::= `maintaining | maintaining_redundantly`
`| loop_invariant | loop_invariant_redundantly`
variant-function ::= *decreasing-keyword spec-expression* `;`
decreasing-keyword ::= `decreasing | decreasing_redundantly`
`| decreases | decreases_redundantly`

15.1 assume statement

[[[needs discussion]]]

15.2 assert statement

[[[needs discussion]]]

15.3 Local ghost declaration

A local ghost declaration is a variable declaration with a `ghost` modifier, entirely contained in an annotation. It introduces a new variable that may be used in subsequent annotations within the remainder of the block in which the declaration appears. A ghost variable is not used in program execution as Java variables are, but is used by the static checker to reason about the execution of the routine body in which the ghost variable is used.

- The variable name may not be already declared as a local variable or local ghost variable or as a formal parameter of the routine in which the declaration appears.

- Each variable declared may have an initializer; the initializer is in the scope of the newly declared variable.
- The modifiers `final`, `uninitialized`, `non_null` may be used on the ghost declaration.

Examples:

```
//@ ghost int i = 0;
//@ ghost int i = 0, j, k = i+3;
//@ ghost float[] a = { 1,2,3 };
//@ ghost Object o;
//@ final ghost non_null Object o = new Object();
//@ ghost \TYPE t = \typeof(t);
```

15.4 set statement

A set statement is the equivalent of an assignment statement but is within an annotation. It is used to assign a value to a ghost variable or to a ghost field. A set statement serves to assist the static checker in reasoning about the execution of the routine body in which it appears.

- the target of the set statement must be a ghost variable or a ghost field
- the right-hand-side of the set statement must be pure (not have side effects)

Examples:

```
//@ set i = 0;
//@ set collection.elementType = \type(int);
```

[[[Questions: must the rhs be pure? Should we allow an arabitrary statement, not just an assignment? such as set ++i; or set i += 5;]]]

15.5 unreachable statement

The `unreachable` statement is an annotation that asserts that the control flow of a routine will never reach that point in the program. It is equivalent to `assert false`. If control flow does reach an `unreachable` statement, tool that checks (by reasoning or at runtime) the behavior of the routine should issue an error of some kind. The following is an example:

```
if (true) {
        ...
} else {
        //@ unreachable;
}
```

15.6 hence_by statement

[[[needs discussion]]]

15.7 loop_invariant statement

[[[needs discussion]]]

15.8 decreases statement

[[[needs discussion]]]

15.9 JML Modifiers for Local Declarations

The following modifiers may be applied to local declarations of either Java variables or ghost variables within routine bodies.

15.9.1 non_null

[[[needs discussion]]]

15.9.2 uninitialized

[[[needs discussion]]]

16 Redundancy

redundant-spec ::= *implications* [*examples*] | *examples*
examples ::= `for_example` *example* [`also` *example*] . . .

16.1 Redundant Implications

implications ::= `implies_that` *spec-case-seq*

16.2 Redundant Examples

The following gives the syntax of examples.

example ::= [[*privacy*] `example`]
[*spec-var-decls*] [*spec-header*] *simple-spec-body*
| [*privacy*] `exceptional_example`
[*spec-var-decls*] *spec-header* [*exceptional-example-body*]
| [*privacy*] `exceptional_example`
[*spec-var-decls*] *exceptional-example-body*
| [*privacy*] `normal_example`
[*spec-var-decls*] *spec-header* [*normal-example-body*]
| [*privacy*] `normal_example`
[*spec-var-decls*] *normal-example-body*
exceptional-example-body ::= *exceptional-spec-clause* [*exceptional-spec-clause*] . . .
normal-example-body ::= *normal-spec-clause* [*normal-spec-clause*] . . .

17 Model Programs

This chapter discusses JML's model programs, which are adapted from the refinement calculus [Back88] [Back-vonWright89a] [Buechi-Weck00] [Morgan94] [Morris87].

17.1 Ideas Behind Model Programs

The basic idea of a model program is that it is a specification that is written as an abstract algorithm. Such an abstract algorithm specifies a method in the sense that the method's execution should be a refinement of the model program.

In JML adopt the semantics of the "greybox approach" to refinement calculus [Buechi-Weck00] [Buechi00]. In this semantics, calls to non-pure methods in a model program must occur in the same states in a correct implementation. That is, the notion of refinement is restricted to not permit all implementations with equivalent behavior, but to require that the implementation make certain method calls in the model program. This turns out to be very nice for describing both higher-order features and callbacks.

Consider the following example (from a survey on behavioral subtyping by Leavens and Dhara [Leavens-Dhara00]). In this example, both the methods are specified using model programs, which are explained below.

```
package org.jmlspecs.samples.dirobserver;

//@ model import org.jmlspecs.models.JMLString;
//@ model import org.jmlspecs.models.JMLObjectSetEnumerator;

public interface Directory extends RODirectory {

  /*@ public model_program {
    @   normal_behavior
    @     requires !in_notifier && n != null && n != "" && f != null;
    @     assignable entries;
    @     ensures entries != null
    @        && entries.equals(\old(entries.extend(new JMLString(n), f)));
    @   for (JMLObjectSetEnumerator e = listeners.elements();
    @        e.hasMoreElements(); ) {
    @     in_notifier = true;
    @     ((DirObserver)e.nextElement()).addNotification(this, n);
    @     in_notifier = false;
    @   }
    @ }
    @*/
  public void addEntry(String n, File f);

  /*@ public model_program {
    @   normal_behavior
    @     requires !in_notifier && n != null && n != "";
```

```
    @     assignable entries;
    @     ensures entries != null
    @        && entries.equals(\old(entries.remove(new JMLString(n))));
    @   for (JMLObjectSetEnumerator e = listeners.elements();
    @        e.hasMoreElements(); ) {
    @     in_notifier = true;
    @     ((DirObserver)e.nextElement()).removeNotification(this, n);
    @     in_notifier = false;
    @   }
    @ }
    @*/
  public void removeEntry(String n);
}
```

Both model programs in the above example are formed from a specification statement, which begins with the keyword `normal_behavior` in these examples, and a for-loop. The key event in the for loop bodies is a method call to a non-pure method (`addNotification` or `removeNotification`). These calls must occur in a state equivalent to the one reached in the model program for the implementation to be legal.

The behavior statements abstract away part of a correct implementation. The `normal_behavior` statements in these examples both have a precondition, a frame axiom, and a postcondition. These mean that the statements that they abstract away from must be able to, in any state satisfying the precondition, finish in a state satisfying the postcondition, while only assigning to the locations (and their dependees) named in the frame axiom. For example, the first behavior statements says that whenever `in_notifier` is false, `n` is not null and not empty, and `f` is not null, then this part of the method can assign to `entries` something that isn't null and that is equal to the old value of `entries` extended with a pair consisting of the string `n` and the file `f`.

The model field `entries`, of type `JMLValueToObjectMap`, is declared in the supertype `RODirectory` [Leavens-Dhara00].

17.2 Details of Model Programs

The following gives the syntax of model programs. See Chapter 15 [Statements and Annotation Statements], page 68, for the parts of the syntax of statements that are unchanged from Java. The *jml-compound-statement* and *jml-statement* syntax is the same as the *compound-statement* and *statement* syntax, except that *model-prog-statement*s are not flagged as errors within the *jml-compound-statement* and *jml-statement*s.

model-program ::= [*privacy*] `model_program` *jml-compound-statement*
jml-compound-statement ::= *compound-statement*
jml-statement ::= *statement*
model-prog-statement ::= *nondeterministic-choice*
| *nondeterministic-if*
| *spec-statement*
| *invariant*
nondeterministic-choice ::= `choose` *alternative-statements*

alternative-statements ::= *jml-compound-statement*
[`or` *jml-compound-statement*] . . .
nondeterministic-if ::= `choose_if` *guarded-statements*
[`else` *jml-compound-statement*]
guarded-statements ::= *guarded-statement*
[`or` *guarded-statement*] . . .
guarded-statement ::= `{`
assume-statement
jml-statement [*jml-statement*] . . .
`}`

The grammar for specification statements appears below. It is unusual, compared to specification statements in refinement calculus, in that it allows one to specify statements that can signal exceptions, or terminate abruptly. The reasons for this are based on verification logics for Java [Huisman01] [Poll-Jacobs00], which have these possibilities. The meaning of an *abrupt-spec-case* is that the normal termination and signaling an exception are forbidden; that is, the equivalent *spec-statement* using `behavior` would have `ensures false;` and `signals (Exception) false;` clauses.

spec-statement ::= [*privacy*] `behavior` *generic-spec-statement-case*
| [*privacy*] `exceptional_behavior` *exceptional-spec-case*
| [*privacy*] `normal_behavior` *normal-spec-case*
| [*privacy*] `abrupt_behavior` *abrupt-spec-case*
generic-spec-statement-case ::= [*spec-var-decls*] *generic-spec-statement-body*
| [*spec-var-decls*] *spec-header* [*generic-spec-statement-body*]
generic-spec-statement-body ::= *simple-spec-statement-body*
| `{|` *generic-spec-statement-case-seq* `|}`
generic-spec-statement-body-seq ::= *generic-spec-statement-case*
[`also` *generic-spec-statement-case*] . . .
simple-spec-statement-body ::= *simple-spec-statement-clause* [*simple-spec-statement-clause*] . . .
simple-spec-statement-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause* | *duration-clause*
| *ensures-clause* | *signals-clause*
| *continues-clause* | *breaks-clause* | *returns-clause*
abrupt-spec-case ::= [*spec-var-decls*] *spec-header*
[*abrupt-spec-body*]
| [*spec-var-decls*] *abrupt-spec-body*
abrupt-spec-body ::= *abrupt-spec-clause* [*abrupt-spec-clause*] . . .
| `{|` *abrupt-spec-case-seq* `|}`
abrupt-spec-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause* | *duration-clause*
| *continues-clause* | *breaks-clause* | *returns-clause*
abrupt-spec-case-seq ::= *abrupt-spec-case* [`also` *abrupt-spec-case*] . . .

continues-clause ::= *continues-keyword* [*target-label*] [*pred-or-not*] `;`

continues-keyword ::= `continues | continues_redundantly`
target-label ::= `->` `(` *ident* `)`
breaks-clause ::= *breaks-keyword* [*target-label*] [*pred-or-not*] `;`
breaks-keyword ::= `breaks | breaks_redundantly`
returns-clause ::= *returns-keyword* [*pred-or-not*] `;`
returns-keyword ::= `returns | returns_redundantly`

18 Specification for Subtypes

code-contract-spec ::= `code_contract` *code-contract-clause* [*code-contract-clause*] . . .

code-contract-clause ::= *accessible-clause* | *callable-clause* | *measured-clause*

[[[Should the section where 'accessible' and 'callable' are discussed as part of the method specs be moved here?]]]

See Section 11.9.12 [Accessible Clauses], page 54, for the syntax and semantics of the *accessible-clause*. See Section 11.9.13 [Callable Clauses], page 55, for the syntax and semantics of the *callable-clause*.

See [Kiczales-Lamping92] [Steyaert-etal96] for the problem and [Ruby-Leavens00] for how JML deals with these problems.

[[[Need an explanation of the callable clause]]]

19 Refinement

This chapter explains JML's notion of refinement files, which uses the following syntax.

refine-prefix ::= `refine` *string-literal* `;`

The *refine-prefix* in a compilation unit says that the declarations in this compilation unit refine the corresponding declarations in the file named by the *string-literal*. The *string-literal* should name a file, complete with a suffix, for example, `"MyType.java-refined"`. The suffix of such a file is used by JML tools to find the file that is the base of a refinement chain, and all other files in the chain are found using the files named in the *refine-prefix* of a previous file in the chain.

One use of this feature is to allow specifications to be written first, in a separate file from the source code. For example, to specify `MyType`, and then write the code for it in a file named '`MyType.java-refined`', and then one would write Java code in '`MyType.java`'. The file '`MyType.java`' would also include the following *refine-prefix*.

```
refine "MyType.java-refined";
```

Another typical use of this feature is to allow one to add specifications to source code that one does not want to modify. To do that, one would use a '`.refines-java`' (or '`.refines-spec`' or '`.refines-jml`') file with the specifications of the corresponding Java file in it. For example, if one wants to specify the type `LibraryType`, without touching the file '`LibraryType.java`' then one could write specifications in the file '`LibraryType.refines-java`', and include in that file the following *refine-prefix*.

```
refine "LibraryType.java";
```

The following gives more details about the checks and meaning of this feature of JML.

19.1 File Name Suffixes

The JML tools recognize several filename suffixes. The following are considered to be *active* suffixes: '`.refines-java`', '`.refines-spec`', '`.refines-jml`', '`.java`', '`.spec`', and '`.jml`'; There are also three *passive* suffixes: '`.java-refined`', '`.spec-refined`', and '`.jml-refined`'. Files with passive suffixes can be used in refinements but should not normally be passed explicitly to the tools directly.

The filename suffixes are ordered so tools can find the base file in a refinement chain. In the above list of active suffixes, those listed earlier are considered to be *more active* than those listed later. Among files with the same prefix, the one whose suffix appears first in the above list of active suffixes is considered primary by JML tools. For example, when searching for the compilation unit containing a type declaration, JML tools look for those files in the package directory with a prefix that matches the type name; they then select the file with the most active suffix as the base of the refinement sequence. When the base is correctly selected, then all the rest of the files in the refinement sequence will be loaded automatically so that all the specifications for that type are available.

To help ensure that the base is correctly selected, the file with the most active suffix must be the base of a refinement sequence, otherwise the JML typechecker issues an error message. Also, the prefix of the base file must be the same as the public type declared in that compilation unit or an error message is issued. However, it is not necessary that

the file being refined have the same prefix as the file at the base of the refinement chain. Furthermore, a file with the same prefix as the base file may not be in a different refinement sequence. For example, '`SomeName.java-refined`' can be refined by '`MyType.java`' as long as there is no refinement sequence with '`SomeName`' as the prefix of the base of another refinement.

The JML tools deal with all files in a refinement chain whenever one of them is selected for processing by the tool. This allows all of the specifications that apply to be consistently dealt with at all times. For example, supose that there are files named '`Foo.refines-java`' and '`Foo.java`', then if a tool selects the '`Foo.java`', e.g., with the command:

```
$ jmlc *.java
```

then it will see both the '`Foo.refines-java`' and the '`Foo.java`' file.

[[[Have to clarify how this interacts with the CLASSPATH]]]

19.2 Type Checking Refinements

There are some restrictions on what can appear in the different files involved in a particular refinement. Since the Java compilers only see the '`.java`' files, executable code (that is not just for use in specifications) should only be placed in the '`.java`' files. In particular the following restrictions are enforced by JML.

- When the same method is declared in more than one file in a refinement sequence, most parts of the method declaration must be identical in all the files. (Two method declarations are considered to be declaring the *same method* if they have the same signature, i.e., same name and static formal parameter types.) However, in addition to the signature of such a method, the return type, the names of the formal parameters, the declared exceptions the method may throw, and the non-JML modifiers `public`, `protected`, `private`, `static`, and `final`, must all match exactly in each such declaration in a refinement chain.
- The `model` modifier must appear in all declarations of a given method or it must appear in none of them. It is not permitted to implement a model method with a non-model method or to refine a non-model method with a model method. Use a `spec_public` or `spec_protected` method if you want to use that method in a specification.
- Some of the JML method modifiers do not always have to match in all declarations of the same method in a refinement chain. One may add `pure`, `non_null`, `spec_public`, or `spec_protected` to any of the declarations for a method in any file. However, if `pure` is added to a method specification, then all subsequent declarations of that method in a refinement sequence must also be declared `pure`. Also, it is, of course, not permitted to add `spec_protected` to a method that has been declared `public` or `spec_public` in other declarations. One can add `non_null` to any formal parameter in any file, although good style suggests that all of these annotations appear on one declaration of that method.
- The specification of a refining method declaration must start with the JML keyword `also`; if it does not an error message is issued. A *refining method declaration* is a declaration that overrides a superclass method or refines the specification of the same method in a refinement chain. In JML, method specifications are inherited by

subclasses and in refinement chains. The `also` keyword indicates that the current specification is refining the specification inherited either from the superclass or from the previous declaration of the method in a refinement sequence. Therefore, it is an error if the specification of a non-refining method begins with `also` (unless it overrides an inherited method).

- If a non-model method has a body, then the body can only appear in a '`.java`' file; an error message is issued if the body of a non-model method appears in a file with any other suffix. Furthermore, the body of a model method may only appear in one file of a refinement sequence. This means that each method of each class can have at most one method body.
- When the same field is declared in more than one file in a refinement sequence, then the signature of each such declaration must be identical in all the files. (Two field declarations are considered to be declaring the *same field* if they have the same name.) The signature of such a field, including its type, the non-JML modifiers `public`, `protected`, `private`, `static`, and `final`, must all match exactly in each such declaration.
- All declarations of a given field must either use the modifier `model` or not. It is not permitted to implement a model field with a non-model field or vice versa. Use a `spec_public` or `spec_protected` field if you want to use the same name. The same comment holds for `ghost` fields.
- Some of the JML field modifiers do not always have to match in all declarations of the same field in a refinement chain. One may add `non_null`, `spec_public`, or `spec_protected` to any of of the declarations for a field in any file. However, it is of course not permitted to add `spec_protected` to a field that has been declared public in other declarations.
- Initializers are not allowed in all field declarations. A non-model field can have an initializer expression but it can only appear in a '`.java`' file because this is where a compiler expects to find it.

  Fields declared using the `ghost` modifier can have an initializer expression in any file, but they may have at most one initializer expression in all the files.

  Model fields cannot have an initializer expression because there is no storage associated with such fields. Use the `initially` clause to specify the initial state of model fields (although the initial state is usually determined from the represents clause).
- Any number of *jml-var-assertion*'s can be declared for any field declaration and these are all conjoined. For example, if a variable `int count` is declared and there are two `initially` clauses, in the same or different files, then these initially clause predicates are conjoined; that is, both must be satisfied initially.
- An initializer block or a static initializer block (with code) may only appear in a '`.java`' file. One can write annotations to specify the effects of such initializers in JML annotations in other files, using the keywords `initializer` and `static_initializer`.

JML uses specification inheritance to impose the specifications of supertypes on their subtypes [Dhara-Leavens96] to support the concept of behavioral subtyping [America87] [Leavens90] [Leavens91] [Leavens-Weihl90] [Leavens-Weihl95] [Liskov-Wing94]. JML also supports a notion of weak behavioral subtyping [Dhara-Leavens94b] [Dhara97].

19.3 Refinement Viewpoints

In refinements, specification inheritance allows the specifier to separate the public, protected, and private specifications into different files. Public specifications give the public behavior and are meant for clients of the class. Protected specifications are meant for programmers creating subclasses and give the protected behavior of the type; they give the behavior of protected methods and fields that are not visible to clients. Similarly, private specifications are meant for implementors of the class and provide the behavior related to private methods and fields of the class; implementors must satisfy the combined public, protected, and private specifications of a method.

[[[Needs work]]]

19.3.1 default constructor refinement

In Java, a default constructor is automatically generated for a class when no constructors are declared in a class. However, in JML, a default constructor is not generated for a class unless the file suffix is '`.java`' (the same constructor is generated as in the Java language). Consider, for example, the following refinement sequence.

```
// ------ file MyClass.jml-refined --------
public class MyClass {
   //@ public model int x;

   /*@ public normal_behavior
     @    ensures x == 0; @*/
   public MyClass();
}

// --------- file MyClass.jml -------------
// refine "MyClass.jml-refined";
public class MyClass {
   protected int x_;
   //@           in x;

   //@ protected represents x <- x_;
}

// ---------- file MyClass.java -----------
// refine "MyClass.jml";
public class MyClass {
   protected int x_;
   public MyClass() { x_ = 0; }
}

// ----------------------------------------
```

In the protected specification declared in '`MyClass.jml`', no constructor is defined. If JML were to generate a default constructor for this class declaration, then the `public` constructor defined earlier in the refinement chain, in '`MyClass.jml-refined`', could have a visibility modifier that conflicts with the one automatically generated for the protected specification (the visibility modifier of an automatically generated default constructor depends on other factors including the visibility of the class). Recall that the signature, including the visibility modifier, must match for every method and constructor declared in a refinement chain. To avoid such conflicts, JML does not generate a default constructor unless the file suffix is '`.java`' (as part of the standard compilation process).

A similar problem can occur when the only constructor is protected or private as in the following example.

```
// ------ file MyClass2.jml-refined --------
public class MyClass2 {
   //@ public model int x;
   //@ public initially x == 0;
}

// --------- file MyClass2.jml -------------
// refine "MyClass2.jml-refined";
public class MyClass2 {
   protected int x_;
   //@           in x;

   //@ protected represents x <- x_;

   /*@ protected normal_behavior
     @    ensures x == 0; @*/
   protected MyClass2();
}

// ---------- file MyClass2.java -----------
// refine "MyClass2.jml";
public class MyClass2 {
   protected int x_;
   protected MyClass2() { x_ = 0; }
}

// ----------------------------------------
```

In this example, no constructor is defined for the public specification in '`MyClass2.jml-refined`'. If a default constructor were generated for this class declaration, then the `protected` constructor defined later in the refinement chain, in '`MyClass2.jml`', would have a visibility modifier that conflicts with the one automatically generated and

JML would emit an error. Thus JML only generates the default constructor for the executable declaration of a class in the '`.java`' file and only when required by the Java language.

Appendix A Grammar Summary

The following is a summary of the context-free grammar for JML. See Chapter 3 [Syntax Notation], page 12, for the notation used. In the first section below, grammatical productions are to be understood lexically. That is, no white space (see Section 4.1 [White Space], page 13) may intervene between the characters of a token.

A.1 Lexical Conventions

microsyntax ::= *lexeme* [*lexeme*] . . .
lexeme ::= *white-space* | *lexical-pragma* | *comment*
| *annotation-marker* | *doc-comment* | *token*
token ::= *ident* | *keyword* | *special-symbol* | *java-literal*
| *informal-description*
white-space ::= *non-nl-white-space* | *end-of-line*
non-nl-white-space ::= a blank, tab, or formfeed character
end-of-line ::= *newline* | *carriage-return* | *carriage-return newline*
newline ::= a newline character
carriage-return ::= a carriage return character
lexical-pragma ::= *nowarn-pragma*
nowarn-pragma ::= `nowarn` [*spaces* [*nowarn-label-list*]] `;`
spaces ::= *non-nl-white-space* [*non-nl-white-space*] . . .
nowarn-label-list ::= *nowarn-label* [*spaces*] [`,` [*spaces*] *nowarn-label* [*spaces*]] . . .
nowarn-label ::= *letter* [*letter*] . . .
comment ::= *C-style-comment* | *C++-style-comment*
C-style-comment ::= `/*` [*C-style-body*] *C-style-end*
C-style-body ::= *non-at-plus-star* [*non-star-slash*] . . .
| `+` *non-at* [*non-star-slash*] . . .
| stars-non-slash [non-star-slash] . . .
non-star-slash ::= *non-star*
| *stars-non-slash*
stars-non-slash ::= `*` [`*`] . . . *non-slash*
non-at-plus-star ::= any character except `@`, `+`, or `*`
non-at ::= any character except `@`
non-star ::= any character except `*`
non-slash ::= any character except `/`
C-style-end ::= [`*`] . . . `*/`
C++-style-comment ::= `//` [`+`] *end-of-line*
| `//` *non-at-plus-end-of-line* [*non-end-of-line*] . . . *end-of-line*
| `//+` *non-at-end-of-line* [*non-end-of-line*] . . . *end-of-line*
non-end-of-line ::= any character except a newline or carriage return
non-at-plus-end-of-line ::= any character except `@`, `+`, newline, or carriage return
non-at-end-of-line ::= any character except `@`, newline, or carriage return
annotation-marker ::= `//@` | `//+@`
| `/*@` | `/*+@` | `@+*/` | `@*/` | `*/`

ignored-at-in-annotation ::= `@`
doc-comment ::= `/**` [`*`] . . . *doc-comment-body* `*/`
doc-comment-ignored ::= *doc-comment*
doc-comment-body ::= [*description*] . . .
description ::= *doc-non-empty-textline*
tagged-paragraph ::= *paragraph-tag* [*doc-non-nl-ws*] . . .
jml-specs ::= *jml-tag* [*method-specification*] *end-jml-tag*
paragraph-tag ::= `@author` | `@deprecated` | `@exception`
| `@param` | `@return` | `@see`
| `@serial` | `@serialdata` | `@serialfield`
| `@since` | `@throws` | `@version`
| `@` *letter* [*letter*] . . .
doc-atsign ::= `@`
doc-nl-ws ::= *end-of-line* [*doc-non-nl-ws*] . . . [`*` [`*`] . . . [*doc-non-nl-ws*] . . .]
doc-non-nl-ws ::= *non-nl-white-space*
doc-non-empty-textline ::= *non-at-end-of-line* [*non-end-of-line*] . . .
jml-tag ::= `<jml>` | `<JML>` | `<esc>` | `<ESC>`
end-jml-tag ::= `</jml>` | `</JML>` | `</esc>` | `</ESC>`
ident ::= *letter* [*letter-or-digit*] . . .
letter ::= `_`, `$`, `a` through `z`, or `A` through `Z`
digit ::= `0` | `1` | `2` | `3` | `4` | `5` | `6` | `7` | `8` | `9`
letter-or-digit ::= *letter* | *digit*
keyword ::= *java-keyword* | *jml-predicate-keyword* | *jml-keyword*
jml-predicate-keyword ::= `\duration` | `\elemtype`
| `\everything` | `\exists`
| `\forall` | `\fresh` | `\invariant_for`
| `\is_initialized` | `\lblneg` | `\lblpos`
| `\lockset` | `\max` | `\min`
| `\nonnullelements` | `\nothing` | `\not_modified`
| `\not_specified` | `\num_of` |`\old`
| `\other` | `\private_data` | `\product`
| `\reach` | `\result` | `\space`
| `\such_that` | `\sum` | `\type`
| `\typeof` | `\TYPE` | `\working_space`
jml-keyword ::= `abrupt_behavior` | `accessible_redundantly` | `accessible`
| `also` | `assert_redundantly`
| `assignable_redundantly` | `assignable`
| `assume_redundantly` | `assume` | `axiom`
| `behavior` | `breaks_redundantly` | `breaks`
| `callable_redundantly` | `callable` | `choose_if`
| `choose` | `code_contract`
| `constraint_redundantly` | `constraint`
| `constructor` | `continues_redundantly` | `continues`
| `decreases_redundantly` | `decreases`
| `decreasing_redundantly` | `decreasing`
| `diverges_redundantly` | `diverges` | `duration_redundantly`
| `duration` | `ensures_redundantly`

```
      | ensures | example | exceptional_behavior
      | exceptional_example | exsures_redundantly | exsures
      | field | forall
      | for_example | ghost
      | implies_that | helper | hence_by_redundantly
      | hence_by | initializer | initially
      | instance | invariant_redundantly | invariant
      | loop_invariant_redundantly | loop_invariant
      | maintaining_redundantly | maintaining
      | measured_by_redundantly | measured_by | method
      | model_program | model | modifiable_redundantly
      | modifiable | modifies_redundantly | modifies
      | monitors_for | monitored | non_null
      | normal_behavior | normal_example | nowarn
      | old | or | post_redundantly | post
      | pre_redundantly | pre
      | pure | readable | refine
      | represents_redundantly | represents | requires_redundantly
      | requires | returns_redundantly | returns
      | set | signals_redundantly | signals
      | spec_protected | spec_public | static_initializer
      | uninitialized | unreachable
      | weakly | when_redundantly | when
      | working_space_redundantly | working_space
```

special-symbol ::= *java-special-symbol* | *jml-special-symbol*

java-special-symbol ::= *java-separator* | *java-operator*

java-separator ::= `(` | `)` | `{` | `}` | '`[`' | '`]`' | `;` | `,` | `.`

java-operator ::= `=` | `<` | `>` | `!` | `~` | `?` | `:`
| `==` | `<=` | `>=` | `!=` | `&&` | '`||`' | `++` | `--`
| `+` | `-` | `*` | `/` | `&` | '`|`' | `^` | `%` | `<<` | `>>` | `>>>`
| `+=` | `-=` | `*=` | `/=` | `&=` | '`|=`' | `^=` | `%=` | `<<=` | `>>=` | `>>>=`

jml-special-symbol ::= `==>` | `<==` | `<==>` | `<=!=>` | `->` | `<-` | `<:` | `..` | '`{|`' | '`|}`'

java-literal ::= *integer-literal* | *floating-point-literal* | *boolean-literal*
| *character-literal* | *string-literal* | *null-literal*

integer-literal ::= *decimal-integer-literal* | *hex-integer-literal* | *octal-integer-literal*

decimal-integer-literal ::= *decimal-numeral* [*integer-type-suffix*]

decimal-numeral ::= `0` | *non-zero-digit* [*digits*]

digits ::= *digit* [*digit*] . . .

digit ::= `0` | *non-zero-digit*

non-zero-digit ::= `1` | `2` | `3` | `4` | `5` | `6` | `7` | `8` | `9`

integer-type-suffix ::= `l` | `L`

hex-integer-literal ::= *hex-numeral* [*integer-type-suffix*]

hex-numeral ::= `0x` *hex-digit* [*hex-digit*] . . . | `0X` *hex-digit* [*hex-digit*] . . .

hex-digit ::= *digit* | `a` | `b` | `c` | `d` | `e` | `f`
| `A` | `B` | `C` | `D` | `E` | `F`

octal-integer-literal ::= *octal-numeral* [*integer-type-suffix*]

octal-numeral ::= `0` *octal-digit* [*octal-digit*] . . .

octal-digit ::= `0` | `1` | `2` | `3` | `4` | `5` | `6` | `7`
floating-point-literal ::= *digits* `.` [*digits*] [*exponent-part*] [*float-type-suffix*]
| `.` *digits* [*exponent-part*] [*float-type-suffix*]
| *digits exponent-part* [*float-type-suffix*]
| *digits* [*exponent-part*] *float-type-suffix*
exponent-part ::= *exponent-indicator signed-integer*
exponent-indicator ::= `e` | `E`
signed-integer ::= [*sign*] *digits*
sign ::= `+` | `-`
float-type-suffix ::= `f` | `F` | `d` | `D`
boolean-literal ::= `true` | `false`
character-literal ::= `'` *single-character* `'` | `'` *escape-sequence* `'`
single-character ::= any character except `'`, `\`, carriage return, or newline
escape-sequence ::= `\b` // backspace
| `\t` // tab
| `\n` // newline
| `\r` // carriage return
| `\'` // single quote
| `\"` // double quote
| `\\` // backslash
| *octal-escape*
| *unicode-escape*
octal-escape ::= `\` *octal-digit* [*octal-digit*]
| `\` *zero-to-three octal-digit octal-digit*
zero-to-three ::= `0` | `1` | `2` | `3`
unicode-escape ::= `\u` *hex-digit hex-digit hex-digit hex-digit*
string-literal ::= `"` [*string-character*] . . . `"`
string-character ::= *escape-sequence*
| any character except `"`, `\`, carriage return, or newline
null-literal ::= `null`
informal-description ::= `(*` *non-star-close* [*non-star-close*] . . . `*)`
non-star-close ::= *non-star*
| *stars-non-close*
stars-non-close ::= `*` [`*`] . . . *non-close*
non-close ::= any character except `)`

A.2 Compilation Units

compilation-unit ::= [*package-definition*]
package-definition ::= `package` *name* `;`
import-definition ::= [`model`] `import` *name-star* `;`
name ::= *ident* [`.` *ident*] . . .
name-star ::= *ident* [`.` *ident*] . . . [`.` `*`]

A.3 Type Definitions

type-definition ::= [*doc-comment*] *modifiers class-or-interface-def*
| `;`
class-or-interface-def ::= *class-definition* | *interface-definition*
type-spec ::= *type* [*dims*] | `\TYPE` [*dims*]
type ::= *reference-type* | *builtInType*
reference-type ::= *name*
modifiers ::= [*modifier*] . . .
modifier ::= `public` | `private` | `protected`
| `spec_public` | `spec_protected`
| `abstract` | `static` |
| `model` | `ghost` | `pure`
| `final` | `synchronized`
| `instance` | `helper`
| `transient` | `volatile`
| `native` | `strictfp`
| `const`[1]
| `non_null`
class-definition ::= `class` *ident* [`extends` *name* [`weakly`]]
interface-definition ::= `interface` *ident* [*interface-extends*] *class-block*
interface-extends ::= `extends` *name-weakly-list*
implements-clause ::= `implements` *name-weakly-list*
name-weakly-list ::= *name* [`weakly`] [`,` *name* [`weakly`]] . . .
class-block ::= `{` [*field*] . . . `}`

A.4 Field declarations

field ::= [*doc-comment*] . . . *modifiers member-decl*
| *modifiers jml-declaration*
| [*method-specification*] [`static`] *compound-statement*
| *method-specification* `static_initializer`
| *method-specification* `initializer`
| `axiom` *predicate* `;`
| `;`
member-decl ::= *variable-decls* | *method-decl*
| *class-definition* | *interface-definition*
variable-decls ::= [`field`] *type-spec variable-declarators* `;` [*jml-data-group-clause*] . . .
variable-declarators ::= *variable-declarator* [`,` *variable-declarator*] . . .
variable-declarator ::= *ident* [*dims*] [`=` *initializer*]
initializer ::= *expression* | *array-initializer*
array-initializer ::= `{` [*initializer-list*] `}`
initializer-list ::= *initializer* [`,` *initializer*] . . . [`,`]

[1] `const` is reserved but not used in Java

method-decl ::= *method-specification*
| *method-or-constructor-keyword*
method-or-constructor-keyword ::= `method` | `constructor`
method-head ::= *ident* `(` [*param-declaration-list*] `)`
method-body ::= *compound-statement* | `;`
throws-clause ::= `throws` *name* [`,` *name*] ...
param-declaration-list ::= *param-declaration* [`,` *param-declaration*] ...
param-declaration ::= [*param-modifier*] ... *type-spec ident* [*dims*]
param-modifier ::= `final` | `non_null`

A.5 Methods and constructors

A.6 Other elements of type declarations

A.7 Type Specifications

jml-declaration ::= *invariant* | *history-constraint*
| *represents-decl* | *initially-clause*
| *monitors-for-clause* | *readable-if-clause*
invariant ::= *invariant-keyword predicate* `;`
invariant-keyword ::= `invariant` | `invariant_redundantly`
history-constraint ::= *constraint-keyword predicate*
constraint-keyword ::= `constraint` | `constraint_redundantly`
constrained-list ::= *method-name-list* | `\everything`
method-name-list ::= *method-name* [`,` *method-name*] ...
method-name ::= *method-ref* [`(` [*param-disambig-list*] `)`]
method-ref ::= *method-ref-start* [`.` *method-ref-rest*] ...
| `new` *reference-type*
method-ref-start ::= `super` | `this` | *ident* | `\other`
method-ref-rest ::= `this` | *ident* | `\other`
param-disambig-list ::= *param-disambig* [`,` *param-disambig*] ...
param-disambig ::= *type-spec* [*ident* [*dims*]]
represents-decl ::= *represents-keyword store-ref-expression l-arrow-or-eq spec-expression* `;`
| *represents-keyword store-ref-expression* `\such_that` *predicate* `;`
represents-keyword ::= `represents` | `represents_redundantly`
l-arrow-or-eq ::= `<-` | `=`
initially-clause ::= `initially` *predicate* `;`
readable-if-clause ::= `readable` *ident* `if` *predicate* `;`
monitors-for-clause ::= `monitors_for` *ident l-arrow-or-eq spec-expression-list* `;`

A.8 Method Specifications

method-specification ::= *specification* | *extending-specification*
extending-specification ::= `also` *specification*
specification ::= *spec-case-seq* [*redundant-spec*]
| *redundant-spec*
spec-case-seq ::= *spec-case* [`also` *spec-case*] . . .
spec-case ::= *lightweight-spec-case* | *heavyweight-spec-case*
| *model-program* | *code-contract-spec*
privacy ::= `public` | `protected` | `private`
lightweight-spec-case ::= *generic-spec-case*
generic-spec-case ::= [*spec-var-decls*] [*spec-header*] *simple-spec-body*
| [*spec-var-decls*] [*spec-header*] `{|` *generic-spec-case-seq* `|}`
generic-spec-case-seq ::= *generic-spec-case* [`also` *generic-spec-case*] . . .
spec-header ::= *requires-clause*
simple-spec-body ::= *simple-spec-body-clause* [*simple-spec-body-clause*] . . .
simple-spec-body-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause*
| *duration-clause* | *ensures-clause* | *signals-clause*
heavyweight-spec-case ::= *behavior-spec-case*
| *exceptional-behavior-spec-case*
| *normal-behavior-spec-case*
behavior-spec-case ::= [*privacy*] `behavior` *generic-spec-case*
normal-behavior-spec-case ::= [*privacy*] `normal_behavior` *normal-spec-case*
normal-spec-case ::= [*spec-var-decls*] *normal-spec-body*
| [*spec-var-decls*] [*spec-header*] `{|` *normal-spec-case-seq* `|}`
normal-spec-case-seq ::= *normal-spec-case* [`also` *normal-spec-case*] . . .
normal-spec-body ::= *normal-spec-clause* [*normal-spec-clause*] . . .
normal-spec-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause*
| *duration-clause* | *ensures-clause*
exceptional-behavior-spec-case ::= [*privacy*] `exceptional_behavior` *exceptional-spec-case*
exceptional-spec-case ::= [*spec-var-decls*] *exceptional-spec-body*
| [*spec-var-decls*] [*spec-header*] `{|` *exceptional-spec-case-seq* `|}`
exceptional-spec-case-seq ::= *exceptional-spec-case* [`also` *exceptional-spec-case*] . . .
exceptional-spec-body ::= *exceptional-spec-clause* [*exceptional-spec-clause*] . . .
exceptional-spec-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause*
| *duration-clause* | *signals-clause*
spec-var-decls ::= *forall-var-decls* [*old-var-decls*]
| *old-var-decls*
forall-var-decls ::= *forall-var-decl* [*forall-var-decl*] . . .

forall-var-decl ::= `forall` *quantified-var-decl* `;`
old-var-decls ::= `old-var-decl` [*old-var-decl*] . . .
old-var-decl ::= `old` *type-spec spec-variable-declarators* `;`
requires-clause ::= *requires-keyword pred-or-not* `;`
requires-keyword ::= `requires` | `pre`
| `requires_redundantly` | `pre_redundantly`
pred-or-not ::= *predicate* | `\not_specified`
ensures-clause ::= *ensures-keyword pred-or-not* `;`
ensures-keyword ::= `ensures` | `post`
| `ensures_redundantly` | `post_redundantly`
signals-clause ::= *signals-keyword*
signals-keyword ::= `signals` | `signals_redundantly`
| `exsures` | `exsures_redundantly`
diverges-clause ::= *diverges-keyword pred-or-not* `;`
diverges-keyword ::= `diverges` | `diverges_redundantly`
when-clause ::= *when-keyword pred-or-not* `;`
when-keyword ::= `when` | `when_redundantly`
assignable-clause ::= *assignable-keyword conditional-store-ref-list* `;`
assignable-keyword ::= `assignable` | `assignable_redundantly`
| `modifiable` | `modifiable_redundantly`
| `modifies` | `modifies_redundantly`
conditional-store-ref-list ::= *conditional-store-ref*
conditional-store-ref ::= *store-ref* [`if` *predicate*]
| *other-ref* [`if` *predicate*]
other-ref ::= `\other` [*store-ref-name-suffix*] . . .
working-space-clause ::= *working-space-keyword* `\not_specified` `;`
| *working-space-keyword spec-expression* [`if` *predicate*] `;`
working-space-keyword ::= `working_space` | `working_space_redundantly`
duration-clause ::= *duration-keyword* `\not_specified` `;`
| *duration-keyword spec-expression* [`if` *predicate*] `;`
duration-keyword ::= `duration` | `duration_redundantly`
measured-clause ::= *measured-by-keyword* `\not_specified` `;`
| *measured-by-keyword spec-expression* [`if` *predicate*] `;`
measured-by-keyword ::= `measured_by` | `measured_by_redundantly`
accessible-clause ::= *accessible-keyword conditional-store-ref-list* `;`
accessible-keyword ::= `accessible` | `accessible_redundantly`
callable-clause ::= *callable-keyword callable-methods-list* `;`
callable-keyword ::= `callable` | `callable_redundantly`
callable-methods-list ::= *method-name-list* | *store-ref-keyword*

A.9 Frame Conditions and Data Groups

jml-data-group-clause ::= *in-group-clause* | *maps-into-clause*
in-group-clause ::= *in-keyword group-list* `;`
in-keyword ::= `in` | `in_redundantly`

group-list ::= *group-name* [`,` *group-name*] . . .
group-name ::= [*group-name-prefix*] *ident*
group-name-prefix ::= `super .` | `this .`
maps-into-clause ::= *maps-keyword member-field-ref* `\into` *group-list* `;`
maps-keyword ::= `maps` | `maps_redundantly`
member-field-ref ::= *ident* `.` *maps-member-ref-expr*
| *maps-array-ref-expr* [`.` *maps-member-ref-expr*]
maps-member-ref-expr ::= *ident* | `*`
maps-array-ref-expr ::= *ident maps-spec-array-dim* [*maps-spec-array-dim*] . . .
maps-spec-array-dim ::= '`[`' *spec-array-ref-expr* '`]`'

A.10 Predicates and Specification Expressions

predicate ::= *spec-expression*
spec-expression-list ::= *spec-expression* [`,` *spec-expression*] . . .
spec-expression ::= *expression*
jml-primary ::= `\result`
| `\old (` *spec-expression* `)`
| `\not_modified (` *store-ref-list* `)`
| `\fresh (` *spec-expression-list* `)`
| `\reach (` *spec-expression* `)`
| `\duration (` *expression* `)`
| `\space (` *spec-expression* `)`
| `\max (` *spec-expression* `)`
| `\working_space (` *expression* `)`
| *informal-description*
| `\nonnullelements (` *spec-expression* `)`
| `\typeof (` *spec-expression* `)`
| `\elemtype (` *spec-expression* `)`
| `\type (` *type* `)`
| `\lockset`
| `\is_initialized (` *reference-type* `)`
| `\invariant_for (` *spec-expression* `)`
| `( \lblneg` *ident spec-expression* `)`
| `( \lblpos` *ident spec-expression* `)`
| *spec-quantified-expr*
set-comprehension ::= `{` *type-spec quantified-var-declarator*
set-comprehension-pred ::= *postfix-expr* `. has (` *ident* `)`
spec-quantified-expr ::= `(` *quantifier quantified-var-decls* `;` [[*predicate*] `;`]
quantifier ::= `\forall` | `\exists` | `\max` | `\min` | `\num_of` | `\product` | `\sum`
quantified-var-decls ::= *type-spec quantified-var-declarator*
quantified-var-declarator ::= *ident* [*dims*]
spec-variable-declarators ::= *spec-variable-declarator*
spec-variable-declarator ::= *ident* [*dims*] [`=` *spec-initializer*]
spec-array-initializer ::= `{` [*spec-initializer*

spec-initializer ::= *spec-expression*
| *spec-array-initializer*
expression-list ::= *expression* [`,` *expression*] . . .
expression ::= *assignment-expr*
assignment-expr ::= *conditional-expr* [*assignment-op assignment-expr*]
assignment-op ::= `=` | `+=` | `-=` | `*=` | `/=` | `%=` | `>>=`
| `>>>=` | `<<=` | `&=` | `'|='` | `^=`
conditional-expr ::= *equivalence-expr*
equivalence-expr ::= *implies-expr* [*equivalence-op implies-expr*] . . .
equivalence-op ::= `<==>` | `<=!=>`
implies-expr ::= *logical-or-expr*
| *logical-or-expr* `<==` *logical-or-expr*
implies-non-backward-expr ::= *logical-or-expr*
logical-or-expr ::= *logical-and-expr* [`'||'` *logical-and-expr*] . . .
logical-and-expr ::= *inclusive-or-expr* [`&&` *inclusive-or-expr*] . . .
inclusive-or-expr ::= *exclusive-or-expr* [`'|'` *exclusive-or-expr*] . . .
exclusive-or-expr ::= *and-expr* [`^` *and-expr*] . . .
and-expr ::= *equality-expr* [`&` *equality-expr*] . . .
equality-expr ::= *relational-expr* [`==` *relational-expr*] . . .
| *relational-expr* [`!=` *relational-expr*] . . .
relational-expr ::= *shift-expr* `<` *shift-expr*
| *shift-expr* `>` *shift-expr*
| *shift-expr* `<=` *shift-expr*
| *shift-expr* `>=` *shift-expr*
| *shift-expr* `<:` *shift-expr*
| *shift-expr* [`instanceof` *type-spec*]
shift-expr ::= *additive-expr* [*shift-op additive-expr*] . . .
shift-op ::= `<<` | `>>` | `>>>`
additive-expr ::= *mult-expr* [*additive-op mult-expr*] . . .
additive-op ::= `+` | `-`
mult-expr ::= *unary-expr* [*mult-op unary-expr*] . . .
mult-op ::= `*` | `/` | `%`
unary-expr ::= `(` *type-spec* `)` *unary-expr*
| `++` *unary-expr*
| `--` *unary-expr*
| `+` *unary-expr*
| `-` *unary-expr*
| *unary-expr-not-plus-minus*
unary-expr-not-plus-minus ::= `~` *unary-expr*
| `!` *unary-expr*
| `(` *builtinType* `)` *unary-expr*
| `(` *reference-type* `)` *unary-expr-not-plus-minus*
| *postfix-expr*
postfix-expr ::= *primary-expr* [*primary-suffix*] . . . [`++`]
| *primary-expr* [*primary-suffix*] . . . [`--`]
| *builtinType* [`'['` `']'`] . . . `.` `class`
primary-suffix ::= `.` *ident*

```
        | . this
        | . class
        | . new-expr
        | . super ( [ expression-list ] )
        | ( [ expression-list ] )
        | '[' expression ']'
        | [ '[' ']' ] ... . class
primary-expr ::= ident | new-expr
        | constant | super | true
        | false | this | null
        | ( expression )
        | jml-primary
        | informal-description
builtInType ::= void | boolean | byte
        | char | short | int
        | long | float | double
constant ::= java-literal
new-expr ::= new type new-suffix
new-suffix ::= ( [ expression-list ] ) [ class-block ]
        | array-decl [ array-initializer ]
        | set-comprehension
array-decl ::= dim-exprs [ dims ]
dim-exprs ::= '[' expression ']' [ '[' expression ']' ] ...
dims ::= '[' ']' [ '[' ']' ] ...
array-initializer ::= { [ initializer [ , initializer ] ... [ , ] ] }
initializer ::= expression
        | array-initializer
store-ref-list ::= store-ref [ , store-ref ] ...
store-ref ::= store-ref-expression
        | informal-description
        | store-ref-keyword
store-ref-expression ::= store-ref-name [ store-ref-name-suffix ] ...
store-ref-name ::= ident | super | this
store-ref-name-suffix ::= . ident | . this | '[' spec-array-ref-expr ']' | . *
spec-array-ref-expr ::= spec-expression
        | spec-expression .. spec-expression
        | *
store-ref-keyword ::= \nothing | \everything | \not_specified | \private_data
```

A.11 JML primitive types

A.12 Statements and Annotation Statements

compound-statement ::= `{` *statement* [*statement*] . . . `}`
statement ::= *compound-statement*
| *local-declaration* `;`
| *ident* `:` *statement*
| *expression* `;`
| `if (` *expression* `)` *statement* [`else` *statement*]
| [*loop-invariant*] . . . [*variant-function*] . . . [*ident*] `:` *loop-stmt*
| `break` [*ident*] `;`
| `continue` [*ident*] `;`
| `return` [*expression*] `;`
| *switch-statement*
| *try-block*
| `throw` *expression* `;`
| `synchronized (` *expression* `)` *statement*
| `;`
| *assert-statement*
| *hence-by-statement*
| *assert-redundantly-statement*
| *assume-statement*
| *set-statement*
| *unreachable-statement*
| *model-prog-statement* // only allowed in model programs
loop-stmt ::= `while (` *expression* `)` *statement*
| `do` *statement* `while (` *expression* `) ;`
| `for (` [*for-init*] `;` [*expression*] `;` [*expression-list*] `)`
for-init ::= *local-declaration* | *expression-list*
local-declaration ::= *local-modifiers variable-decls*
local-modifiers ::= [*local-modifier*] . . .
local-modifier ::= `model | ghost | final | non_null`
switch-statement ::= `switch (` *expression* `) {` [*switch-body*] . . . `}`
switch-body ::= *switch-label-seq* [*statement*] . . .
switch-label-seq ::= *switch-label* [*switch-label*] . . .
switch-label ::= `case` *expression* `:` | `default :`
try-block ::= `try` *compound-statement* [*handler*] . . .
handler ::= `catch (` *param-declaration* `)` *compound-statement*
assert-statement ::= `assert` *expression* [`:` *expression*] `;`
hence-by-statement ::= *hence-by-keyword predicate* `;`
hence-by-keyword ::= `hence_by | hence_by_redundantly`
assert-redundantly-statement ::= `assert_redundantly` *predicate* [`:` *expression*] `;`
assume-statement ::= *assume-keyword predicate* [`:` *expression*] `;`
assume-keyword ::= `assume | assume_redundantly`
set-statement ::= `set` *assignment-expr* `;`
unreachable-statement ::= `unreachable ;`
loop-invariant ::= *maintaining-keyword predicate* `;`
maintaining-keyword ::= `maintaining | maintaining_redundantly`
| `loop_invariant | loop_invariant_redundantly`
variant-function ::= *decreasing-keyword spec-expression* `;`

decreasing-keyword ::= `decreasing` | `decreasing_redundantly`
| `decreases` | `decreases_redundantly`

A.13 Redundancy

redundant-spec ::= *implications* [*examples*] | *examples*
examples ::= `for_example` *example* [`also` *example*] . . .
implications ::= `implies_that` *spec-case-seq*
example ::= [[*privacy*] `example`]
| [*privacy*] `exceptional_example`
| [*privacy*] `exceptional_example`
| [*privacy*] `normal_example`
| [*privacy*] `normal_example`
exceptional-example-body ::= *exceptional-spec-clause* [*exceptional-spec-clause*] . . .
normal-example-body ::= *normal-spec-clause* [*normal-spec-clause*] . . .

A.14 Model Programs

model-program ::= [*privacy*] `model_program` *jml-compound-statement*
jml-compound-statement ::= *compound-statement*
jml-statement ::= *statement*
model-prog-statement ::= *nondeterministic-choice*
| *nondeterministic-if*
| *spec-statement*
| *invariant*
nondeterministic-choice ::= `choose` *alternative-statements*
alternative-statements ::= *jml-compound-statement*
nondeterministic-if ::= `choose_if` *guarded-statements*
guarded-statements ::= *guarded-statement*
guarded-statement ::= `{`
spec-statement ::= [*privacy*] `behavior` *generic-spec-statement-case*
| [*privacy*] `exceptional_behavior` *exceptional-spec-case*
| [*privacy*] `normal_behavior` *normal-spec-case*
| [*privacy*] `abrupt_behavior` *abrupt-spec-case*
generic-spec-statement-case ::= [*spec-var-decls*] *generic-spec-statement-body*
| [*spec-var-decls*] *spec-header* [*generic-spec-statement-body*]
generic-spec-statement-body ::= *simple-spec-statement-body*
| `{|` *generic-spec-statement-case-seq* `|}`
generic-spec-statement-body-seq ::= *generic-spec-statement-case*
simple-spec-statement-body ::= *simple-spec-statement-clause* [*simple-spec-statement-clause*] . . .
simple-spec-statement-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause* | *duration-clause*

| *ensures-clause* | *signals-clause*
| *continues-clause* | *breaks-clause* | *returns-clause*

abrupt-spec-case ::= [*spec-var-decls*] *spec-header*
| [*spec-var-decls*] *abrupt-spec-body*

abrupt-spec-body ::= *abrupt-spec-clause* [*abrupt-spec-clause*] . . .
| `{|` *abrupt-spec-case-seq* `|}`

abrupt-spec-clause ::= *diverges-clause*
| *assignable-clause*
| *when-clause* | *working-space-clause* | *duration-clause*
| *continues-clause* | *breaks-clause* | *returns-clause*

abrupt-spec-case-seq ::= *abrupt-spec-case* [`also` *abrupt-spec-case*] . . .

continues-clause ::= *continues-keyword* [*target-label*] [*pred-or-not*] `;`

continues-keyword ::= `continues` | `continues_redundantly`

target-label ::= `->` `(` *ident* `)`

breaks-clause ::= *breaks-keyword* [*target-label*] [*pred-or-not*] `;`

breaks-keyword ::= `breaks` | `breaks_redundantly`

returns-clause ::= *returns-keyword* [*pred-or-not*] `;`

returns-keyword ::= `returns` | `returns_redundantly`

A.15 Specification for Subtypes

code-contract-spec ::= `code_contract` *code-contract-clause* [*code-contract-clause*] . . .

code-contract-clause ::= *accessible-clause* | *callable-clause* | *measured-clause*

A.16 Refinement

refine-prefix ::= `refine` *string-literal* `;`

depends-keyword ::= `depends` | `depends_redundantly`

depends-decl ::= *depends-keyword*

jml-var-assertions ::= *jml-var-assertion* [`;` *jml-var-assertion*] . . .

jml-var-assertion ::= `readable_if` *predicate*
| `monitored_by` *spec-expression-list*

store-ref ::= `\fields_of` `(` *spec-expression* [`,` *reference-type* [`,` *store-ref-expression*]] `)`

jml-primary ::= `\reach` `(` *spec-expression* `,` *reference-type* [`,` *store-ref-expression*] `)`

Appendix B Modifier Summary

This table summarizes which Java and JML modifiers may be used in various grammatical contexts.

Grammatical construct	Java modifiers	JML modifiers

All modifiers	`public`
	`protected`
	`private`
	`abstract`
	`static final`
	`synchronized`
	`transient`
	`volatile`
	`native`
	`strictfp`
		`spec_public`
		`spec_protected`
		`model ghost`
		`pure instance`
		`helper`
		`non_null`

Class declaration	`public final`
	`abstract`
	`strictfp`
		`pure model`

Interface declaration	`public` `strictfp`	`pure model`

Nested Class declaration	`public`
	`protected`
	`private`
	`static final`
	`abstract`
	`strictfp`
		`spec_public`
		`spec_protected`
		`model pure`

Nested interface declaration	`public`
	`protected`
	`private`
	`static`
	`strictfp`
		`spec_public`
		`spec_protected`
		`model pure`

Local Class (and local model class) declaration	`final`
	`abstract`
	`strictfp`
		`pure model`

Type specification (e.g. invariant)	`public`
	`protected`
	`private`
		-

Field declaration	`public`	`spec_public`
	`protected`	`spec_protected`
	`private final`	`non_null`
	`volatile`	`instance`
	`transient`	
	`static`	
Ghost Field declaration	`public`	`non_null`
	`protected`	`instance`
	`private`	
	`static final`	
Model Field declaration	`public`	`non_null`
	`protected`	`instance`
	`private`	
	`static`	
Method declaration	`public`	`spec_public`
	`protected`	`spec_protected`
	`private`	`pure non_null`
	`abstract`	`helper`
	`final static`	
	`synchronized`	
	`native`	
	`strictfp`	
Constructor declaration	`public`	`spec_public`
	`protected`	`spec_protected`
	`private`	`helper pure`
Model method	`public`	`pure non_null`
	`protected`	`helper`
	`private`	
	`abstract`	
	`static final`	
	`synchronized`	
	`strictfp`	
Model constructor	`public`	`pure helper`
	`protected`	
	`private`	
Java Initialization block	`static` ???	???

JML initializer annotation	`static` ???	???
Formal parameter	`final`	`non_null`
Local variable and local ghost variable declaration	`final`	`ghost non_null uninitialized`

Note that within interfaces, fields are implicitly public, static and final. Ghost and model fields are implicitly public and static, though they may be declared instance (i.e. not static).

Appendix C Type Checking Summary

[[[Hope to generate this automatically]]]

Appendix D Verification Logic Summary

[[[Hope to generate this automatically]]]

Appendix E Differences

The subsections below detail the differences between JML and other tools. and between JML and Java itself.

E.1 Differences Between JML and Other Tools

ESC/Java and JML share a common syntax thanks to the efforts of Raymie Stata who initiated the effort to bring their syntaxes together. After a long process, the syntax of ESC/Java and JML were both changed and JML is nearly a superset of ESC/Java now. This section discusses the current state of affairs, as of JML 3.5 and ESC/Java 1.2.4.

The following differences remain between ESC/Java and JML, but could, we hope be eliminated by changing ESC/Java to be more compatible with JML. (This is starting to happen.)

- ESC/Java has `ghost` variables which are concrete, specification-only variables. In ESC/Java, ghost variables must be public, and they seem to be scoped differently than normal variables, while in JML the scope of a ghost variable declaration is exactly the same as for a normal Java declaration. ESC/Java should change to make the scope of ghost variable declarations the same scoping as normal Java declarations. It would also be convenient to allow local ghost variables, as JML does, and to allow non-public ghost variables.
- ESC/Java changed to support the following syntax for loop invariants and termination functions:

  ```
  //@ loop_invariant ...loop_inv...
  //@ decreases ...termination_fun...
  while (...cond...) do { ...statements... }
  ```

  JML supports this and also the following synonyms (adapted from RESOLVE [Edwards-etal94])

  ```
  //@ maintaining ...loop_inv...
  //@ decreasing ...termination_fun...
  while (...cond...) do { ...statements... }
  ```

  (Completed by Joe Kiniry as of April 2003.)
- ESC/Java has not yet adopted JML's syntax for method specifications that includes the use of `also` as a separate keyword, (instead of ESC/Java's `also_requires`, `also_ensures` keywords, which JML does not support) and the use of `{|` and `|}` as brackets for *spec-case* sequences. ESC/Java also does not parse JML's heavyweight method specifications, and does not understand (and does not need to understand) the `implies_that` and `for_example` sections of a method specification. (Completed by Joe Kiniry and David Cok as of Spring 2003.)
- ESC/Java does not support the following synonyms that JML supports.

  ```
  ESC/Java      JML allowed
  keyword       synonyms
  ----------------------
  requires      pre
  ```

```
modifies      assignable, modifiable
ensures       post
exsures       signals
```

(Completed by Joe Kiniry as of April 2003.)

- ESC/Java does not parse JML's `initially` clauses on declarations.
- ESC/Java does not require semicolons where they are required by JML. Currently semicolons at the end of an annotation comment are optional in ESC/Java.
- JML allows privacy modifiers (public, protected, private) on invariants, and enforces them; for example, public invariants cannot mention protected or private names. (Completed by David Cok as of April 2003.)
- JML has the modifier `pure`, which can be used to declare classes and methods that have no side effects. (Completed by David Cok and Joe Kiniry as of Spring 2003.)
- JML allows `\nothing` and `\everything` in a modifies clause. (Completed by David Cok as of Spring 2003.)
- JML allows final fields in modifies clauses in constructors.
- ESC/Java seems to have a `\max` operator which can be applied to a lockset. This conflicts with JML's `\max` quantifier. (Completed by Joe Kiniry as of July 2003.)
- ESC/Java requires bodies for methods in '`.spec`' files (Completed by David Cok as of Spring 2003.)
- JML treats `non_null` annotations as redundant, and allows them to be specified on arguments in overrides of methods. ESC/Java rejects files when this happens. It's true that it's a specification error to add `non_null` to a parameter in an override, but it should be okay to specify it redundantly. For return types, it should be okay to add it in subtypes. (Planning underway by Joe Kiniry as of January 2003.) ESC/Java does not permit `non_null` annotations on method return types.
- ESC/Java has lexical pragmas of the form `nowarn` *labels*`;`. JML requires a semicolon at the end of such pragmas. (Completed by Joe Kiniry as of Spring 2003.)
- ESC/Java requires a whole syntactic construct in a `//@` comment, whereas in JML the construct can be split over multiple lines. (This is theoretically the case now in JML, but is not enforced as in the JML parser `//@` is simply ignored, as are `/*@` and `@*/`.) (Underway by Joe Kiniry and David Cok as of Spring 2003.)

The following differences between ESC/Java and JML are designed to remain differences, due to the differing goals of the two languages. While it would be nice if ESC/Java could parse and ignore all of JML's syntax, but that is not currently planned. The first two differences below allow users of both tools to deal with this problem.

- JML supports annotation forms `//+@` and `/*+@ ... @+*/`, so that annotations that JML understands but ESC/Java doesn't can be written.
- JML allows specifications to appear within javadoc comments in the form

```
/** ...
 <pre><JML>
  ...
 </JML></pre>
*/
```

which, of course, will appear before the method signature. ESC/Java already has such comments of the form `<esc>` ... `</esc>`. ESC/Java will not recognize the `<JML>`, `<ESC>`, or `<jml>` forms so that specification forms that JML understands but ESC/Java doesn't can be written in documentation comments.

- JML uses the file suffixes '`.jml`', '`.jml-refined`', '`.java-refined`', '`.refines-jml`', and '`.refines-java`' files as well as '`.java`' files.
- JML has dependent (abstract) model variables.
- JML has model classes, interfaces, or methods.
- JML has in and maps-into data group declarations.
- JML has history constraints (`constraint`).
- JML has other refinement calculus and annotation primitives besides the `assert` and `assume` that are in ESC/Java, that can appear only in model programs.
- JML allows pure method invocations in assertions.
- JML allows new expressions (with pure constructors) in assertions.
- JML has `working_space` and `duration` clauses for space and time specifications.
- JML has `measured_by` clauses to give termination functions for recursive methods and specifications.
- JML has a notion of refinement of specifications with a `refine` keyword.
- JML has model import declarations.
- JML has when clauses in specifications for concurrency.
- JML method specifications may have an implications (`implies_that`) and examples (`for_example`) section.

E.2 Differences Between JML and Java

Appendix F Deprecated

The subsections below briefly describe the deprecated features of JML. A feature is *deprecated* if it is supported in the current release, but slated to be removed from a subsequent release. Such features should not be used.

F.1 Deprecated Syntax

The following syntax is deprecated. (Note that incompatible changes and syntax that is no longer supported is not included in this list.)

The syntax of *jml-var-assertions* was used in declarations. The former `initially` clause in a *jml-var-assertion* has been deleted entirely, as it caused a grammar conflict. It has been replaced by the `initially` clause at the top-level of a type declaration.

The `readable_if` and `monitored_by` clauses have been replaced by their own clauses, which appear at the top-level in a type declaration.

The `\fields_of` store reference has been replaced by data groups and the `x.*` expression.

The `\reach` expression has been simplified because the more complicated syntax (expressions with more than one argument) can be replaced by expressions that reference data groups if necessary. Also, one of the main purposes of the complex syntax was to create sets of objects for the `\fields_of` expression, a construct that is going to be removed in a future release.

The `depends` clause has also been replaced by data group clauses, i.e., the `maps` and `in` clauses.

depends-keyword ::= `depends` | `depends_redundantly`
depends-decl ::= *depends-keyword*
store-ref-expression l_arrow-or-eq store-ref-list `;`
jml-var-assertions ::= *jml-var-assertion* [`;` *jml-var-assertion*] . . .
jml-var-assertion ::= `readable_if` *predicate*
| `monitored_by` *spec-expression-list*
store-ref ::= `\fields_of` `(` *spec-expression* [`,` *reference-type* [`,` *store-ref-expression*]] `)`
jml-primary ::= `\reach` `(` *spec-expression* `,` *reference-type* [`,` *store-ref-expression*] `)`

Appendix G What's Missing

What is missing from this reference manual?

The following constructs are not discussed at all:

- `\other`
- `\private_data`
- `\such_that`
- `abrupt_behavior`
- `breaks` and `breaks_redundantly`
- `callable` and `callable_redundantly`
- `choose` and `choose_if`
- `continues` and `continues_redundantly`
- `in` and `in_redundantly`
- `maps` and `maps_redundantly`
- `example` and `exceptional_example`
- `forall`
- `implies_that`
- `hence_by` and `hence_by_redundantly`
- `measured_by` and `measured_by_redundantly`
- `model_program`
- `old` – In fact, it is unclear if `old` is a separate keyword from `\old`.
- `represents` and `represents_redundantly`
- `returns` and `returns_redundantly`
- `weakly` xxx
- The fact that `if` is a keyword is not discussed at all.

Other stuff not to forget - DRCok

- \not_specified
- \private_data
- \nothing
- \everything
- nowarn annotation
- methods and constructors without bodies in java files
- methods and constructors with bodies in specification files
- methods and constructors in annotation expressions - purity - modifies clauses - various checking
- anonymous and block-level classes
- field, method, constructor keywords
- exceptions in annotation expressions
- equivalence of \TYPE and java.lang.Class
- inheritance of non_null

Bibliography

[America87]
Pierre America. Inheritance and Subtyping in a Parallel Object-Oriented Language. In Jean Bezivin and others (eds.), *ECOOP '87, European Conference on Object-Oriented Programming, Paris, France*. Lecture Notes in Computer Science, Vol. 276 (Springer-Verlag, NY), pages 234-242.

[Arnold-Gosling-Holmes00]
Ken Arnold, James Gosling, and David Holmes. *The Java Programming Language Third Edition*. The Java Series. Addison-Wesley, Reading, MA, 2000.

[ANSI95] *Working Paper for Draft Proposed International Standard for Information Systems — Programming Language Java*. CBEMA, 1250 Eye Street NW, Suite 200, Washington DC 20005, April 28, 1995. (Obtained by anonymous ftp to research.att.com, directory dist/c++std/WP.)

[Back88] R. J. R. Back. A calculus of refinements for program derivations. *Acta Informatica*, 25(6):593-624, August 1988.

[Back-vonWright89a]
R. J. R. Back and J. von Wright. Refinement Calculus, Part I: Sequential Nondeterministic Programs. In J. W. de Bakker, et al, (eds.), *Stepwise Refinement of Distributed Systems, Models, Formalisms, Correctness, REX Workshop*, Mook, The Netherlands, May/June 1989, pages 42-66. Volume 430 of *Lecture Notes Computer Science*, Spring-Verlag, 1989.

[Back-vonWright98]
Ralph-Johan Back and Joakim von Wright. *Refinement Calculus: A Systematic Introduction*. Springer-Verlag, 1998.

[Borgida-etal95]
Alex Borgida, John Mylopoulos, and Raymond Reiter. On the Frame Problem in Procedure Specifications. *IEEE Transactions on Software Engineering*, 21(10):785-798, October 1995.

[Buechi-Weck00]
Martin Büchi and Wolfgang Weck. The Greybox Approach: When Blackbox Specifications Hide Too Much. Technical Report 297, Turku Centre for Computer Science, August 1999.
'`http://www.tucs.abo.fi/publications/techreports/TR297.html`'.

[Buechi00] Martin Büchi. Safe Language Mechanisms for Modularization and Concurrency. Ph.D. Thesis, Turku Center for Computer Science, May 2000. TUCS Dissertations No. 28.

[Burdy-etal03]
Lilian Burdy, Yoonsik Cheon, David Cok, Michael Ernst, Joe Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. Dept. of Computer Science, University of Nijmegen, TR NIII-R0309, 2003.
'`ftp://ftp.cs.iastate.edu/pub/leavens/JML/jml-white-paper.pdf`'.

[Cheon-Leavens02]
Yoonsik Cheon and Gary T. Leavens. A Simple and Practical Approach to Unit Testing: The JML and JUnit Way. In *ECOOP 2002 – Object-Oriented Programming, 16th European Conference, Malaga, Spain*, pages 231–255. Springer-Verlag, June 2002. Also Department of Computer Science, Iowa State University, TR #01-12a, November 2001, revised March 2002 which is available from the URL
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR01-12/TR.pdf`'.

[Cheon-Leavens02b]
Yoonsik Cheon and Gary T. Leavens. A Runtime Assertion Checker for the Java Modeling Language (JML). In Hamid R. Arabnia and Youngsong Mun (eds.), *Proceedings of the International Conference on Software Engineering Research and Practice (SERP '02), Las Vegas, Nevada, USA*, pages 322–328. CSREA Press, June 2002. Also Department of Computer Science, Iowa State University, TR #02-05, March 2002 which is available from the URL
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR02-05/TR.pdf`'.

[Cheon-etal03]
Yoonsik Cheon, Gary T. Leavens, Murali Sitaraman, and Stephen Edwards. Model Variables: Cleanly Supporting Abstraction in Design By Contract. Department of Computer Science, Iowa State University, TR 03-10, March 2003.
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR03-10/TR.pdf`'.

[Cheon03] Yoonsik Cheon. A Runtime Assertion Checker for the Java Modeling Language. Department of Computer Science, Iowa State University, TR 03-09, April, 2003.
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR03-09/TR.pdf`'.

[Clifton-etal00]
Curtis Clifton, Gary T. Leavens, Craig Chambers, and Todd Millstein. MultiJava: Modular Open Classes and Symmetric Multiple Dispatch for Java. In *OOPSLA 2000 Conference on Object-Oriented Programming, Systems, Languages, and Applications, Minneapolis, Minnesota* (ACM SIGPLAN Notices, 35(10):130-145, Oct., 2000).

[Cohen90] Edward Cohen. *Programming in the 1990s: An Introduction to the Calculation of Programs*. Springer-Verlag, New York, N.Y., 1990.

[Dhara-Leavens94b]
Krishna Kishore Dhara and Gary T. Leavens. Weak Behavioral Subtyping for Types with Mutable Objects. In S. Brookes and M. Main and A. Melton and M. Mislove (eds.), *Mathematical Foundations of Programming Semantics, Eleventh Annual Conference*, '`http://www.elsevier.nl/locate/entcs/volume1.html`'. Volume 1 of *Electronic Notes in Computer Science*, Elsevier, 1995.

[Dhara-Leavens96]
Krishna Kishore Dhara and Gary T. Leavens. Forcing Behavioral Subtyping Through Specification Inheritance. In *Proceedings 18th International Conference on Software Engineering*, Berlin, Germany, pages 258-267. IEEE 1996.

An extended version is Department of Computer Science, Iowa State University, TR #95-20b, December 1995, which is available from the URL '`ftp://ftp.cs.iastate.edu/pub/techreports/TR95-20/TR.ps.Z`'.

[Dhara97] Krishna Kishore Dhara. Behavioral Subtyping in Object-Oriented Languages. Ph.D. Thesis, Department of Computer Science, Iowa State University. Technical Report TR #97-09, May 1997. Available from the URL '`ftp://ftp.cs.iastate.edu/pub/techreports/TR97-09/TR.ps.gz`'.

[Dijkstra76]
Edsger W. Dijkstra. *A Discipline of Programming* (Prentice-Hall, Englewood Cliffs, N.J., 1976).

[Edwards-etal94]
Stephen H. Edwards, Wayne D. Heym, Timothy J. Long, Murali Sitaraman, and Bruce W. Weide. Part II: Specifying Components in RESOLVE. *ACM SIGSOFT Software Engineering Notes*, 19(4):29-39 (Oct. 1994).

[Fitzgerald-Larsen98]
John Fitzgerald and Peter Gorm Larsen. *Modelling Systems: Practical Tools and Techniques in Software Development*. Cambridge University Press, Cambridge, UK, 1998.

[Gosling-etal00]
James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java Language Specification Second Edition*. The Java Series. Addison-Wesley, Boston, MA, 2000.

[Gries-Schneider95]
David Gries and Fred B. Schneider. Avoiding the Undefined by Underspecification. In Jan van Leeuwen, editor, *Computer Science Today: Recent Trends and Developments*, volume 1000 of *Lecture Notes in Computer Science*, pages 366–373. Springer-Verlag, New York, N.Y., 1995.

[Guttag-Horning-Wing85b]
John V. Guttag and James J. Horning and Jeannette M. Wing. The Larch Family of Specification Languages. *IEEE Software*, 2(5):24-36 (Sept. 1985).

[Guttag-Horning93]
John V. Guttag and James J. Horning with S.J. Garland, K.D. Jones, A. Modet and J.M. Wing. *Larch: Languages and Tools for Formal Specification* (Springer-Verlag, NY, 1993). (The ISBN numbers are 0-387-94006-5 and 3-540-94006-5.) The traits in Appendix A (the "Handbook") are found on-line at the following URL
'`http://www.research.digital.com/SRC/larch/`'.

[Hall90] Anthony Hall. Seven Myths of Formal Methods. *IEEE Software*, 7(5):11-19 (Sept. 1990).

[Hayes93] I. Hayes (ed.), *Specification Case Studies*, second edition (Prentice-Hall, Englewood Cliffs, N.J., 1990).

[Hesselink92]
Wim H. Hesselink. *Programs, Recursion, and Unbounded Choice* (Cambridge University Press, Cambridge, UK, 1992).

[Hoare69] C. A. R. Hoare. An Axiomatic Basis for Computer Programming. *Comm. ACM*, 12(10):576-583 (Oct. 1969).

[Hoare72a]
C. A. R. Hoare. Proof of correctness of data representations. *Acta Informatica*, 1(4):271-281 (1972).

[Huisman01]
Marieke Huisman. Reasoning about JAVA programs in higher order logic with PVS and Isabelle. IPA dissertation series, 2001-03. Ph.D. dissertation, University of Nijmegen, 2001.

[ISO96] International Standards Organization. *Information Technology - Programming Languages, Their Environments and System Software Interfaces - Vienna Development Method - Specification Language - Part 1: Base language*. International Standard ISO/IEC 13817-1, December, 1996.

[Jones90] Cliff B. Jones. *Systematic Software Development Using VDM*. International Series in Computer Science. Prentice Hall, Englewood Cliffs, N.J., second edition, 1990.

[Jones95e] C.B. Jones, Partial functions and logics: A warning. *Information Processing Letters*, 54(2):65-67, 1995.

[Kiczales-Lamping92]
Gregor Kiczales and John Lamping. Issues in the Design and Documentation of Class Libraries. In Andreas Paepcke (ed.), *OOPSLA '92 Proceedings* (ACM SIGPLAN Notices, 27(10):435-451, Oct., 1992).

[Krone-Ogden-Sitaraman03]
Joan Krone, William F. Ogden, Murali Sitaraman. Modular Verification of Performance Constraints. Technical Report RSRG-03-04, Department of Computer Science, Clemson University, May, 2003. Available from '`http://www.cs.clemson.edu/~resolve/reports/RSRG-03-04.pdf`'

[Lamport89]
Leslie Lamport. A Simple Approach to Specifying Concurrent Systems. *CACM*, 32(1):32-45 (Jan. 1989).

[LeavensLarchFAQ]
Gary T. Leavens. Larch frequently asked questions. Version 1.110. Available in '`http://www.cs.iastate.edu/~leavens/larch-faq.html`', May 2000.

[Leavens-Baker99]
Gary T. Leavens and Albert L. Baker. Enhancing the pre- and postcondition technique for more expressive specifications. In Jeannette M. Wing, Jim Woodcock, and Jim Davies, editors, *FM'99 — Formal Methods: World Congress on Formal Methods in the Development of Computing Systems, Toulouse, France, September 1999, Proceedings*, volume 1709 of *Lecture Notes in Computer Science*, pages 1087–1106. Springer-Verlag, 1999.

[Leavens-Baker-Ruby99]
Gary T. Leavens, Albert L. Baker, and Clyde Ruby. JML: a Notation for Detailed Design. In Haim Kilov, Bernhard Rumpe, and Ian Simmonds (editors), *Behavioral Specifications for Businesses and Systems*, chapter 12, pages 175-188.

[Leavens-Baker-Ruby02]
Gary T. Leavens, Albert L. Baker, and Clyde Ruby. Preliminary Design of JML: A Behavioral Interface Specification Language for Java Iowa State University, Department of Computer Science, TR #98-06r, August 2002, which is available from the URL
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR98-06/TR.pdf`'.

[Leavens-Dhara00]
Gary T. Leavens and Krishna Kishore Dhara. Concepts of Behavioral Subtyping and a Sketch of Their Extension to Component-Based Systems. In Gary T. Leavens and Murali Sitaraman (eds.), *Foundations of Component-Based Systems*, Cambridge University Press, 2000, pp. 113-135.
'`http://www.cs.iastate.edu/~leavens/FoCBS-book/06-leavens-dhara.pdf`'

[Leavens-Weihl90]
Gary T. Leavens and William E. Weihl. Reasoning about Object-oriented Programs that use Subtypes (extended abstract). In N. Meyrowitz (ed.), *OOPSLA ECOOP '90 Proceedings* (ACM SIGPLAN Notices, 25(10):212-223, Oct., 1990).

[Leavens-Weihl95]
Gary T. Leavens and William E. Weihl. Specification and Verification of Object-Oriented Programs Using Supertype Abstraction. *Acta Informatica*, 32(8):705-778 (Nov. 1995).

[Leavens-Wing97a]
Gary T. Leavens and Jeannette M. Wing. Protective interface specifications. In Michel Bidoit and Max Dauchet, editors, *TAPSOFT '97: Theory and Practice of Software Development, 7th International Joint Conference CAAP/FASE, Lille, France*, volume 1214 of *Lecture Notes in Computer Science*, pages 520–534. Springer-Verlag, New York, N.Y., 1997.

[Leavens90]
Gary T. Leavens. Modular Verification of Object-Oriented Programs with Subtypes. Department of Computer Science, Iowa State University (Ames, Iowa, 50011), TR 90-09, July 1990. Available from the URL
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR90-09/TR.ps.Z`'.

[Leavens91]
Gary T. Leavens. Modular Specification and Verification of Object-Oriented Programs. *IEEE Software*, 8(4):72-80 (July 1991).

[Leavens96b]
Gary T. Leavens. An Overview of Larch/C++: Behavioral Specifications for C++ Modules. In Haim Kilov and William Harvey (editors), *Specification of Behavioral Semantics in Object-Oriented Information Modeling* (Kluwer Academic Publishers, 1996), Chapter 8, pages 121-142. An extended version is

Department of Computer Science, Iowa State University, TR #96-01c, July 1996, which is available from the URL
'ftp://ftp.cs.iastate.edu/pub/techreports/TR96-01/TR.ps.Z'.

[Leavens97c]
Gary T. Leavens. *Larch/C++ Reference Manual*. Version 5.14. Available in 'ftp://ftp.cs.iastate.edu/pub/larchc++/lcpp.ps.gz' or on the World Wide Web at the URL
'http://www.cs.iastate.edu/~leavens/larchc++.html', October 1997.

[Ledgard80]
Henry. F. Ledgard. A Human Engineered Variant of BNF. *ACM SIGPLAN Notices*, 15(10):57-62 (October 1980).

[Leino-Nelson-Saxe00]
K. Rustan M. Leino, Greg Nelson, and James B. Saxe. ESC/Java User's Manual. SRC Technical Note 2000-02, October, 2000.

[Leino-etal00]
K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended Static Checking. Web page at
'http://research.compaq.com/SRC/esc/Esc.html'.

[Leino95] K. Rustan M. Leino. Towards Reliable Modular Programs. PhD thesis, California Institute of Technology, January 1995. Available from the URL
'ftp://ftp.cs.caltech.edu/tr/cs-tr-95-03.ps.Z'.

[Leino95b] K. Rustan M. Leino. A myth in the modular specification of programs. KRML 63, November 1995. Obtained from the author (rustan@pa.dec.com).

[Leino98] K. Rustan M. Leino. Data groups: Specifying the modification of extended state. *OOPSLA '98 Conference Proceedings*, ACM SIGPLAN Notices, Vol 33, Num 10, October 1998, pp. 144-153.

[Lerner91] Richard Allen Lerner. Specifying Objects of Concurrent Systems. School of Computer Science, Carnegie Mellon University, CMU-CS-91-131, May 1991. Available from the URL
'ftp://ftp.cs.cmu.edu/afs/cs.cmu.edu/project/larch/ftp/thesis.ps.Z'.

[Liskov-Guttag86]
Barbara Liskov and John Guttag. *Abstraction and Specification in Program Development* (MIT Press, Cambridge, Mass., 1986).

[Liskov-Wing93b]
Barbara Liskov and Jeannette M. Wing. Specifications and their use in defining subtypes. In Andreas Paepcke, editor, *OOPSLA '93 Proceedings*, volume 28, number 10 of *ACM SIGPLAN Notices*, pages 16-28. ACM Press, October 1993.

[Liskov-Wing94]
Barbara Liskov and Jeannette M. Wing. A Behavioral Notion of Subtyping. *ACM TOPLAS*, 16(6):1811-1841 (Nov. 1994).

[Meyer92a]
Bertrand Meyer. Applying "design by contract". *Computer*, 25(10):40–51, October 1992.

[Meyer92b]
Bertrand Meyer. *Eiffel: The Language*. Object-Oriented Series. Prentice Hall, New York, N.Y., 1992.

[Meyer97] Bertrand Meyer. *Object-oriented Software Construction*. Prentice Hall, New York, N.Y., second edition, 1997.

[Morgan94]
Carroll Morgan. *Programming from Specifications*, second edition (Prentice-Hall, 1994).

[Morris87] Joseph˜M. Morris. A theoretical basis for stepwise refinement and the programming calculus. *Science of Computer Programming*, 9(3):287-306, December 1987.

[Mueller-Poetzsch-Heffter00]
Peter Müller and Arnd Poetzsch-Heffter. Modular Specification and Verification Techniques for Object-Oriented Software Components. In Gary T. Leavens and Murali Sitaraman (eds.), *Foundations of Component-Based Systems*, pages 137-159. Cambridge University Press, 2000.

[Mueller-Poetzsch-Heffter00a]
Peter Müller and Arnd Poetzsch-Heffter. A Type System for Controlling Representation Exposure in Java. In S. Drossopoulou, et al. (eds.), *Formal Techniques for Java Programs*, 2000. Technical Report 269, Fernuniversität Hagen, Available from
'`http://www.informatik.fernuni-hagen.de/pi5/publications.html`'

[Mueller-Poetzsch-Heffter01a]
Peter Müller and Arnd Poetzsch-Heffter. Universes: A Type System for Alias and Dependency Control. Technical Report 279, Fernuniversität Hagen, 2001. Available from
'`http://www.informatik.fernuni-hagen.de/pi5/publications.html`'

[Mueller-Poetzsch-Heffter-Leavens02]
Peter Müller, Arnd Poetzsch-Heffter, and Gary T. Leavens. Modular Specification of Frame Properties in JML. Technical Report TR 02-02, Department of Computer Science, Iowa State University, Ames, Iowa, 50011, February 2002. Available from
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR02-02/TR.pdf`'

[Mueller02]
Peter Müller. Modular Specification and Verification of Object-Oriented Programs. Volume 2262 of *Lecture Notes in Computer Science*, Springer-Verlag, 2002.

[Nelson89] Greg Nelson. A Generalization of Dijkstra's Calculus. *ACM Transactions on Programming Languages and Systems*, 11(4):517-561 (Oct. 1989).

[Noble-Vitek-Potter98]
James Noble, Jan Vitek, and John Potter. Flexible Alias Protection. In Eric Jul (ed.), *ECOOP '98 – Object-Oriented Programming, 12th European Conference,*

Brussels, Belgium, pages volume 1445 of *Lecture Notes in Computer Science*, pages 158-185. Springer-Verlag, New York, N.Y., 1998.

[Parnas72] D. L. Parnas. On the Criteria to be Used in Decomposing Systems into Modules. *Comm. ACM*, 15(12) (Dec., 1972).

[Poetzsch-Heffter97]
Arnd Poetzsch-Heffter. Specification and Verification of Object-Oriented Programs. Habilitationsschrift, Technische Universitaet Muenchen, 1997. Available from the URL
'`http://wwweickel.informatik.tu-muenchen.de/persons/poetzsch/habil.ps.gz`'.

[Poll-Jacobs00]
E. Poll and B.P.F. Jacobs. A Logic for the Java Modeling Language JML. Computing Science Institute Nijmegen, Technical Report CSI-R0018. Catholic University of Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, November 2000.

[Raghavan-Leavens00]
Arun D. Raghavan and Gary T. Leavens. Desugaring JML Method Specifications. Technical Report 00-03a, Department of Computer Science, Iowa State University, Ames, Iowa, 50011, April, 2000, revised July 2000. Available in
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR00-03/TR.ps.gz`'.

[Rosenblum95]
David S. Rosenblum. A practical approach to programming with assertions. *IEEE Transactions on Software Engineering*, 21(1):19–31, January 1995.

[Ruby-Leavens00]
Clyde Ruby and Gary T. Leavens. Safely Creating Correct Subclasses without Seeing Superclass Code. In *OOPSLA 2000 Conference on Object-Oriented Programming, Systems, Languages, and Applications, Minneapolis, Minnesota*, pp. 208-228. Volume 35, number 10 of *ACM SIGPLAN Notices*, October, 2000. Also technical report 00-05d, Department of Computer Science, Iowa State University, Ames, Iowa, 50011. April 2000, revised April, June, July 2000. Available in
'`ftp://ftp.cs.iastate.edu/pub/techreports/TR00-05/TR.ps.gz`'.

[Spivey92] J. Michael Spivey. *The Z Notation: A Reference Manual*, second edition, (Prentice-Hall, Englewood Cliffs, N.J., 1992).

[Steyaert-etal96]
Patrick Steyaert, Carine Lucas, Kim Mens, and Theo D'Hondt. Issues in the Design and Documentation of Class Libraries. In *OOPSLA '96 Proceedings* (ACM SIGPLAN Notices, 31(10):268-285, Oct., 1996).

[Tan94] Yang Meng Tan. Formal Specification Techniques for Promoting Software Modularity, Enhancing Documentation, and Testing Specifications. MIT Lab. for Comp. Sci., TR 619, June 1994. Also published as *Formal Specification Techniques for Engineering Modular C Programs*. International Series in Software Engineering (Kluwer Academic Publishers, Boston, 1995).

[Watt91] David A. Watt. *Programming Language Syntax and Semantics*. Prentice Hall, International Series in Computer Science, New York, 1991.

[Wills92b] Alan Wills. Specification in Fresco. In Susan Stepney and Rosalind Barden and David Cooper (eds.), *Object Orientation in Z*, chapter 11, pages 127-135. Springer-Verlag, Workshops in Computing Series, Cambridge CB2 1LQ, UK, 1992.

[Wing83] Jeannette Marie Wing. *A Two-Tiered Approach to Specifying Programs* Technical Report TR-299, Mass. Institute of Technology, Laboratory for Computer Science, 1983.

[Wing87] Jeannette M. Wing. Writing Larch Interface Language Specifications. *ACM Transactions on Programming Languages and Systems*, 9(1):1-24 (Jan. 1987).

[Wing90a] Jeannette M. Wing. A Specifier's Introduction to Formal Methods. *Computer*, 23(9):8-24 (Sept. 1990).

Example Index

Concept Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

Z

Table of Contents