

# How the Design of JML Accommodates Both Runtime Assertion Checking and Formal Verification

Gary T. Leavens<sup>a</sup>, Yoonsik Cheon<sup>b</sup>, Curtis Clifton<sup>a</sup>,  
Clyde Ruby<sup>a</sup>, David R. Cok<sup>c</sup>

<sup>a</sup>*Department of Computer Science, Iowa State University, Ames, Iowa, USA*

<sup>b</sup>*Department of Computer Science, University of Texas at El Paso,  
El Paso, Texas, USA*

<sup>c</sup>*Eastman Kodak Company, Research & Development Laboratories,  
Rochester, New York, USA*

---

## Abstract

Specifications that are used in detailed design and in the documentation of existing code are primarily written and read by programmers. However, most formal specification languages either make heavy use of symbolic mathematical operators, which discourages use by programmers, or limit assertions to expressions of the underlying programming language, which makes it difficult to write exact specifications. Moreover, using assertions that are expressions in the underlying programming language can cause problems both in runtime assertion checking and in formal verification, because such expressions can potentially contain side effects. The Java Modeling Language, JML, avoids these problems. It uses a side-effect free subset of Java's expressions to which are added a few mathematical operators (such as the quantifiers `\forall` and `\exists`). JML also hides mathematical abstractions, such as sets and sequences, within a library of Java classes. The goal is to allow JML to serve as a common notation for both formal verification and runtime assertion checking; this gives users the benefit of several tools without the cost of changing notations.

*Key words:* specification languages, runtime assertion checking, formal methods, program verification, programming by contract, Java language, JML language

---

---

*Email address:* leavens@cs.iastate.edu (Gary T. Leavens).

## 1 Introduction

The Java Modeling Language, JML [1,2], is the result of a cooperative, international effort aimed at providing a common notation and semantics for the specification of Java code at the detailed-design level [3]. JML is being designed cooperatively so that many different tools can use a common notation for Hoare-style behavioral interface specifications. In this paper we explain the features of JML's design that make its assertions easily understandable by programmers and suitable for both runtime assertion checking and formal verification.

### 1.1 Background

By a *Hoare-style* specification we mean one that uses pre- and postconditions to specify the behavior of methods [4–7]. A *behavioral interface specification language* (BISL) is a specification language that specifies both the syntactic interface of a module and its behavior [8–11]. JML, the interface specification languages in the Larch family [8–11] and RESOLVE/C++ [12,13] are BISLs. Most design by contract languages and tools, such as Eiffel [14,15] and APP [16], are also BISLs, because they place specifications inside programming language code. By contrast, neither Z [17–19] nor VDM [20,21,5,22] is a BISL; they have no way to specify interface details for a particular programming language. OCL [23,24] is a BISL for the UML, but the UML itself is language-independent; this poses problems for a Java programmer, because the UML does not have standard notations for all details of Java method signatures. For example, the UML's syntax for specifying the signatures of operations has no standard notation for declaring that a Java method is `strictfp` or for declaring the exceptions that a method may throw [25, pp. 128-129] [26, p. 516].<sup>1</sup> Also the OCL has no standard constraints that correspond to JML's exceptional postconditions. Because BISLs like JML specify both interface and behavior, they are good at specifying detailed designs that include such Java details. This makes JML well suited to the task of documenting reusable components, libraries, and frameworks written in Java.

---

<sup>1</sup> Larman notes that the UML has some nonstandard ways to specify the exceptions that a method may throw, by either using Java's own syntax directly or by using a "property string".

## 1.2 Tool Support

Because BISLs are easily integrated with code, they lend themselves to tool support for activities related to detailed design, coding, testing, and maintenance. An important goal of JML is to enable a wide spectrum of such tools. Besides tools that enforce JML's semantics (e.g., type checking), the most important JML tools help with the following tasks.

**Runtime checking and testing.** The Iowa State group provides (via [www.jmlspecs.org](http://www.jmlspecs.org)):

- the `jmlc` runtime assertion checking compiler [27], which generates class files from JML-annotated Java sources,<sup>2</sup> and
- the `jmlunit` tool [28], which uses the runtime assertion checker to generate test oracle code for JUnit tests.

**Documentation.** David Cok provides the `jmldoc` tool, also available through [www.jmlspecs.org](http://www.jmlspecs.org), which generates HTML documentation similar to that produced by Javadoc [29], but containing specifications as well. The generated documentation is useful for browsing specifications or publishing on the web.

**Static analysis and verification.** The following tools are prepared by our partners at Compaq and the University of Nijmegen:

- The ESC/Java tool [30–32] statically checks Java code for likely errors. ESC/Java understands a subset of JML annotations.
- The ESC/Java2 tool [?] extends ESC/Java to understand all of JML and to check most of it.
- The LOOP tool [33–36] assists in the formal verification of the correctness of implementations from JML specifications, using the theorem prover PVS.

In addition, the Daikon dynamic invariant detector [37,38] outputs invariants for Java programs in a subset of JML, and the Korat automated testing tool [39] uses the `jmlunit` tool to exercise the test data it derives.

In this paper, we discuss how JML meets the needs of tools for runtime assertion checking, documentation, static analysis, and verification. We focus on runtime assertion checking and formal verification, which we consider to be the extremes of the spectrum of tools that a BISL might support. The tasks of runtime assertion checking and formal verification have widely differing needs:

- Runtime assertion checking places a high premium on executability. Many

---

<sup>2</sup> Besides this runtime assertion checking work at Iowa State, which relies on adding instrumentation to compiled code, Steven Edwards's group at Virginia Tech is working on a wrapper-class based approach to runtime assertion checking that will allow instrumentation of programs for which source code is not available.

specification languages intended for runtime assertion checking, such as Eiffel [14,15] and APP [16], only allow assertions that are completely executable. This is sensible for a language that is intended only to support runtime assertion checking and not formal verification.

- On the other hand, formal theorem proving and reasoning place a high premium on the use of standard mathematical notations. Thus, most specification languages intended for formal reasoning or verification, such as VDM, the members of the Larch family, and especially Z, feature a variety of symbolic mathematical notations. Many expressive mathematical notations, such as quantifiers, are impossible, in general, to execute at runtime. Again, including such notations is sensible for a language intended only to support formal theorem proving and reasoning and not runtime assertion checking.

### 1.3 Problems

We begin by describing some problems that arise when addressing the needs of the range of tools exemplified by runtime assertion checking and formal verification. Like the tools, the problems encompass a broad range, including issues of notation, logic, and expressiveness.

#### 1.3.1 Notational Problem

It is often said that syntax does not matter; however, our experience with Larch/Smalltalk [40] and Larch/C++ [41–45] showed that programmers object to learning a specialized mathematical notation (the Larch Shared Language). This is similar to the problems found by Finney [46], who did a preliminary experiment demonstrating that the symbolic notation in Z specifications may make them hard to read. Conversely, in executable languages like Eiffel and APP, programmers feel comfortable with the use of the programming language’s expressions in assertions. Such an assertion language is therefore more appealing for purposes of documentation than highly symbolic mathematical notations.

To summarize, the first problem that we address in this paper is how to provide a good syntax for specification expressions. *Specification expressions* are the syntactic forms that are used to denote values in assertions. By a *good* syntax we mean one that is close enough to programming language expressions that programmers feel comfortable with it and yet has all of the features necessary to support both runtime assertion checking and formal verification.

### 1.3.2 *Undefinedness Problem*

Expressions in a programming language may abruptly terminate (e.g., throw exceptions) and may go into infinite loops; consequently, they may have undefined values from a strictly mathematical point of view. Programming languages typically provide features to control what subexpressions must be evaluated, which can be used to avoid such undefinedness. For example, Java provides short-circuit versions of boolean operators (such as `&&` and `||`) that allow programmers to suppress evaluation of some subexpressions.

We want both programmers and mathematicians to use JML's notations; hence, JML's specification expressions should not only look like Java's expressions and use Java's semantics, but should also validate the standard laws of logic. However, because of a potential for undefinedness, Java expressions do not satisfy all the standard rules of logic; for example, in Java the conjunction  $E_1 \ \&\& \ E_2$  is not equal to  $E_2 \ \&\& \ E_1$ , although in logic they would be equal. To resolve this conflict, we are willing to accept a slightly different semantics for assertion evaluation as long as programmers are not too surprised by it.

Thus, the second problem we address in this paper is how to find a semantics for expressions used in assertions that validates standard laws of logic and yet does not surprise programmers and is still useful for runtime assertion checking.

### 1.3.3 *Side Effects Problem*

Another important semantic issue is that expressions in a programming language like Java (and most others, including Eiffel) can contain side effects. Side effects have a very practical problem related to runtime assertion checking. It is generally assumed that assertions may be evaluated or skipped with no change in the outcome of a computation, but an assertion with side effects has the potential to alter the computation's outcome. For example, an assertion with side effects might mask the presence of a bug that would otherwise be revealed or cause bugs that are not otherwise present. Because one of the principal uses of runtime assertion checking is debugging and isolating bugs, it is unacceptable for side effects from assertion checking to alter the outcome of a computation.

Thus, the third problem that we address in this paper is how to prevent side effects in assertions while still retaining as much of the syntax of normal programming language expressions as possible.

### 1.3.4 *Mathematical Library Problem*

Most specification languages come with a library of mathematical concepts such as sets and sequences. Such concepts are especially helpful in specifying collection types. For example, to specify a `Stack` type, one would use a mathematical sequence to describe, abstractly, the states that a stack object may take [47]. VDM, OCL, Z, and the interface specification languages of the Larch family all have libraries of such mathematical concepts. They also are standard in theorem provers such as PVS.

However, as discussed in Section 1.3.1, we want to limit the barriers that Java programmers must overcome to use JML. Thus, the fourth problem that we address in this paper is how to provide a library of mathematical concepts in a way that does not overwhelm programmers, and yet is useful for formal verification.

### 1.4 *Other Goals of JML*

In addition to providing solutions to the preceding four problems, the design of JML is guided and constrained by several other goals. One of the most important of these goals is to allow users to write specifications that document detailed designs of existing code. This motivates the choice of making JML a BISL, as described above. Moreover, we would like JML to be useful for documenting code regardless of whether it was designed according to any particular design method or discipline. This is important because the cost of specification is high enough that it is not always justified until one knows that the design and the code have stabilized enough to make the documentation potentially useful to other people.

In general, JML's design adheres to the goal of being able to document existing designs; however, there is one significant aspect of JML's design that departs from this goal—JML imposes the specifications of supertypes on subtypes, a property termed *specification inheritance*, in order to achieve behavioral subtyping [48].

JML's use of specification inheritance is justified by another of our goals: we want JML to support *modular reasoning*, that is, reasoning about the behavior of a compilation unit using just the specifications of the compilation units that it references (as opposed to the details of their implementations). Modular reasoning is important because without it, the difficulty of understanding an object-oriented program increases much more rapidly than the size of the program, and thus the benefits of the abstraction mechanisms in object-oriented languages are lost. Consequently, modular reasoning is also important for formal verification, because then the scope of the verification

problem is limited.

Specification inheritance, and the resulting behavioral subtyping, allows modular reasoning to be sound, by allowing one to reason based on the static types of references. Subsumption in Java allows a reference to a subtype object to be substituted for a supertype reference. The requirements of behavioral subtyping [48–54] guarantee that all such substituted objects will obey the specifications inherited from the static type of the reference [48,55,56].

Because modular reasoning provides benefits to programmers and verifiers, we favor specification inheritance over the conflicting goal of being able to document existing designs that do not follow behavioral subtyping. In any case, it is possible to work around the requirements of behavioral subtyping for cases in which a subtype does not obey the inherited specifications of its supertype(s). One simply underspecifies each supertype enough to allow all of the subtypes that are desired [53,54]. Note that this work-around does not involve changing the code or the design, but only the specification, so it does not interfere with the goal of documenting existing code.

### *1.5 Outline*

The remainder of this paper is organized as follows. The next section discusses our solution to the notational problem described above. Having described the notation in general terms, Section 3 provides more background on JML. The subsequent three sections treat the remaining problems discussed above. The paper ends with a discussion of related work and some conclusions.

## **2 Solving the Notational Problem**

To solve the notational problem described in Section 1.3.1, JML generally follows Eiffel, basing the syntax of specification expressions on Java’s expression syntax. However, because side effects are not desired in specification expressions, JML’s specification expressions do not include Java expressions that can cause obvious side effects, i.e., assignment expressions and Java’s increment and decrement operators (`++` and `--`).

Furthermore, to make JML suitable for formal verification efforts, JML includes a number of operators that are not present in Java [2, Section 3]. The syntax of these operators comes in two flavors: those that are symbolic and those that are textual.

We did not want to introduce excess notation that would cause difficulties

for programmers when reading specifications, so JML adds just five symbolic operators. Four of these are logical operators: forward and reverse implication, written `==>` and `<==`, respectively, and logical equivalence and inequivalence, written `<==>` and `<!=>`, respectively. The inclusion of symbols for logical operators is inspired by the calculational approach to formal methods [57–59]. The other symbolic operator is `<:`, which is used to compare types to see if they are in a subtype relationship [31].

All the other operators added to Java and available in JML’s specification expressions use a textual notation consisting of a backslash (`\`) followed by an English word or phrase. For example, the logical quantifiers in JML are written as `\forall` and `\exists` [2].

Besides these quantifiers, JML also has several other operators using this backslash syntax. One of the most important is `\old()`, which is used in method postconditions to indicate an expression whose value is computed in the pre-state of a method call. For example, `\old(i-1)` denotes the value of `i-1` evaluated in the pre-state of a method call. This notation is borrowed from the `old` operator in Eiffel. Other JML expressions using the backslash syntax include `\fresh(o)`, which says that `o` was not allocated in the pre-state of a method call, but is allocated (and not null) in the post-state, and `\result`, which denotes the normal result returned by a method.

The backslashes in the syntax of these operators serve a very important purpose—they prevent the rest of the operator’s name from being interpreted as a Java identifier. This allows JML to avoid reserving Java identifiers in specification expressions. For example, `result` can be used as a program variable and is distinguished from `\result`. This trick is useful in allowing JML to specify arbitrary Java programs. Indeed, because a goal of JML is to document existing code, it cannot add new reserved words to Java.

### 3 Background on JML

In this section we provide additional background on JML that will be useful in understanding our solutions to the remaining problems.

#### 3.1 Semantics of Specification Expressions

Just as JML adopts much of Java’s expression syntax, it attempts to keep JML’s semantics similar to Java’s. In particular, the semantics of specification expressions is a reference semantics. That is, when the name of a variable or



field is used in an expression, it denotes either a primitive value (such as an integer) or a reference to an object. References themselves are values in the semantics, which allows one to directly express aliasing or the lack of it. For example, the expression `arg != fieldVal` says that `arg` and `fieldVal` are not aliased. Java also allows one to compare the states of objects using the `equals` method. For example, in the postcondition of a `clone` method, one might write the following to say that the result returned by `clone` is a newly allocated object that has the same state as the receiver (`this`):

```
\fresh(\result) && this.equals(\result);
```

Note that the exact meaning of the `equals` method for a given type is left to the designer of that type, as in Java. Thus, if one only knows that `o` is an `Object`, it is hard to conclude much about `x` from `o.equals(x)`.

Because JML uses this reference semantics, specifiers must show the same care as Java programmers when choosing between the `==` and `equals` equality tests. And like Eiffel, but unlike Larch-style interface specification languages, JML does not need “state functions” to be applied to extract the value of an expression from a reference. Values are implicitly extracted as needed by methods and operators. Besides being easier for programmers, this lends some succinctness to the notation.

Currently, JML adopts all of the Java semantics for integer arithmetic. Thus types such as `int` use two’s complement arithmetic and are finite. Although Java programmers are, in theory, aware of the nature of integer arithmetic, JML’s adoption of Java’s semantics causes some misunderstandings; for example, some published JML specifications are inconsistent because of this semantics [60]. Chalin has suggested adding new primitive value types for infinite precision arithmetic to JML; in particular, he suggests a type `\bigint` for infinite precision integers [61,60]. He is currently implementing and experimenting with this idea.

### 3.2 Method and Type Specifications

To explain JML’s semantics for method specifications, we use the example in Fig. 1. JML uses special comments, called *annotations*, to hold the specification of behavior; these are added to the interface information contained in the Java code. A specifier writes these annotation comments by inserting an at-sign (@) following the usual characters that signify the start of a comment. In multi-line annotation comments, at-signs at the beginnings of lines are ignored.

Fig. 1 starts with a “model import” directive, which says that JML will con-

---

```

/*@ model import org.jmlspecs.models.*;
public interface Stack {
    /*@ public model instance JMLObjectSequence absVal;

    /*@ public instance invariant absVal != null;

    /*@ public normal_behavior
        @   requires true;
        @   assignable absVal;
        @   ensures absVal.equals(\old(absVal.insertFront(x))); @*/
    void push(Object x);

    /*@   public normal_behavior
        @   requires !absVal.isEmpty();
        @   assignable absVal;
        @   ensures absVal.equals(\old(absVal.trailer()))
        @       && \result == \old(absVal.first());
        @ also
        @   public exceptional_behavior
        @   requires absVal.isEmpty();
        @   assignable \nothing;
        @   signals (Exception e)
        @       e instanceof IllegalStateException; @*/
    Object pop();

    /*@ ensures \result <==> absVal.isEmpty();
    /*@ pure @*/ boolean isEmpty();
}

```

Fig. 1. The specification and code for the interface `Stack`.

---

sider all types in the named package, `org.jmlspecs.models`, to be imported for purposes of the specification. This allows the JML tools to find the type `JMLObjectSequence` (see the third line) in that package.

The type `JMLObjectSequence` is used as the type of the model instance field, named `absVal`. In this declaration, the `model` keyword says that the field is not part of the Java code, but is used solely for purposes of specification. The `instance` keyword says that the field is imagined, for purposes of specification, to be a non-static field in every class that implements this interface.<sup>3</sup>

Following the declaration of the two model instance fields is an invariant. It

---

<sup>3</sup> Omitting `instance` makes fields static and final, which is Java's default for fields declared in interfaces.

says that the field `absVal` is never null.

Following the invariant are the declarations and specifications of three methods. In JML, a method's specifications are typically written, as they are in Fig. 1, before the header of the method they specify. This makes the scope of the formal parameters of a method a bit strange, because it extends backward into the method's specification. However, it works best with Java tools, which expect comments related to a method, such as Javadoc comments, to precede the method's header.

Consider the specification of the first method, `push`. This shows the general form of a “normal behavior” specification case. A *specification case* includes a precondition, indicated by the keyword `requires`, and some other specification clauses. A specification case is satisfied if, whenever the precondition is satisfied, the other clauses are also satisfied. Additionally, in a normal behavior specification case, the method must not throw an exception when the precondition is satisfied. The specification case given for `push` includes, besides the `requires` clause, a frame axiom, introduced by the keyword `assignable`, and a normal postcondition, following the keyword `ensures`.

As with specification languages in the Larch family, a precondition that is just true can be omitted. In the Larch family, an omitted frame axiom means “`assignable \nothing;`”, which is a very strong specification that says that the method has no side effects. Following a suggestion of Erik Poll, we decided that such a specification was too strong for a default. So in JML, an omitted frame axiom allows assignment to all locations. This agrees with most of the defaults for omitted clauses in JML, which impose no restrictions.

JML also allows specifiers to write “exceptional behavior” specification cases, which say that, when the precondition is satisfied, the method must not return normally but must instead throw an exception. An example appears in the specification of the `pop` method. This specification has two specification cases connected with `also`. The meaning of the `also` is that the method must satisfy both of these specification cases [62,63]. Thus, when the value of the model instance field `absVal` is not empty, a call to `pop` must return normally and must satisfy the given `ensures` clause. But when the value of the model instance field `absVal` is empty, a call to `pop` must throw an `IllegalStateException`. This kind of case analysis can be desugared into a single specification case, which can be given a semantics in the usual way [34,44,64,65].

The specification cases given for `push` and `pop` are heavyweight specification cases [2, Section 1]. Such specification cases are useful when one wants to give a relatively exact specification, especially for purposes of formal verification. For runtime assertion checking or documentation, one may want to specify only part of the behavior of a method. This can be done using JML's lightweight

specification cases, which are indicated by the absence of a behavior keyword (like `normal_behavior`). Fig. 1 gives an example of a lightweight specification case in the specification of the method `isEmpty`.

## 4 Dealing with Undefinedness

As discussed in Section 1.3.2, a fundamental problem in using the underlying language for specification expressions is dealing with expressions that have undefined values. In Java, undefinedness in expressions is typically signaled by the expression throwing an exception. For example, when one divides an integer by 0, the expression throws an `ArithmeticException`. Exceptions may also be thrown by methods called from within specification expressions.

Specification languages have adopted several different approaches to dealing with undefinedness in expressions [66,67]. We wanted a semantics that would not be surprising to either Java programmers or to those doing formal verification. Typically, a Java programmer would try to write the specification in a way that “protects” the meaning of the expression against any source of undefinedness [68]. This can be accomplished by using the short-circuit boolean operators; for example, a specifier might write `denom > 0 && num/denom > 1` to be sure that the division would be defined whenever it was carried out.

However, we would like specifications to be meaningful even if they are not protective. Hence, the semantics of JML does not rely on the programmer writing protective specifications but, instead, ensures that every expression has some value. To do this, we adopted the “underspecified total functions” approach favored in the calculational style of formal methods [59,67]. That is, an expression that would not have a value in Java is given an arbitrary, but unspecified, value. For example, `num/0` has some integer value, although this approach does not say what the value is, and says only that it must be uniformly substituted in any surrounding expression.

An advantage of this substitution approach is that it validates the rules for standard logic. For example, in JML,  $E_1 \ \&\& \ E_2$  is equivalent to  $E_2 \ \&\& \ E_1$ . Consider what happens if  $E_1$  throws an exception; in that case, one may choose some unspecified boolean value for  $E_1$ , say  $b$ . This means that  $E_1 \ \&\& \ E_2$  equals  $b \ \&\& \ E_2$ , which is equal to  $E_2 \ \&\& \ b$ , as can be seen by a simple case analysis on  $E_2$ 's value. The case where  $E_2$  throws an exception is similar. Furthermore, if programmers write protective specifications, they will never be surprised by the details of this semantics.

The JML assertion checking compiler takes advantage of the semantics of undefinedness to attempt, as much as possible, to detect possible assertion

violations [27]. That is, assertion checking attempts to use a value that will make the overall assertion false, whenever the undefinedness of some sub-expression allows it to do so. In this way, the assertion checker can both follow the rules of standard logic and detect places where specifications are not sufficiently protective. This is a good example of how JML caters to the needs of both runtime assertion checking and formal verification.

## 5 Preventing Side Effects in Assertions

As discussed in Section 1.3.3, it is important to prevent side effects in assertions, for both practical and theoretical reasons.

JML is designed to prevent side effects in assertions statically. It does this using an effect-checking type system [69,70]. At the heart of the system is the `pure` modifier. Only methods and constructors that are declared to be `pure` can be used in assertions, and methods and constructors declared `pure` must be side-effect free. In this section we first explain the details of this semantics, and then discuss its ramifications.

### 5.1 JML's Purity Restrictions

JML's semantic restrictions on pure methods and constructors are as follows:

- A *pure method* implicitly has a specification that includes the following specification case [2, Section 2.3.1]:  
`assignable \nothing;`  
This ensures that a correct implementation of the method has no side effects.
- “A *pure constructor* implicitly has a specification that only allows it to assign to the instance fields of the class in which it appears” (including inherited instance fields) [2, Section 2.3.1]. This ensures that, if the constructor is correctly implemented, then a `new` expression that calls it has no side effects.

To explain the first restriction, it helps to first explain the semantics of JML's assignable clause [2, Section 2.1.3.1]. The assignable clause of a method  $m$  describes the set of existing, non-local storage locations that may be assigned by  $m$ 's execution. Local variables in a method, such as  $m$ 's formal parameters and variables declared in  $m$ 's body, can be assigned regardless of  $m$ 's assignable clause. Similarly, fields of objects allocated by  $m$  itself, and thus not existing in  $m$ 's pre-state, can be freely assigned during the  $m$ 's execution regardless of its assignable clause. Other locations, which exist in the pre-state, and which are

not local to  $m$ , can only be assigned if they are mentioned in  $m$ 's assignable clause (perhaps implicitly via a data group).

Therefore, the first restriction implies that a pure method may not perform any input or output, nor may it assign to existing, non-local storage. Similarly, by the second restriction, a pure constructor may not do any I/O and may not assign to non-local storage other than the instance fields of the object the constructor is initializing. A pure constructor is allowed to assign to the instance fields of the object being constructed, because in an expression such as `new T()`, the newly-created object does not exist in that expression's pre-state.

Note that, in JML, saying that a method may not assign to existing, non-local storage means precisely that—even benevolent side effects are prohibited [2, Section 2.1.3.1]. A *benevolent side effect* is a change in the internal state of an existing object in a way that is not externally visible [4]. Prohibiting even benevolent side effects is necessary for sound modular reasoning about method implementations [71]. It is also a useful restriction for reasoning about supertypes from their specifications [72] and for reasoning about concurrent programs.

In the current version of JML, the purity restrictions described above are enforced conservatively. The most conservative aspect of purity checking is that pure methods and pure constructors may only invoke other methods and constructors that are pure. This is somewhat overly conservative, but is simple to implement. A less conservative rule would allow assignments to fields in objects that are created after the start of a pure method's execution, as such assignments are not covered by the assignable clause. In any case, the purity of a method  $m$  can be checked modularly by using the assignable clauses of the methods that  $m$  calls.

The type system of JML is an important advance over languages like Eiffel, which trust programmers to avoid side effects in assertions rather than statically checking this property. However, as we will see in the following subsection, JML's purity restrictions give rise to some practical problems.

Many of these practical problems arise from the interaction between purity checking and specification inheritance. Because a pure method has an implicit specification that prohibits side effects during its execution, all methods and constructors that override a pure method or constructor must also be pure. That is, in JML, purity is inherited. This inheritance of purity is necessary to make purity checking (and reasoning) modular in the presence of subtyping and dynamic dispatch.

An important consequence of inheritance of purity is that a method cannot be correctly specified as pure if any overriding method has side effects. In

particular, a method in `Object` can be specified as pure only if every override of that method, in every Java class, obeys JML's purity restrictions.

## 5.2 Practical Problems with JML's Purity Restrictions

An initial practical problem is how to decide which methods in Java's libraries should be specified as `pure`. One way to start to answer this question is to use a static analysis to conservatively estimate which methods in Java's libraries have side effects. A conservative analysis could count a method as having side effects if it assigns to non-local storage or calls native methods (which may do I/O), either directly or indirectly. All other methods can safely be specified as `pure`, provided they are not overridden by methods that the analysis says have side effects. Researchers from Purdue have provided a list of such methods to us, using their tools from the Open Virtual Machine project.<sup>4</sup> We plan to integrate this technology into the JML tools eventually.

Declaring a method to be `pure` entails a very strong specification, namely that the method and all possible overriding methods have no side effects. Thus, finding that a method, and all known methods that override it, obey JML's purity restrictions is not the same as deciding that the method *should* be specified as pure. Such a decision affects not just all existing overrides of the method, but all future implementations and overrides. How is one to make such a decision?

This problem is particularly vexing because there are many methods that are intuitively side-effect free, but that do not obey JML's purity restrictions. Methods with benevolent side effects are common examples. Two examples from the protocol of `Object` will illustrate the importance of this problem.

First, consider computing a hash code for an instance of a class. Because this may be computationally costly, an implementation may desire to compute the hash code the first time it is asked for and then cache the result in a private field of the object. When the hash code is requested on subsequent occasions, the cached result is returned without further computation. For example, this is done in the `hashCode` method of Java's `String` class. However, in JML, storing the computed hash code into the cache is considered to be a side effect. So `String`'s `hashCode` method cannot be specified as pure.

Second, consider computing object equality. In some implementations, an object's fields might be lazily initialized or computed only on first access. If the `equals` method happens to be the first such method to be called on such an object, it will trigger the delayed computation. We found such an example in our

---

<sup>4</sup> See <http://www.ovmj.org/>.

work on the MultiJava compiler [73,74]; in this compiler, the class `CClassType` has such delayed computations, and its override of `Object`'s `equals` method can trigger a previously delayed computation with side effects. It seems very difficult to rewrite this method to be side-effect free, because to do so one would probably need to change the compiler's architecture. (Similar kinds of lazy initialization of fields occur in implementations of the Singleton pattern, although these usually do not affect the `equals` method.)

We have shown two cases where methods in the protocol of `Object` are overridden by methods that cannot be pure. By purity and specification inheritance, these examples imply that neither `hashCode` nor `equals` can be specified as pure in `Object`. `Object` is typically used in Java as the type of the elements in a collection. Hence, in the specification of a collection type, such as a hash table, one cannot use the `hashCode` or `equals` methods on elements. Without changes, this would make JML unsuitable for specifying collection types.

(This problem is mostly a problem for collection types, because one can specify many subclasses of `Object` with pure `hashCode` and `equals` methods. Specifications operating on instances of such subclasses can use these methods without violating JML's type system.)

### 5.3 Solving the Problems

The desire to use intuitively side-effect free methods in specifications, even if they are not pure according to JML's semantics, is strong enough that we considered changing the semantics of the `assignable` clause in order to allow benevolent side effects. However, we do not know how to do that and still retain sound modular reasoning [71]. In any case, the use of such methods in runtime assertion checking would still be problematic because of the side effects they might cause. In addition, we would like to prevent problems when a programmer wrongly believes that side effects are benevolent; it is not clear whether an automatic static analysis could prevent such problems, and even if so, whether such a tool could be modular.

Thus far, the only viable solution we have identified is to refactor specifications by adding pure *model* (i.e., specification-only) methods that are to be used in specifications in place of program methods that cannot be pure. That is, whenever one has an intuitively side-effect free program method,  $m$ , that is not pure according to JML's semantics, one creates a pure model method  $m'$ , which returns the same result as  $m$  but without its side effects. Then one replaces calls to  $m$  by calls to  $m'$  in assertions.

We are currently experimenting with this solution. The most important part of this experiment is to replace uses of `Object`'s `equals` method, which cannot be



---

```

/*@ public normal_behavior
  @ assignable objectState;
  @ ensures \result <==> this.isEqualTo(obj);
  @*/
public boolean equals(Object obj);

/*@ public normal_behavior
  @ requires obj != null;
  @ assignable \nothing;
  @ ensures (* \result is true iff obj is equal to this *);
  @ also
  @ public normal_behavior
  @ requires obj != null && \typeof(this) == \type(Object);
  @ assignable \nothing;
  @ ensures \result <==> this == obj;
  @ also
  @ public normal_behavior
  @ requires obj == null;
  @ assignable \nothing;
  @ ensures \result <==> false;
public pure model boolean isEqualTo(Object obj) {
  return this == obj;
}
@*/

```

Fig. 2. The refactored specification for `Object`'s `equals` method and the pure model method `isEqualTo`. The text between `(*` and `*)` in the first specification case of `isEqualTo`'s specification is an “informal description”, which formally is equivalent to writing `true` [44].

---

pure, with calls to a new pure model method in `Object`, called `isEqualTo`. The specifications of these methods are shown in Fig. 2. The `assignable` clause in the specification of the `equals` method permits benevolent side effects; it is also specified to return the same result as would a call to `isEqualTo`. Thus, whenever someone overrides `equals`, they should also override the `isEqualTo` method. When an override of `equals` is specified as pure, then an override of `isEqualTo` in the same class can be specified in terms of this pure `equals` method, and the implementation of the model `isEqualTo` method can simply call `equals` as well. However, an implementation of `equals` can never call `isEqualTo`, because program code cannot call model methods (since model methods can only be used in specifications). Therefore, to avoid code duplication when `equals` is not declared to be pure but the two methods share some common implementation code, one can introduce a (non-model) pure, private method that both `equals` and `isEqualTo` can call.

We have also applied this refactoring to all the collection classes in `java.util` (and in other packages) that we had previously specified, in order to check that the solution is viable. So far the results seem satisfactory. However, as of March 2004, this restructuring is not part of the JML release, because the JML tools are not yet able to handle some of the details of this approach. In particular, the runtime assertion checker is not yet able to compile the model methods added to `Object` without having all of `Object`'s source code available. (And we cannot legally ship Sun's source code for `Object` in the JML release.) However, we are working on solutions to this problem that will allow us to obtain more experience with this approach and to do more case studies.

#### 5.4 *Future Work on Synchronized Methods and Purity*

JML currently permits synchronized methods to be declared `pure` if they meet all the criteria described in Section 5.1. Given that obtaining a lock is a side effect that can affect control flow in a program, does allowing synchronized methods to be pure violate the intent of JML's purity restrictions? That is the question we investigate in this section.

##### 5.4.1 *Background*

Java has language-level support for mutual exclusion [75, §10.3]. A method may be declared `synchronized`, which means that the thread making a call to that method must first obtain a lock on the method's receiver object. The receiver object for a method call  $o.m(e)$  is  $o$ , and for a static method call of the form  $C.g(e)$  is the class object for the class in which the method is located, namely  $C.class$ . A thread that is attempting to obtain a lock will wait until no other thread holds it; however, if the thread already holds the lock, it will proceed without interruption and without changing any storage. That is, if the thread holds the lock already, it can enter a synchronized method without any side effects.

Java has various ways to test whether a thread holds a lock. The most explicit of these is the side-effect free method `Thread.holdsLock`. Thus even a sequential program can observe side effects from locking.

##### 5.4.2 *The Problem*

A synchronized method is not, in general, side-effect free. The locking used in synchronization is a modification of the state of a program execution, and can alter control flow in concurrently executing threads. Thus it would seem that

synchronized methods violate the intent of JML's purity restrictions, because calling them can, in general, cause side effects.

On the other hand, if we followed this observation to its logical conclusion and prohibited synchronized methods from being declared to be pure, JML would have several problems. The first problem is that prohibiting pure synchronized methods would be inconvenient, violating the ease-of-use requirement. For example, the class `java.util.Vector` is commonly used and has many synchronized methods that could otherwise be pure, such as `firstElement` and `elementAt`; such methods, or similar model methods, are necessary to access the state of a vector in assertions. A more important problem is that during runtime assertion checking, assertions need to be evaluated in a thread-consistent state. In a multithreaded program, an object that is shared by several threads can only be guaranteed to be in a consistent state when it is locked. If assertions, such as pre- and postconditions and the methods called within them, are evaluated without locking the shared objects involved, then other threads may modify the internal state of the object during assertion evaluation, leading to nonsensical or inconsistent results.

Hence we have a dilemma: obtaining a lock is a side effect, but methods called during assertion checking must, in general, be guarded by a lock if they are to return meaningful and consistent results.

#### 5.4.3 Possible Approaches

The only way out of the dilemma appears to be to consider special cases in which either obtaining a lock does not cause side effects or in which the side effects due to locking cannot be observed. The key observation is that a thread that calls a synchronized method does not obtain locks it already holds. That is, a synchronized method will act in a pure manner if it is invoked by a thread that already owns a lock on the method's receiver. In particular, no side effects occur during a call,  $o.m(e)$  to a pure synchronized method that originates from within another synchronized method whose receiver is  $o$ , because the other method already holds  $o$ 's lock. It follows that synchronized methods will not have side effects if they are called during assertion checking on behalf of another synchronized method on the same object. This condition could be enforced statically, and might be useful for model methods, which cannot be called directly by Java program methods.

However, checking whether a thread holds a lock is not, in general, statically decidable. So one possible semantics for JML is to require that all pure synchronized methods (implicitly) satisfy the following specification:

```
requires Thread.holdsLock(this);
```

The runtime assertion checker could check this precondition and raise an assertion violation error if the calling thread does not hold the receiver’s lock. This check could be done before the calling thread attempt to obtain the receiver’s lock, for example by calling the synchronized method from a non-synchronized method that performs this check first. This approach would guarantee that the synchronized method would not have the side effect of obtaining the receiver’s lock.

Unfortunately, the above precondition is still too strong for concurrent data abstractions, because having the receiver’s lock does not, in general, imply having the locks of its component objects that might be exposed to outside inspection or manipulation. We need a way to state and enforce constraints on aliasing. For this, we are considering a variant of the Universe type system [76,77], which would allow us to enforce such alias constraints statically in JML. The idea is to statically guarantee that all paths to an object pass through a single “owner” object. With this kind of type system, we could weaken the above precondition to state that the current thread must either hold the lock on the receiver’s owner object, or on the object itself.

## 6 Mathematical Libraries

As described in Section 1.3.4, we need to provide a library of mathematical concepts with JML in a way that does not overwhelm programmers, and yet is useful for formal verification.

### 6.1 *Hiding the Mathematics*

It is sometimes convenient to use mathematical concepts such as sets and sequences in specification, particularly for collection classes [78,79,11]. For example, the specification of `Stack` in Fig. 1 uses the type `JMLObjectSequence`, which is part of JML’s `org.jmlspecs.models` package. This package contains types that are intended for such mathematical modeling. Besides sequences, these include sets, bags, relations, and maps, and a few other convenience types.

Most types in the `org.jmlspecs.models` package have only pure methods and constructors.<sup>5</sup> For example, `JMLObjectSequence`’s `insertFront` method re-

---

<sup>5</sup> The `org.jmlspecs.models` package does have some types that have non-pure methods. These are various kinds of iterators and enumerators. The methods of these iterators and enumerators that have side effects cannot be used in specification expressions.

turns a sequence object that is like the receiver, but with its argument placed at the front; the receiver is not changed in any way. `JMLObjectSequence`'s `trailer` method similarly returns a sequence containing all but the first element of the receiver, without changing the receiver. Because such methods are pure, they can be used during runtime assertion checking without changing the underlying computation.

JML gains two advantages from having these mathematical modeling types in a Java package, as opposed to having them be purely mathematical concepts. First, these types all have Java implementations and thus can be used during runtime assertion checking. Second, using these types in assertions avoids the introduction of special mathematical notation; instead, normal Java expressions (method calls) are used to do things like concatenating sequences or intersecting sets. This is an advantage for our main audience, which consists of programmers and not mathematicians.

## 6.2 Use by Theorem Provers

The second part of the mathematical libraries problem described in Section 1.3.4 is that the library of mathematical modeling types should be useful for formal verification. The types in the `org.jmlspecs.models` package are intended to correspond (loosely) to the libraries of mathematical concepts found in theorem provers, such as PVS. As we gain experience, we can add additional methods to these types to improve their correspondence to these mathematical concepts. It is also possible to add new packages of such types tailored to specific theorem provers or to other notations, such as OCL.

When translating specification expressions into theorem prover input, the LOOP tool currently treats all methods in the same way — it does not make a special case for pure methods in the `org.jmlspecs.models` package. This makes the resulting proof obligations more complex than is desirable. Since the types in the models package are known, one should be able, as a special case, to replace the general semantics of such a method call with a call to some specific function from the theorem prover's library of mathematical concepts. To facilitate this, it may be that these model types should all be declared to be final, which is currently not the case.

## 7 Related Work

We have already discussed how JML differs from conventional formal specification languages, such as Z [17–19], VDM [20,21,5,22], the Larch family [8–11],

and RESOLVE [12,13]. To summarize, the main difference is that JML's specification expressions are based on a subset of the Java programming language, a design that is more congenial to Java programmers.

The Alloy Annotation Language (AAL) offers a syntax similar to JML for annotating Java programs [80]. AAL supports extensive compile-time checking based on static analysis techniques. Unlike similar static analysis tools such as ESC/Java [81], AAL also supports method calls and relational expressions in assertions. However, AAL's assertion language is based on a simple first-order logic with relational operators [82] and not on a subset of Java expressions. We believe that a Java-based syntax is more likely to gain acceptance among Java programmers. However, JML could adopt some of AAL's features for specifying sets of objects using regular expressions. These would be helpful in using JML's frame axioms, where they would allow JML to more precisely describe locations that can be assigned to in the method. (Another option that would have similar benefits would be to use the approach taken in DemeterJ [83].)

We have also discussed how JML differs from design by contract languages, such as Eiffel [14,15], and tools, such as APP [16]. Summarizing, JML provides better support for more exact specifications and formal verification by

- extending the set of specification expressions with more expressive mathematical constructs, such as quantifiers,
- ensuring that specification expressions do not contain side effects, and
- providing a library of types corresponding to mathematical concepts.

JML's specification-only (model) declarations and frame axioms also contribute to its ability to specify types more precisely than is easily done with design by contract tools.

We know of several other design by contract tools for Java [84–89]. The approaches vary from a simple assertion mechanism similar to the `assert` macros of C and C++ to full-fledged contract enforcement capabilities. Jass [84], iContract [88], and JContract [89] focus on the practical use of design by contract in Java. Handshake and jContractor focus on implementation techniques such as library-based on-the-fly instrumentation of contracts [85,87]. Contract Java focuses on properly blaming contract violations [86,90]. These notations and tools suffer from the same problems as Eiffel. That is, none of them guarantee the lack of side effects in assertions, handle undefinedness in a way that would facilitate formal verification and reasoning, support more expressive mathematical notations such as quantifiers, or provide a set of immutable types designed for use in specifications. In sum, they all focus on runtime checking, and thus it is difficult to write exact specifications for formal verification and reasoning.

## 8 Conclusion

JML synthesizes the best from the worlds of design by contract and more mathematical specification languages. Because of its expressive mathematical notations, its specification-only (model) declarations, and library of mathematical modeling types, one can more easily write more exact specifications in JML than in a design by contract language, such as Eiffel. These more detailed specifications, along with JML's purity checking, allow JML to be useful for formal verification. Thus, JML's synthesis of features allows it to serve many roles in the Java formal methods community.

Our experience so far is that this approach has had a modest impact. Release 4.1 of JML has been downloaded over 400 times. JML has been used in at least 5 universities for teaching some aspects of formal methods. It is used somewhat extensively in the Java Smart Card industry and has been used in at least one company outside of that industry (Fulcrum).

In the future, we would like to extend the range of tools that JML supports to include tools for model checking and specification of concurrent Java programs [91]. We invite others to join us in this effort to furnish Java programmers with a single notation that can be used by many tools.

## Acknowledgments

The work of Leavens, Cheon, Clifton, and Ruby was supported in part by the US National Science Foundation, under grants CCR-0097907 and CCR-0113181.

Thanks to Robyn Lutz, Sharon Ryan, and Janet Leavens for comments on earlier drafts of this paper. Thanks to all who have contributed to the design and implementation of JML including Al Baker, Erik Poll, Bart Jacobs, Joe Kiniry, Rustan Leino, Raymie Stata, Michael Ernst, Gary Daugherty, Arnd Poetzsch-Heffter, Peter Müller, Alexandru D. Salcianu, and others acknowledged in [2]. “Design by Contract” is a trademark of Interactive Software Engineering.

## References

- [1] G. T. Leavens, A. L. Baker, C. Ruby, JML: A notation for detailed design, in: H. Kilov, B. Rumpe, I. Simmonds (Eds.), *Behavioral Specifications of Businesses and Systems*, Kluwer Academic Publishers, Boston, 1999, pp. 175–188.

- [2] G. T. Leavens, A. L. Baker, C. Ruby, Preliminary design of JML: A behavioral interface specification language for Java, Tech. Rep. 98-06v, Iowa State University, Department of Computer Science, this is an obsolete version. (May 2003).  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR98-06/TR.ps.gz>
- [3] G. T. Leavens, K. R. M. Leino, E. Poll, C. Ruby, B. Jacobs, JML: notations and tools supporting detailed design in Java, in: OOPSLA 2000 Companion, Minneapolis, Minnesota, ACM, 2000, pp. 105–106.  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR00-15/TR.ps.gz>
- [4] C. A. R. Hoare, An axiomatic basis for computer programming, *Commun. ACM* 12 (10) (1969) 576–583.
- [5] C. B. Jones, *Systematic Software Development Using VDM*, 2nd Edition, International Series in Computer Science, Prentice Hall, Englewood Cliffs, N.J., 1990.
- [6] H. B. M. Jonkers, Upgrading the pre- and postcondition technique, in: S. Prehn, W. J. Toetenel (Eds.), *VDM '91 Formal Software Development Methods 4th International Symposium of VDM Europe Noordwijkerhout, The Netherlands, Volume 1: Conference Contributions*, Vol. 551 of *Lecture Notes in Computer Science*, Springer-Verlag, New York, NY, 1991, pp. 428–456.
- [7] H. B. M. Jonkers, *Ispec: Towards practical and sound interface specifications*, in: W. Grieskamp, T. Santen, B. Stoddart (Eds.), *Integrated Formal Methods, Second International Conference, IFM 2000, Dagstuhl Castle, Germany, November 1-3, 2000*, Vol. 1945 of *Lecture Notes in Computer Science*, Springer-Verlag, 2000, pp. 116–135.
- [8] J. V. Guttag, J. J. Horning, S. Garland, K. Jones, A. Modet, J. Wing, *Larch: Languages and Tools for Formal Specification*, Springer-Verlag, New York, NY, 1993.
- [9] L. Lamport, A simple approach to specifying concurrent systems, *Commun. ACM* 32 (1) (1989) 32–45.
- [10] G. T. Leavens, Larch frequently asked questions, version 1.110. Available in <http://www.cs.iastate.edu/~leavens/larch-faq.html> (May 2000).  
URL <http://www.cs.iastate.edu/~leavens/larch-faq.html>
- [11] J. M. Wing, Writing Larch interface language specifications, *ACM Trans. Prog. Lang. Syst.* 9 (1) (1987) 1–24.
- [12] S. H. Edwards, W. D. Heym, T. J. Long, M. Sitaraman, B. W. Weide, Part II: Specifying components in RESOLVE, *ACM SIGSOFT Software Engineering Notes* 19 (4) (1994) 29–39.
- [13] W. F. Ogden, M. Sitaraman, B. W. Weide, S. H. Zweben, Part I: The RESOLVE framework and discipline — a research synopsis, *ACM SIGSOFT Software Engineering Notes* 19 (4) (1994) 23–28.



- [14] B. Meyer, Eiffel: The Language, Object-Oriented Series, Prentice Hall, New York, NY, 1992.
- [15] B. Meyer, Object-oriented Software Construction, 2nd Edition, Prentice Hall, New York, NY, 1997.
- [16] D. S. Rosenblum, Towards a method of programming with assertions, in: Proceedings of the 14th International Conference on Software Engineering, 1992, pp. 92–104.
- [17] J. Spivey, An introduction to Z and formal specifications, Software Engineering Journal 4 (1) (1989) 40–50.
- [18] J. M. Spivey, The Z Notation: A Reference Manual, International Series in Computer Science, Prentice-Hall, New York, NY, 1989, ISBN 013983768X.
- [19] J. Woodcock, J. Davies, Using Z: Specification, Refinement, and Proof, Prentice Hall International Series in Computer Science, 1996.  
URL <http://www.comlab.ox.ac.uk/usingz.html>
- [20] J. Bicarregui, J. S. Fitzgerald, P. A. Lindsay, R. Moore, B. Ritchie, Proof in VDM: A Practitioner’s Guide, Springer-Verlag, New York, NY, 1994.
- [21] J. Fitzgerald, P. G. Larsen, Modelling Systems: Practical Tools in Software Development, Cambridge, Cambridge, UK, 1998.
- [22] I. S. Organization, Information technology – programming languages, their environments and system software interfaces – Vienna Development Method – specification language – part 1: Base language, ISO/IEC 13817-1 (Dec. 1996).
- [23] J. Warmer, A. Kleppe, The Object Constraint Language: Precise Modeling with UML, Addison Wesley Longman, Reading, Mass., 1999.
- [24] J. Warmer, A. Kleppe, OCL: The constraint language of the UML, Journal of Object-Oriented Programming 12 (1) (1999) 10–13,28.
- [25] G. Booch, J. Rumbaugh, I. Jacobson, The Unified Modeling Language User Guide, Object Technology Series, Addison Wesley Longman, Reading, Mass., 1999.
- [26] C. Larman, Applying UML and Patterns: An Introduction to Object-Oriented Analysis and Design and the Unified Process, second edition Edition, Prentice Hall PTR, Upper Saddle River, NJ, 2002.
- [27] Y. Cheon, G. T. Leavens, A runtime assertion checker for the Java Modeling Language (JML), in: H. R. Arabnia, Y. Mun (Eds.), Proceedings of the International Conference on Software Engineering Research and Practice (SERP ’02), Las Vegas, Nevada, USA, June 24-27, 2002, CSREA Press, 2002, pp. 322–328.  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR02-05/TR.pdf>

- [28] Y. Cheon, G. T. Leavens, A simple and practical approach to unit testing: The JML and JUnit way, in: B. Magnusson (Ed.), ECOOP 2002 — Object-Oriented Programming, 16th European Conference, Málaga, Spain, Proceedings, Vol. 2374 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, 2002, pp. 231–255.
- [29] L. Friendly, The design of distributed hyperlinked programming documentation, in: S. Fraïssè, F. Garzotto, T. Isakowitz, J. Nanard, M. Nanard (Eds.), Proceedings of the International Workshop on Hypermedia Design (IWH'D'95), Montpellier, France, 1–2 June 1995, Springer, 1995, pp. 151–173.  
URL [citeseer.nj.nec.com/friendly95design.html](http://citeseer.nj.nec.com/friendly95design.html)
- [30] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, R. Stata, Extended static checking for Java, in: Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation (PLDI'02), Vol. 37, 5 of SIGPLAN, ACM Press, New York, 2002, pp. 234–245.
- [31] K. R. M. Leino, G. Nelson, J. B. Saxe, ESC/Java user's manual, Technical note, Compaq Systems Research Center (Oct. 2000).
- [32] K. R. M. Leino, J. B. Saxe, R. Stata, Checking Java programs via guarded commands, Technical Note 1999-002, Compaq Systems Research Center, Palo Alto, CA (May 1999).  
URL <http://gatekeeper.dec.com/pub/DEC/SRC/technical-notes/abstracts/src-tn-1999-002.html>
- [33] M. Huisman, Reasoning about Java programs in higher order logic with PVS and Isabelle, Ipa dissertation series, 2001-03, University of Nijmegen, Holland (Feb. 2001).
- [34] M. Huisman, B. Jacobs, Java program verification via a Hoare logic with abrupt termination, in: T. Maibaum (Ed.), Fundamental Approaches to Software Engineering (FASE 2000), Vol. 1783 of LNCS, Springer-Verlag, 2000, pp. 284–303, an earlier version is technical report CSI-R9912.
- [35] B. Jacobs, J. Kiniry, M. Warnier, Java program verification challenges, in: F. S. de Boer, M. M. Bonsangue, S. Graf, W.-P. de Roever (Eds.), FMCO 2002: Formal Methods for Component Objects, Proceedings, Vol. 2852 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, 2003, pp. 202–219.
- [36] B. Jacobs, J. van den Berg, M. Huisman, M. van Berkum, U. Hensel, H. Tews, Reasoning about Java classes (preliminary report), in: OOPSLA '98 Conference Proceedings, Vol. 33(10) of ACM SIGPLAN Notices, ACM, 1998, pp. 329–340.
- [37] M. Ernst, J. Cockrell, W. G. Griswold, D. Notkin, Dynamically discovering likely program invariants to support program evolution, IEEE Transactions on Software Engineering 27 (2) (2001) 99–123.
- [38] J. W. Nimmer, M. D. Ernst, Static verification of dynamically detected program invariants: Integrating Daikon and ESC/Java, in: Proceedings of RV'01, First Workshop on Runtime Verification, Elsevier, 2001.

URL <http://people.csail.mit.edu/people/mernst/pubs/invariants-verify-rv2001.pdf>

- [39] C. Boyapati, S. Khurshid, D. Marinov, Korat: Automated testing based on Java predicates, in: Proceedings International Symposium on Software Testing and Analysis (ISSTA), ACM, 2002, pp. 123–133.  
URL <http://doi.acm.org/10.1145/566172.566191>
- [40] Y. Cheon, G. T. Leavens, The Larch/Smalltalk interface specification language, ACM Transactions on Software Engineering and Methodology 3 (3) (1994) 221–253.
- [41] Y. Cheon, G. T. Leavens, A quick overview of Larch/C++, Journal of Object-Oriented Programming 7 (6) (1994) 39–49.
- [42] G. T. Leavens, An overview of Larch/C++: Behavioral specifications for C++ modules, in: H. Kilov, W. Harvey (Eds.), Specification of Behavioral Semantics in Object-Oriented Information Modeling, Kluwer Academic Publishers, Boston, 1996, Ch. 8, pp. 121–142, an extended version is TR #96-01d, Department of Computer Science, Iowa State University, Ames, Iowa, 50011.
- [43] G. T. Leavens, Larch/C++ Reference Manual, version 5.41. 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> (Apr. 1999).  
URL <ftp://ftp.cs.iastate.edu/pub/larchc++/lcpp.ps.gz>
- [44] G. T. Leavens, A. L. Baker, Enhancing the pre- and postcondition technique for more expressive specifications, in: J. M. Wing, J. Woodcock, J. Davies (Eds.), FM'99 — Formal Methods: World Congress on Formal Methods in the Development of Computing Systems, Toulouse, France, September 1999, Proceedings, Vol. 1709 of Lecture Notes in Computer Science, Springer-Verlag, 1999, pp. 1087–1106.  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR97-19/TR.ps.gz>
- [45] G. T. Leavens, Y. Cheon, Preliminary design of Larch/C++, in: U. Martin, J. Wing (Eds.), Proceedings of the First International Workshop on Larch, July, 1992, Workshops in Computing, Springer-Verlag, New York, NY, 1993, pp. 159–184.
- [46] K. Finney, Mathematical notation in formal specification: Too difficult for the masses?, IEEE Transactions on Software Engineering 22 (2) (1996) 158–159.
- [47] C. A. R. Hoare, Notes on data structuring, in: O.-J. Dahl, E. Dijkstra, C. A. R. Hoare (Eds.), Structured Programming, Academic Press, Inc., New York, NY, 1972, pp. 83–174.
- [48] K. K. Dhara, G. T. Leavens, Forcing behavioral subtyping through specification inheritance, in: Proceedings of the 18th International Conference on Software Engineering, Berlin, Germany, IEEE Computer Society Press, 1996, pp. 258–267, a corrected version is Iowa State University, Dept. of Computer Science TR #95-20c.

- [49] P. America, Inheritance and subtyping in a parallel object-oriented language, in: J. Bezivin, et al. (Eds.), ECOOP '87, European Conference on Object-Oriented Programming, Paris, France, Springer-Verlag, New York, NY, 1987, pp. 234–242, lecture Notes in Computer Science, volume 276.
- [50] P. America, Designing an object-oriented programming language with behavioural subtyping, in: J. W. de Bakker, W. P. de Roever, G. Rozenberg (Eds.), Foundations of Object-Oriented Languages, REX School/Workshop, Noordwijkerhout, The Netherlands, May/June 1990, Vol. 489 of Lecture Notes in Computer Science, Springer-Verlag, New York, NY, 1991, pp. 60–90.
- [51] G. T. Leavens, K. K. Dhara, Concepts of behavioral subtyping and a sketch of their extension to component-based systems, in: G. T. Leavens, M. Sitaraman (Eds.), Foundations of Component-Based Systems, Cambridge University Press, 2000, Ch. 6, pp. 113–135.  
 URL <http://www.cs.iastate.edu/~leavens/FoCBS-book/06-leavens-dhara.pdf>
- [52] G. T. Leavens, D. Pigozzi, A complete algebraic characterization of behavioral subtyping, *Acta Informatica* 36 (2000) 617–663.  
 URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR96-15/TR.ps.gz>
- [53] G. T. Leavens, Verifying object-oriented programs that use subtypes, Tech. Rep. 439, Massachusetts Institute of Technology, Laboratory for Computer Science, the author's Ph.D. thesis. (Feb. 1989).
- [54] B. Liskov, J. Wing, A behavioral notion of subtyping, *ACM Trans. Prog. Lang. Syst.* 16 (6) (1994) 1811–1841.
- [55] G. T. Leavens, W. E. Weihl, Reasoning about object-oriented programs that use subtypes (extended abstract), in: N. Meyrowitz (Ed.), OOPSLA ECOOP '90 Proceedings, Vol. 25(10) of ACM SIGPLAN Notices, ACM, 1990, pp. 212–223.
- [56] G. T. Leavens, W. E. Weihl, Specification and verification of object-oriented programs using supertype abstraction, *Acta Informatica* 32 (8) (1995) 705–778.
- [57] E. Cohen, Programming in the 1990s: An Introduction to the Calculation of Programs, Springer-Verlag, New York, NY, 1990.
- [58] E. W. Dijkstra, C. S. Scholten, Predicate Calculus and program semantics, Springer-Verlag, NY, 1990.
- [59] D. Gries, F. B. Schneider, A Logical Approach to Discrete Math, Texts and Monographs in Computer Science, Springer-Verlag, New York, NY, 1994.
- [60] P. Chalin, Improving JML: For a safer and more effective language, Tech. Rep. 2003-001.1, Computer Science Department, Concordia University (March 2003).
- [61] P. Chalin, Back to basics: Language support and semantics of basic infinite integer types in JML and Larch, Tech. Rep. CU-CS 2002-003.1, Computer Science Department, Concordia University (Oct. 2002).  
 URL <http://www.cs.concordia.ca/~faculty/chalin/papers/TR-CU-CS-2002-003.1.pdf>

- [62] A. Wills, Capsules and types in Fresco: Program validation in Smalltalk, in: P. America (Ed.), ECOOP '91: European Conference on Object Oriented Programming, Vol. 512 of Lecture Notes in Computer Science, Springer-Verlag, New York, NY, 1991, pp. 59–76.
- [63] J. M. Wing, A two-tiered approach to specifying programs, Tech. Rep. TR-299, Massachusetts Institute of Technology, Laboratory for Computer Science (1983).
- [64] B. Jacobs, E. Poll, A logic for the Java modeling language JML, in: Fundamental Approaches to Software Engineering (FASE'2001), Genova, Italy, 2001, Vol. 2029 of Lecture Notes in Computer Science, Springer-Verlag, 2001, pp. 284–299.  
URL <http://www.cs.kun.nl/~erikpoll/publications/jmllogic.html>
- [65] A. D. Raghavan, G. T. Leavens, Desugaring JML method specifications, Tech. Rep. 00-03c, Iowa State University, Department of Computer Science, this is an obsolete version. (Aug. 2001).  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR00-03/TR.ps.gz>
- [66] H. Barringer, J. H. Cheng, C. B. Jones, A logic covering undefinedness in program proofs, *Acta Informatica* 21 (3) (1984) 251–269.
- [67] D. Gries, F. B. Schneider, Avoiding the undefined by underspecification, in: J. van Leeuwen (Ed.), *Computer Science Today: Recent Trends and Developments*, no. 1000 in Lecture Notes in Computer Science, Springer-Verlag, New York, NY, 1995, pp. 366–373.
- [68] G. T. Leavens, J. M. Wing, Protective interface specifications, *Formal Aspects of Computing* 10 (1998) 59–75.
- [69] D. K. Gifford, J. M. Lucassen, Integrating functional and imperative programming, in: *ACM Conference on LISP and Functional Programming*, ACM, 1986, pp. 28–38.
- [70] J.-P. Talpin, P. Jouvelot, The type and effect discipline, *Information and Computation* 111 (2) (1994) 245–296.  
URL <http://cri.ensmp.fr/~talpin/ic94.ps.gz>
- [71] K. R. M. Leino, A myth in the modular specification of programs, Tech. Rep. KRML 63, Digital Equipment Corporation, Systems Research Center, 130 Lytton Avenue Palo Alto, CA 94301, obtain from the author, at [leino@microsoft.com](mailto:leino@microsoft.com) (Nov. 1995).
- [72] C. Ruby, G. T. Leavens, Safely creating correct subclasses without seeing superclass code, in: *OOPSLA 2000 Conference on Object-Oriented Programming, Systems, Languages, and Applications*, Minneapolis, Minnesota, Vol. 35(10) of ACM SIGPLAN Notices, 2000, pp. 208–228.  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR00-05/TR.ps.gz>
- [73] C. Clifton, MultiJava: Design, implementation, and evaluation of a Java-compatible language supporting modular open classes and symmetric multiple

dispatch, Tech. Rep. 01-10, Department of Computer Science, Iowa State University, Ames, Iowa, 50011, the author's masters thesis (Nov. 2001).  
URL <ftp://ftp.cs.iastate.edu/pub/techrepts/TR01-10/TR.pdf>

- [74] C. Clifton, G. T. Leavens, C. Chambers, T. Millstein, MultiJava: Modular open classes and symmetric multiple dispatch for Java, in: OOPSLA '00: Proceedings of the 15th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications, Vol. 35(10) of ACM SIGPLAN Notices, ACM, New York, NY, 2000, pp. 130–145.  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR00-06/TR.ps.gz>
- [75] K. Arnold, J. Gosling, D. Holmes, The Java Programming Language Third Edition, Addison-Wesley, Reading, MA, 2000.
- [76] P. Müller, Modular Specification and Verification of Object-Oriented Programs, Vol. 2262 of Lecture Notes in Computer Science, Springer-Verlag, 2002, the author's Ph.D. Thesis. Available from <http://www.informatik.fernuni-hagen.de/import/pi5/publications.html>.
- [77] P. Müller, A. Poetzsch-Heffter, G. T. Leavens, Modular specification of frame properties in JML, *Concurrency, Computation Practice and Experience*. 15 (2003) 117–154.  
URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR02-02/TR.pdf>
- [78] C. A. R. Hoare, Proof of correctness of data representations, *Acta Informatica* 1 (4) (1972) 271–281.
- [79] B. Liskov, J. Guttag, Abstraction and Specification in Program Development, The MIT Press, Cambridge, Mass., 1986.
- [80] S. Khurshid, D. Marinov, D. Jackson, An analyzable annotation language, in: Proceedings of OOPSLA '02 Conference on Object-Oriented Programming, Languages, Systems, and Applications, Vol. 37(11) of SIGPLAN Notices, ACM, New York, NY, 2002, pp. 231–245.
- [81] D. L. Detlefs, K. R. M. Leino, G. Nelson, J. B. Saxe, Extended static checking, SRC Research Report 159, Compaq Systems Research Center, 130 Lytton Ave., Palo Alto (Dec 1998).
- [82] D. Jackson, Alloy: A lightweight object modeling notation, *ACM Transactions on Software Engineering and Methodology* 11 (2) (2002) 256–290.
- [83] K. Lieberherr, D. Orleans, J. Ovlinger, Aspect-oriented programming with adaptive methods, *Commun. ACM* 44 (10) (2001) 39–41.  
URL <http://portal.acm.org/toc.cfm?id=383845&type=issue&coll=portal&dl=ACM&idx=J79&part=magazine&WantType=Magazines&title=CACM>
- [84] D. Bartetzko, C. Fischer, M. Moller, H. Wehrheim, Jass - Java with assertions, in: Workshop on Runtime Verification held in conjunction with the 13th Conference on Computer Aided Verification, CAV'01, 2001, published in *Electronic Notes in Theoretical Computer Science*, K. Havelund and G. Rosu (eds.), 55(2), 2001.

- [85] A. Duncan, U. Holzle, Adding contracts to Java with Handshake, Tech. Rep. TRCS98-32, Department of Computer Science, University of California, Santa Barbara, CA (Dec. 1998).
- [86] R. B. Findler, M. Felleisen, Contract soundness for object-oriented languages, in: OOPSLA '01 Conference Proceedings, Object-Oriented Programming, Systems, Languages, and Applications, October 14-18, 2001, Tampa Bay, Florida, USA, 2001, pp. 1–15.
- [87] M. Karaorman, U. Holzle, J. Bruno, jContractor: A reflective Java library to support design by contract, in: P. Cointe (Ed.), Meta-Level Architectures and Reflection, Second International Conference on Reflection '99, Saint-Malo, France, July 19–21, 1999, Proceedings, Vol. 1616 of Lecture Notes in Computer Science, Springer-Verlag, 1999, pp. 175–196.
- [88] R. Kramer, iContract – the Java design by contract tool, TOOLS 26: Technology of Object-Oriented Languages and Systems, Los Alamitos, California (1998) 295–307.
- [89] Parasoft Corporation, Using design by contract™ to automate Java™ software and component testing, available from [http://www.parasoft.com/jsp/products/tech\\_papers.jsp?product=Jcontract](http://www.parasoft.com/jsp/products/tech_papers.jsp?product=Jcontract), as of Feb. 2003.
- [90] R. B. Findler, M. Latendresse, M. Felleisen, Behavioral contracts and behavioral subtyping, in: Proceedings of Joint 8th European Software Engineering Conference (ESEC) and 9th ACM SIGSOFT International Symposium on the Foundations of Software Engineering (FSE), September 10-14, 2001, Vienna, Austria, 2001.
- [91] E. Abraham-Mumm, F. de Boer, W. de Roever, , M. Steffen, A tool-supported proof system for multithreaded java, in: F. de Boer, M. Bonsangue, S. Graf, W.-P. de Roever (Eds.), FMCO 2002: Formal Methods for Component Objects, Proceedings, Lecture Notes in Computer Science, Springer-Verlag, 2003.