

Karlsruhe Reports in Informatics 2014,10

Edited by Karlsruhe Institute of Technology,
Faculty of Informatics
ISSN 2190-4782

Formal Specification with JML

Marieke Huisman, Wolfgang Ahrendt, Daniel Bruns, Martin Hentschel

2014



Fakultät für Informatik

Please note:

This Report has been published on the Internet under the following
Creative Commons License:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de>.

Formal Specification with JML

Marieke Huisman* Wolfgang Ahrendt† Daniel Bruns‡

Martin Hentschel§

July 14, 2014

Abstract

This text is a general, self contained, and tool independent introduction into the *Java Modeling Language*, JML. It is a preview of a chapter planned to appear in a book about the KeY approach and tool to the verification of Java software. JML is the dominating starting point of KeY style Java verification. However, this paper does not in any way depend on any tool nor verification methodology. Other chapters in this book talk about the usage of JML in KeY style verification. Here, we only refer to KeY in very few places, without relying on it. This introduction is written for all readers with an interest in formal specification of software in general, and anyone who wants to learn about the JML approach to specification in particular. The authors appreciate any comments or questions that help to improve the text.

*`m.huisman@utwente.nl`, University of Twente, Enschede, The Netherlands

†`ahrendt@chalmers.se`, Chalmers University of Technology, Gothenburg, Sweden

‡`bruns@kit.edu`, Karlsruhe Institute of Technology, Karlsruhe, Germany

§`hentschel@cs.tu-darmstadt.de`, Technical University of Darmstadt, Darmstadt, Germany

Contents

1	Introduction	5
2	Method Contracts (Part 1)	6
2.1	Clauses of a Contract	7
2.2	Defensive versus Offensive Method Implementations	10
2.3	Specifications and Implementations	10
3	Expressions	11
3.1	Quantified Boolean Expressions	11
3.2	Numerical Comprehensions	13
3.3	Evaluation in the Prestate	13
4	Method Contracts (Part 2)	14
4.1	Visibility of Specifications	14
4.2	Specification Cases	15
4.3	Semantics of Normal Behavior Specification Cases	15
4.4	Specifications for Constructors	17
4.5	Notions of Purity	18
5	Class Level Specifications	18
5.1	Visible States	19
5.2	Invariants	20
5.3	Initially Clauses	24
5.4	History Constraints	24
5.5	Static Class Specifications	26
5.6	Inheritance of Specifications	26
6	Nonnull versus Nullable Object References	28
7	Exceptional Behavior	29
8	Specification Only Class Members	32
8.1	Model Fields and Model Methods	32
8.2	Ghost Variables	33
8.3	Ghost Variables vs Model Fields	36
9	Integer Semantics	37
10	Specification for Verification	38
10.1	Framing	39
10.2	Loop Invariants	39
10.3	Assertions and Block Contracts	42
11	Conclusion	43
11.1	Tool Support for JML	44
11.2	Other Program Annotation Languages	45

1 Introduction

The *Java Modeling Language*, JML, is an increasingly popular specification language for Java software, that has been developed as a community effort since 1999. The nature of such a project entails that language details change, sometimes rapidly, over time and there is no ultimate reference for JML. Fortunately, for the items that we address in this introduction, the syntax and semantics are for the greatest part already settled in [Leavens et al., 2013]. Basic design decisions have been described in [Leavens et al., 2006],¹ that outlines these three overall goals:

- “JML must be able to document the interfaces and behavior of existing software, regardless of the analyses and design methods to create it. [...]
- The notation used in JML should be readily understandable by Java programmers, including those with only standard mathematical training. [...]
- The language must be capable of being given a rigorous formal semantics, and must also be amenable to tool support.”

This essentially means two things to the specification language: Firstly, it needs to express properties about the special aspects of the Java language, e.g., inheritance, object initialization, or abrupt termination. Secondly, the specification language itself heavily relies on Java; its syntax extends Java’s syntax and its semantics extend Java’s semantics. The former makes it convenient to talk about such features in a natural way instead of defining auxiliary constructs or instrumenting the code as in other specification methodologies. The latter can also come in handy since, with a reasonable knowledge of Java, little theoretical background is needed in order to use JML. This has been one of the major aims in the design of JML. It however bears the problem that reasoning about specifications in a formal and abstract way becomes more difficult as even simple expressions are evaluated w.r.t. the complex semantics of Java.

Assertions in source code to prove correctness of the implementation have already been proposed long time ago [Floyd, 1967]. However, assertions were not widely used in practice—the `assert` statement in Java only first appeared in version 1.4, in 2002. Other programming languages adopted assertions earlier: Bertrand Meyer introduced the concept of *Design by Contract* (DbC) in 1986 with the Eiffel language [Meyer, 1992, 1997]. DbC is a programming methodology where the behavior of program components is described as a *contract* between the provider and the clients of the component. The client only has to study the component’s contract, and this should tell him or her exactly what he or she can expect from the component. The provider is free to choose any implementation, as long as it respects the component’s contract. Design by Contract has become a popular methodology for object-oriented languages. In this case, the components are the program’s classes. Contracts naturally correspond with the object-oriented paradigm to hide (or encapsulate) the internal state of an object.

The Eiffel compiler came with a special option to check validity of a contract at runtime. Subsequently, the same ideas were applied to reason about other programming languages (including Modula III, C++, and Smalltalk, that were all handled in the Larch project [Guttag and Horning, 1993, Leavens and Cheon, 1993]). With the growing popularity of Java, several people decided to develop a specification language

¹This 2006 journal publication is a revised version of a technical report that first appeared in 1998.

for Java. Gary Leavens and his students at Iowa State University used their experience from the Larch project, and started work on a DbC specification language for Java in 1998. They proposed a specification language, and simultaneously developed a JML runtime assertion checker, that could be used to validate the contracts at runtime. At more or less the same time, Rustan Leino and his team at the DEC/Compaq research centre started working on a tool for static code analysis. For the *Extended Static Checker for Java*, ESC/Java [Leino et al., 2000], they developed a specification language that was more or less a subset of JML. A successor, ESC/Java2 [Cok and Kiniry, 2005], finally adopted JML as it is now. Several projects have been targeting tool supported formal verification of Java programs: the LOOP project [van den Berg and Jacobs, 2001], the Krakatoa project [Marché et al., 2004], and of course KeY. While in KeY originally specifications had been written in OCL, from version 0.99 (released in 2005) on, JML has been the primary input language².

Ever since, the community has worked on adopting a single JML language, with a single semantics—and this is still an ongoing process. Over the years, JML has become a very large language, containing many different specification constructs, some of which are only sensible in a single analysis technique. Because of the language being so large, not for all constructs the semantics is actually understood and agreed upon, and moreover all tools that support JML in fact only support a subset of it. There have been several suggestions of providing a formal semantics [Jacobs and Poll, 2001, Engel, 2005, Darvas and Müller, 2007, Bruns, 2009], but as of 2014, there is no final consensus. Moreover, JML suffers from the lack of support for current Java versions; currently there are no specifications for Java 5 features, such as enums or generic types. Dedicated expressions to deal with enhanced foreach loops have been proposed in [Cok, 2008].

2 Method Contracts (Part 1)

Specifications, whether they are formulated in natural language or some formalism, can express properties about system artifacts on various levels of granularity, like for instance the overall system, some intermediate level, like architectural components, or, on an even finer level of granularity, source code units. JML is designed for *unit* specification. In Java, those units are:

- methods, where JML specifies the effect of a single method invocation;
- classes, where JML merely specifies constraints on the internal structure of an object; and
- interfaces, where JML specifies the external behavior of an object.

Specifications of these units serve as *contracts* for their implementers, fixing what they can rely upon, and what they have to deliver in return, following the aforementioned Design by Contract paradigm.

We start by introducing method specifications in this section. While we go along, we will also introduce more general concepts, such as JML expressions, that are later used for class and interface specifications as well.

²Support for OCL has been dropped in KeY 2.0.

2.1 Clauses of a Contract

Contracts of methods are an agreement between the *caller* of the method and the *callee*, describing what guarantees they provide to each other. More specifically, it describes what is expected from the code that calls the method, and it provides guarantees about what the method will actually do. While in our terminology, ‘contract’ refers to the complete behavioral specification, written JML specifications usually consist of *specification cases*. These specification cases are made up of several *clauses*.

The expectations on the caller are called the *preconditions* of the method. Typically, these will be conditions on the method’s parameters, e.g., an argument should be a nonnull reference, but the precondition can also describe that the method should only be called when the object is in a particular state. In JML, each precondition is preceded by the keyword **requires**, and the conjunction of all requires clauses forms the method’s precondition. We would like to emphasise that it is not the method implementer’s responsibility to check or handle a violation of the precondition. Instead, this is the responsibility of the caller, and the whole point of contracts is to make this distribution of responsibilities explicit, and checkable. Having said that, it can be a difficult design decision when the caller should be responsible for ‘good’ parameters and prestates, and when the called method should check and handle this itself. We refer to Sect. 2.2 for a further discussion of *defensive* versus *offensive* specifications and implementations.

The guarantees provided by the method are called the *postcondition* of the method. They describe how the object’s state is changed by the method, or what the expected return value of the method is. A method only guarantees its postcondition to hold whenever it is called in a state that respects the precondition. If it is called in a state that does not satisfy the precondition, then no guarantee is made at all. In JML, every postcondition expression is preceded by the keyword **ensures**, and the conjunction of all ensures clauses forms the method’s postcondition.

JML specifications are written as special comments in the Java code, starting with `/*@` or `/**@`. The `@` symbol allows the JML parser to recognise that the comment contains a JML specification. Sometimes, JML specifications are also called *annotations*, because they annotate the program code. The preconditions and postconditions are basically just Java expressions (of boolean type). This is done on purpose: if the specifications are written in a language that the programmer is already familiar with, they are easier for him to write and to read. JML extends Java’s syntax; almost every side effect free Java expression, i.e., that does not modify the state and has no observable interaction with the outside world, (cf. [Gosling et al., 2013]) is also a valid JML expression. See Sect. 3 for a detailed discussion of JML expressions.

Example 1. Fig. 1 contains an example of a basic JML specification. It contains specification cases for the methods in an interface `Student`, modeling a typical student at some university.

We discuss the different aspects of this example in full detail.

- To specify a certain method with JML, **requires** and **ensures** clauses are placed immediately before that method, within a JML comment, starting with `/*@` or `/**@`. For instance, the method `changeStatus` is specified in JML using two pre- and two postconditions.
- The `@` symbol is not only used at the beginning of a JML comment, but possibly also at the beginning of each line of the JML specification, and before the `*/`. This is not necessary, but helps to highlight the JML specifications better.

```

1 public interface Student {
2
3     public static final int bachelor = 0;
4     public static final int master = 1;
5
6     public /*@ pure @*/ String getName();
7
8     /*@ ensures \result == bachelor || \result == master;
9     public /*@ pure @*/ int getStatus();
10
11     /*@ ensures \result >= 0;
12     public /*@ pure @*/ int getCredits();
13
14     /*@ ensures getName().equals(n);
15     public void setName(String n);
16
17     /*@ requires c >= 0;
18     @ ensures getCredits() == \old(getCredits()) + c;
19     @*/
20     public void addCredits(int c);
21
22     /*@ requires getCredits() >= 180;
23     @ requires getStatus() == bachelor;
24     @ ensures getCredits() == \old(getCredits());
25     @ ensures getStatus() == master;
26     @*/
27     public void changeStatus();
28
29
30 }

```

Listing 1: First JML example specification Student

In general, a @ is ignored within a JML annotation if it is the first (non-white) character in the line, or if it is the last character before ‘*/’.

- Requires resp. ensures clauses always consist of the keyword **requires** resp. **ensures**, followed by a boolean expression.
- For method `getName`, we specify that it is a **pure** method, i.e., it may not have any (visible) side effects, and it must terminate unconditionally (possibly with an exception). Only pure methods may be used in specification expressions, because these should not have side effects, and always terminate.
- Method `getStatus` is also specified as being pure. In addition, we specify that its result may only be one of two values: `bachelor` or `master`. To denote the return value of the method, the reserved JML keyword `\result` is used.
- For method `getCredits` we also specify that it is pure, and in addition we specify that its return value must be non-negative; a student thus never can have a negative amount of credits.
- Method `setName` is non-pure, i.e., it may have side effects. Its postcondition is expressed in terms of the pure methods `getName` and `equals`: it ensures that after termination the result of `getName` is equal to the parameter `n`.
- Method `addCredits`'s precondition states a condition on the method parameters, namely that only a positive number of credits can be added. The postcondition specifies how the credits change. Again, this postcondition is expressed in terms of a pure method, namely `getCredits`. Notice the use of the keyword `\old`. An expression `\old(E)` in the postcondition actually denotes the value of expression `E` in the state where the method call started, the *prestate* of the method. Thus the postcondition of `addCredits` expresses that the number of credits only increases: after evaluation of the method, the value of `getCredits` is equal to the old value of `getCredits`, i.e., before the method was called, plus the parameter `c`.
- Method `changeStatus`'s precondition specifies that this method only may be called when the student is in a particular state, namely he has obtained a sufficient amount of credits to pass from the Bachelor status to the Master status. Moreover, the method may only be called when the student is still having a Bachelor status. The postcondition expresses that the number of credits is not changed by this operation, but the status is. Notice that the two preconditions and the two postconditions of `changeStatus` are written as separate **requires** and **ensures** clauses, respectively. Implicitly, these are assumed to be joined by conjunction, thus the specification is equivalent to the following specification:

```
/*@ requires getCredits() >= 180 &  
  @         getStatus() == bachelor;  
  @ ensures getCredits() == \old(getCredits()) &  
  @         getStatus() == master;  
  @*/  
public void changeStatus();
```

The reader might have wondered why not all method specifications in `Student` have a pre- and a postcondition. Implicitly though, they have. For every specification clause, there is a default. For pre- and postconditions this is the predicate **true**, i.e., no constraints are placed on the caller of the method, or on the method's implementation.

Example 2. Thus for example the specification of method `getStatus` actually is the following:

```
/*@ requires true;
   @ ensures status == bachelor || status == master;
   */
public int getStatus() {
    return status;
}
```

2.2 Defensive versus Offensive Method Implementations

An important point about method contracts is that they can be used to avoid defensive programming. Consider the specification of method `addCredits` in Listing 1,

This method assumes that its argument is nonnegative, and otherwise it is not going to function correctly. When one uses a defensive programming style, one would first test the value of the argument and throw an exception if this was negative. This clutters up the code, and in many cases it is not necessary. Instead, using specifications, one can use an ‘offensive’ coding style. The specification states what the method requires from its caller. It only guarantees to function correctly if the caller also fulfills its part of the contract. When validating the application, one checks that every call of the method is indeed within the bounds of its specification, and thus the explicit test in the code is not necessary. Thus, making good use of specifications can avoid adding many parameter checks in the code. Such checks are only necessary when the parameters cannot be controlled—for example, because they are given via an external user.

2.3 Specifications and Implementations

Notice that the method specifications are written independently of possible implementations. Classes that implement this interface may choose different implementations, as long as it respects the specification. Method specifications do not always have to specify the exact behavior of a method; they give minimal requirements that the implementation should respect.

Example 3. Considering the specification in Listing 1 again, the method specification for `changeStatus` prescribes that the credits may not be changed by this method. However, method `addCredits` is free to update the status of the student. So for example, an implementation that silently updates the status from Bachelor to Master whenever appropriate is according to the specification. Notice that the specification case is repeated here for understandability and that it is not required and recommended to copy specifications of interfaces in the classes that realize them.

```
/*@ requires c >= 0;
   @ ensures getCredits() == \old(getCredits()) + c;
   */
public void addCredits(int c) {
    credits = credits + c;
    if (credits >= 180) {status = master;}
}
```

Notice also that both `addCredits` and `changeStatus` would be free to change the name of the student, according to the specification, even though we would typically not

expect this to happen. A way to avoid this, is to add explicitly conditions `getName().equals(\old(getName()))` to all postconditions. Later, in Sect. 10.1, we will see how **assignable** clauses can be used to explicitly disallow these unwanted changes in a more convenient way.

3 Expressions

We have already seen that standard Java expressions can be used in JML specifications. These expressions have to be side effect free, thus for example assignments, or increment/decrement operators, are not allowed. As also mentioned above, JML expressions may contain method calls to pure methods.

In addition, JML defines several specification-specific constructs, to be used in expressions. The use of the **\result** and **\old** keywords has already been demonstrated in Listing 1, and the official language specification contains a few more of these. Besides Java's logical operators, such as conjunction **&**, disjunction **|** and negation **!**, also other logical operators are allowed in JML specifications, e.g., implication **==>**, and logical equivalence **<==>**. Since expressions are not supposed to have side effects or terminate exceptionally, in JML in many cases the difference between logical operators such as **&** and **|**, and short circuit operators, such as **&&**, and **||** is not important. However, sometimes the short circuit operators have to be used to ensure an expression is welldefined. For instance, `y != 0 & x/y == 5` may not be a welldefined expression, while `y != 0 && x/y == 5` is.

3.1 Quantified Boolean Expressions

However, for specifying interesting properties, purely propositional boolean expressions are too limited. How could one for instance express any of the following properties with just propositional connectors?

- An array `arr` is sorted.
- The variable `m` holds the maximum entry of array `arr`.
- All `Account` objects in an array `allAccounts` are stored at the index corresponding to their respective `accountNumber` field.

Given that the arrays in these examples have a statically unknown length, propositional connectives are not enough to express any of the above. What we need here is *quantification*. For that, boolean JML expressions are extended by the following constructs³.

- **(\forallall T x; b)**
“for all x of type T , b holds”
- **(\forallall T x; a; b)**
“for all x of type T fulfilling a , b holds”
- **(\existsexists T x; b)**
“there exists an x of type T such that b holds”
- **(\existsexists T x; a; b)**
“there exists an x of type T fulfilling a , such that b holds”

³The JML keywords **\forallall** and **\existsexists** correspond to \forall and \exists in text-book notation.

Here, T is a Java (primitive or reference) type, x is any name (hereby declared to be of type T), and a and b are boolean JML expressions. The a is called *range predicate*. The two forms using a range predicate are not strictly needed, as they can be expressed without. $(\forall \text{forall } T \ x; \ a; \ b)$ is logically equivalent to $(\forall \text{forall } T \ x; \ a \implies b)$, and $(\exists \text{exists } T \ x; \ a; \ b)$ is logically equivalent to $(\exists \text{exists } T \ x; \ a \ \&\& \ b)$. However, the range predicates have a certain pragmatics not shared by their logical counterparts. In $(\forall \text{forall } T \ x; \ a; \ b)$, as well as in $(\exists \text{exists } T \ x; \ a; \ b)$, the boolean expression a is used intuitively to restrict range of x further than T does.

Example 4. Using quantifiers, we can specify that an array should be sorted, for instance in a precondition for a logarithmic lookup method that assumes sorting.

```
//@ requires (\forallall int i, j;
//@           0 <= i & i < j & j < a.length;
//@           a[i] <= a[j]);
public int lookup(int elem) {...
```

The first argument ‘`int i, j`’ is the declaration of the variables over that the quantification ranges. The (optional) second argument ‘`0 <= i & i < j & j < a.length`’ defines the range of the values for this variable, and the third argument is the actually universally quantified formula (‘`a[i] <= a[j]`’ in this case).

Example 5. An alternative, but less preferred, way to phrase the specification in Example 4 is the following:

```
//@ requires (\forallall int i, j;
//@           0 <= i & i < j & j < a.length ==> a[i] <= a[j]);
public int lookup(int elem) {...
```

Besides supporting readability, the range predicate form helps certain JML tools to ‘execute’ quantified formulas where possible. This is less important for theorem provers, like KeY. But a runtime verification tool would need to operationalise the precondition, by looping through all i, j fulfilling $0 \leq i \ \& \ i < j \ \& \ i < a.length$, instead of looping through all i, j between `Integer.MIN_VALUE` and `Integer.MAX_VALUE`.

Example 6. To specify that a method returns the index of an integer array `arr` holding the maximum entry, we can write the following postcondition.

```
//@ ensures (\forallall int i; 0 <= i && i < arr.length; \result >= arr[i]);
```

But is that enough? (The reader may briefly reflect before reading on.) An implementation always returning `Integer.MAX_VALUE` would satisfy the above postcondition⁴. We therefore need an additional postcondition:

```
//@ ensures arr.length > 0 ==>
//@         (\exists int i; 0 <= i && i < arr.length; \result == arr[i]);
```

Example 7. The following boolean JML expressions says that all `Account` objects in an array `allAccounts` are stored at the index corresponding to their respective `accountNumber` field.

```
(\forallall int i; 0 <= i && i < allAccounts.length;
    allAccounts[i].accountNumber == i)
```

Such an expression could for instance be used in an invariant, see Sect. 5.2.

⁴See also Sect. 9 for a discussion on Java integers.

3.2 Numerical Comprehensions

In addition to the boolean quantified expressions, JML offers so called *generalized quantifiers* `\sum`, `\product`, `\min`, `\max`, and `\num_of`. Those are actually numerical comprehensions (or higher order functions) with bound variables. The postcondition in Example 6 can alternatively be given as:

```
//@ ensures \result == (\max int i; 0 <= i && i < arr.length; arr[i]);
```

Notice that is syntactically similar to a quantified formula: the `\max` operator binds a variable `i`, and a boolean guard expression restricts it to be within the range of the array's indices. The type of the `\max` expression is the type of its body; here it is the type of `arr[i]`. The intuitive semantics is obviously that the result is the maximum of all `arr[i]` where `i` is in the array range. However, the `\max` construct is not total, i.e., it is not always a welldefined expression. In case `arr` has zero length, for instance, there is no maximum. A similar case appears with a noncompact range, e.g., the set of all mathematical integers (represented by the JML type `\bigint`, see Sect. 9): `(\max \bigint i; true; i)`.

Another comprehension operator is the summation operator `\sum`, of which we make use in Example 8 on page 31 since the exact number of summands is not known:

```
(\sum int i; 0 <= i && i < s1.length; s1[i].getCredits())
```

In the common mathematical notation, this expression can be given as $\sum_{i=0}^{s1.length-1} s1[i].g\dots$. More generally, sum comprehensions in JML can have several bound variables that range over sets of values. The general pattern is `(\sum T x; P; Q)` where T is a type, P a boolean expression and Q an integer expression corresponds to $\sum_{x \in \{y \in T | P\}} Q$. Likewise the `\product` operator is used to express product comprehensions. Since addition (as multiplication) is commutative and associative, there is no particular order in which elements are summed up. Sums with empty ranges have the value 0 by definition, empty products have value 1.

Expressions using the `\num_of` operator, that gives the cardinality of a finite set, can be expressed in terms of sums: `(\num_of T x; P)` is syntactic sugar for `(\sum T x; P; 1)`.

However, like the maximum, sum comprehensions are not always welldefined. For instance, the expression `(\sum \bigint i; 0 <= i; i)` corresponds to $\sum_{i=0}^{\infty} i$, the value of which is undefined since it diverges. In some tools—including KeY—effective reasoning about these comprehensions is therefore restricted to closed integer intervals, for which sums, etc., are always defined. In particular, KeY only interprets sums of the shape `(\sum int i; ℓ <= i && i < u ; Q)`, where the lower bound ℓ is included and the upper bound u is excluded. This restricted form using intervals has the advantage of having a simple induction schema to define these comprehensions, that lays the foundation to reasoning about sums and products. More details about this are discussed in Sect. ??.

3.3 Evaluation in the Prestate

As indicated in the introductory example, JML allows to mark any expression e in a postcondition with `\old(e)`, which means that e is not evaluated in the current (post)state of the method, but in its prestate. In most cases, `\old(e)` is a subexpression of some bigger expression, and it is important to be aware that all parts of the expression not included in `\old(...)` construct are evaluated in the current (post)state. This is fairly obvious in many examples, like `ensures getCredits() == \old(getCredits()) + c;`

in Fig. 1. For a slightly more subtle example, consider an ATM scenario, where an `insertedCard` (represented by an object with a boolean field `invalid`) is ‘confiscated’ after too many failed attempts to enter the correct PIN, specified by

```
//@ ...
//@ ensures insertedCard == null;
//@ ensures \old(insertedCard).invalid;
//@ ...
```

We encourage the reader, before reading on, to reflect on the difference between `\old(insertedCard).invalid` and `\old(insertedCard.invalid)`.

Writing `\old(insertedCard.invalid)` would mean that the method implementation has to guarantee that the `invalid` field of the old `insertedCard` object was true *before* the method’s execution. This makes no sense, as a method implementation can never influence its prestate. However, `\old(insertedCard).invalid` makes much more sense, as an implementation can, for instance, set the `invalid` field of `insertedCard` object to `true` and afterward set `insertedCard` to `null`. To demand the invalidation of the former `insertedCard` object in the poststate, `\old(insertedCard).invalid` refers to the *current* field of the object *formerly* referred to by `insertedCard`.

4 Method Contracts (Part 2)

Now that the reader is familiar with the particular features of JML expressions, we are ready to continue the presentation of method contracts.

4.1 Visibility of Specifications

So far, the specifications that we have seen have not specified anything about the values of an object’s instance variables. Typically, these are declared private, which limits also their use within specifications. Basically JML uses the same access rules like Java which means that elements used within specifications have to be visible to it and that a specification itself also has a visibility. The access modifiers `public`, `protected` and `private` are explicitly used to define specifications visibility. If none of these modifiers is used a specification has the default (package) visibility.

In addition to the Java access rules, JML forbids the usage of elements within specifications that are less visible than the specification itself. The reason of this restriction is simply that a reader of a specification may need the whole knowledge of all used elements to understand it. As a consequence, it is for instance not possible to use private variables directly within protected or public specifications. However, it is possible to change their visibility only for the specification layer via `spec_protected` or `spec_public`.⁵

Example 8. If we specify the instance variables of class `CStudent` to be `spec_public`, then its constructor can also be specified as in Listing 2.

A second restriction of specification visibility to keep in mind is that specifications that constrain a field must have at least the visibility as the field itself, so they cannot be more hidden. The reason is that otherwise a user of a field would not see the constraints to maintain. This is especially important for invariants and constraints, discussed in Sect. 5.2 and Sect. 5.4.

⁵A model field could be used as alternative, which also allows to use different fields in different implementations.

```

class CStudent implements Student {

    /*@ spec_public @*/ private String name;
    /*@ spec_public @*/ private int credits;
    /*@ spec_public @*/ private int status;
    ....

    /*@ requires c >= 0;
       @ ensures credits == c;
       @ ensures status == bachelor;
       @ ensures name = n;
       @*/
    public CStudent (int c, String n) {
        credits = c;
        name = n;
        status = bachelor;
    }
}

```

Listing 2: Class CStudent with `spec_public` variables

4.2 Specification Cases

When specifying a method, it is often useful, and sometimes necessary, to describe the behavior separately for different parts of the prestate/input space. The structuring mechanism for that is the *specification case*, each of which is specific for a particular precondition. Specification cases are combined by the `also` keyword. The above method contracts consisted of only one specification case. We now give an example where two specification cases are given for one method.

Example 9. Listing 3 shows the specification of a class implementing a set of integers, with a limited capacity that is fixed at the time when the integer set object is constructed.

Here, the method `add` is specified by two specification cases, one for the case where the set is not full, and the element to be added is not contained (`size < limit && !contains(elem)`), and one for the case where the set is full (`size == limit`) or the element to be added is already contained (`contains(elem)`).

Note that it is possible to specify `add` with only one specification case. Confer to [Raghavan and Leavens, 2000] for a procedure to produce flat specifications.

Listing 3 is furthermore an example for extensive usage of quantification. Moreover, it demonstrates the power of pure methods. Without the ability to use `contains` in the specification of the other methods, all the occurrences of `contains` would need to be replaced by the existentially quantified JML expression specifying `contains`, resulting in a much more complicated specification. We will extend on this example when discussing class invariants.

4.3 Semantics of Normal Behavior Specification Cases

An important question is when a method specification is actually satisfied. And in particular, if a method does not terminate, does it then satisfy its specification? The

```

1 public class LimitedIntegerSet {
2     public final int limit;
3     /*@ spec_public @*/ private int arr[];
4     /*@ spec_public @*/ private int size = 0;
5
6     public LimitedIntegerSet(int limit) {
7         this.limit = limit;
8         this.arr = new int[limit];
9     }
10
11     /*@ requires size < limit && !contains(elem);
12        @ ensures \result == true;
13        @ ensures contains(elem);
14        @ ensures (\forallall int e;
15           @             e != elem;
16           @             contains(e) <==> \old(contains(e)));
17        @ ensures size == \old(size) + 1;
18        @
19        @ also
20        @
21        @ requires (size == limit) || contains(elem);
22        @ ensures \result == false;
23        @ ensures (\forallall int e;
24           @             contains(e) <==> \old(contains(e)));
25        @ ensures size == \old(size);
26        @*/
27     public boolean add(int elem) { /*...*/ }
28
29     /*@ ensures !contains(elem);
30        @ ensures (\forallall int e;
31           @             e != elem;
32           @             contains(e) <==> \old(contains(e)));
33        @ ensures \old(contains(elem))
34           ==> size == \old(size) - 1;
35        @ ensures !\old(contains(elem))
36           ==> size == \old(size);
37        @*/
38     public void remove(int elem) { /*...*/ }
39
40     /*@ ensures \result == (\exists int i;
41        @             0 <= i && i < size;
42        @             arr[i] == elem);
43        @*/
44     public /*@ pure @*/ boolean contains(int elem) { /*...*/ }
45
46     // other methods
47 }

```

Listing 3: Specifying (limited size) integer set

specifications as we have seen here implicitly state that the method must always terminate, i.e., they specify a *total correctness* condition, cf. Hoare [1969]. If method `m` is specified as follows:

```
/*@ requires P;  
  @ ensures Q;  
  */  
public ... m(...) { ...
```

this means the following:

If method `m` is executed in a pre-state where `P` holds, *then* execution of method `m` from this pre-state terminates, *and*—if it terminates normally⁶—in the final state the postcondition `Q` holds.

To specify that a method may not terminate under some precondition, one can add an explicit `diverges` clause. A `diverges` clause specifies under which conditions a method may not terminate, for example to express that for certain parameters a method may not terminate. As we have seen above, the default is `false`, i.e., a method must always terminate.

```
/*@ requires P;  
  @ ensures Q;  
  @ diverges x < 0;  
  */  
public ... m(int x) { ...
```

Sometimes we wish to exclude the case that a method may terminate because of an exception. In this case, the respective specification case is preceded by the keyword `normal_behavior`, that states that the method execution must terminate normally, and in the final state the postcondition must hold.

The JML reference manual Leavens et al. [2013] further distinguishes between so called *lightweight* and *heavyweight* specifications. Heavyweight specification cases are preceded by one of the keywords `behavior`, `normal_behavior`, or `exceptional_behavior` (see Sect. 7); all others are lightweight. The difference is that, in lightweight specifications, there are no standardized defaults—except for `diverges` which default is always `false`. Instead, every tool is free to choose its own semantics. KeY takes the choice of applying the same defaults as for heavyweight specifications.

The visibility of a lightweight specification case is always the one of the method they specify.

4.4 Specifications for Constructors

Constructors can be considered as special methods. In the pre-state of a constructor, the object does not yet exist. Thus a precondition of a constructor can only put constraints on the constructor parameters, it cannot require anything about the internal state of the object—as the object does not exist yet when the constructor is called. However, the postcondition of the constructor can specify constraints on the state of the object. Typically, it will relate the object state to the constructor’s parameters.

⁶A method is said to terminate normally if either it reached the end of its body, in a normal state, or it terminated because of a `return` instruction. Below, in Sect. 7 we discuss how we can specify methods that terminate because of an exception.

Example 10. Suppose we have a class `CStudent`, implementing the `Student` interface. It could have the following constructor:

```
/*@ requires c >= 0;
   @ ensures getCredits() == c;
   @ ensures getStatus() == bachelor;
   @ ensures getName() = n;
   */
CStudent (int c, String n) {
    credits = c;
    name = n;
    status = bachelor;
}
```

Thus, to repeat, it would be incorrect to specify e.g., `requires getCredits() >= 0;` or `requires getStatus() == bachelor`—these specifications are meaningless at the moment that the constructor is invoked.

4.5 Notions of Purity

Above in Sect.2.1, we have said that only pure methods may be used in a method specification, and purity was defined as terminating unconditionally and having no visible side effects. ‘No visible side effects’ means that the state that was allocated on the heap before the method call may not be changed. Thus, this does not exclude that a method creates a new object and initialises it. In the same way, constructors are pure if they only operate on fields of the object they initialize, not touching the state that was allocated before the call to the constructor. If it, however, changes other parts of the state it is not pure. For clarity, this notion of purity in standard JML is sometimes known as *weak purity*. This is in contrast to *strict purity*, that requires that the heap is not changed in any way. While weakly and strictly pure methods have the same observable behavior, reasoning about hidden changes in weakly pure methods can make a proof more complicated. In KeY’s extension to JML, strict purity is indicated by the modifier `strictly_pure`.

Apart from that, there are situations where methods are technically speaking not pure, but from a client point of view may be considered to be so. Consider for an example the function that computes a hashcode. The first time this function is called on an object, a field of the object will be written, so that the next calls can be evaluated by looking up this field. Because of this, different notions of purity and *observational purity* exist in the literature [Barnett et al., 2004, 2005c, Darvas and Müller, 2005, Darvas and Leino, 2007, Naumann, 2007, Cok and Leavens, 2008]. For the scope of this chapter, it is sufficient to define purity simply as not having any side effects.

While pure methods must terminate under any circumstance, they may still raise exceptions or have nontrivial precondition. In these cases, the value of a pure method invocation is not always welldefined. Therefore, it is a best practice to have `true` as the precondition of pure methods and to rule out any exceptions.

5 Class Level Specifications

Consider again the specification of `Student` in Listing 1. If we look carefully at the specifications and the description that we give about the student’s credits, we notice

that implicitly we assume some properties about the value of `getCredits` that hold throughout. For example, we wrote above:

“a student thus never can have a negative amount of credits”

and also

“the number of credits only increases.”

But if we would like to make explicit that we assume that these properties always hold, we would have to add this to *all* the specifications in `Student`, and thus in particular also to all methods that do not relate at all to the number of credits. Thus for example, we would get the following specification:

```
/*@ requires getCredits() >= 0;
   @ ensures \result == bachelor || \result == master;
   @ ensures getCredits() >= 0;
   @*/
/*@ pure @*/ public int getStatus();
```

Clearly, this is not desired, because specifications would get very large, and besides describing the intended behavior of that particular method, they also describe properties over the lifetime of the object.

Therefore, JML provides also class level specifications, such as invariants, history constraints and initially clauses. These specify properties over the internal state of an object, and how the state can evolve during the object’s lifetime.

5.1 Visible States

To define in which state invariants hold, JML uses the notion of *visible states* [Poetzsch-Heffter, 1997], that are states reached throughout the execution of a code fragment. In other contexts, e.g., older versions of KeY, the semantics of invariants are based on *observed states* [Beckert et al., 2007, Sect. 8.2]. These two approaches are based on different paradigms. A principle difference is that visible states are not necessarily meant to be visible to an observer, but rather to semantical objects of the program. The targets of visibility, i.e., the objects for which a state is visible, are determined from the running execution and its receivers. The rationale behind this is, that it is primarily intended to impose strong invariants, i.e., that are obliged to hold in *every* intermediate state, but secondarily to allow temporary violations of invariants if the ‘violated object’ is a current method receiver (or of a type on that a static method is invoked). Following the observable state approach, on the other hand, invariants that hold at the beginning of a method invocation also hold at the end. This means that the exact pre- and poststates are the only states *observable*. Visible states are intermediate states in this sense. Following to [Leavens et al., 2013, Sect. 8.2] they are defined as follows (a formalization of this can be found in [Bruns, 2009]).

Definition 1 (Visible states). A state is visible to an object o if it is reached at one of the following moments in a program’s execution:

1. at the end of a constructor invocation that is initializing o ,
2. at the beginning of a finalizer invocation that is finalizing o ,
3. at the beginning and the end of a nonstatic method invocation with o as the receiver, i.e., a method like $o.m$ is called,

4. at the beginning and the end of a static method invocation of a class C where o is an instance of type C , or
5. when none of the aforementioned invocations are in progress.

The crucial one seems to be the last item. This could be seen as overly strict as it seems us to require to check the invariants of nearly every object in every state reached throughout execution. But if we consider a situation in which a class declaration contains public fields, it is desirable to secure they are not arbitrarily changed.

It may appear at a first look as if the first three cases of Definition 1 are reducible to the last case. This is not correct: Any method may invoke another one with identical receiver (*reentrant call*, see also Example 13). In this example, the second case applies, while the fifth does not. Even the first case (constructor) has to be treated separately because a constructor might invoke another constructor. The poststate of this second invocation is visible to the object that is being initialized. Although there is no problem in defining formal semantics, this is a serious problem in practice. It would imply for virtually every constructor to break its invariant. If not declared explicitly, like in our example, Java enforces calls to `super()` to happen first on every constructor invocation. Even if the superclass constructor establishes its own invariant, it has no knowledge of fields in subclasses that need to be assigned to establish the invariant of the *subclass*. But according to Def. 1, its poststate (that is an inner state of the subclass constructor) is visible for the to initialize object—even though it is yet not fully initialized. In [Bruns, 2009], several alternative definitions of visible states are proposed that avoid this issue.

According to [Leavens et al., 2013], a state is visible for a type T (i.e., class or interface) if it occurs after static initialization of T and it is a visible state for some object of static type T . Leaving aside static initialization, this means that *every* state reached by the virtual machine is visible to every class and interface. This is because there is always an infinite number of instances for that the fifth case of Def. 1 applies. Therefore, a static invariant in type C must be respected in every state after C has been initialized.

5.2 Invariants

One of the most important and widely-used specification elements in object-orientation are type *invariants*, also called class or object invariants⁷, is a predicate over the object state that holds in all visible states of an object. These can be seen as conditions to constrain the state an instance can be in. In addition, since the poststate of a constructor is visible to the initialized object, any constructor has to ensure that the invariant is established.

Example 11. Listing 4 shows three possible invariants that can be added to interface `Student`. These specify that credits are never nonnegative; a student's status is always either Bachelor or Master, and nothing else; and if a student's status is Master, he or she has earned more than 180 credits.

Of course, instead of specifying invariants, one could also add these specifications to all pre- and postconditions explicitly. However, this means that if you add a method to a class, you have to remember to add these pre- and postconditions yourself. Moreover, invariants are also inherited by subclasses (and by implementations of interfaces).

⁷Not to be confused with loop invariants. These will be discussed in Sect. 10.2.

```

interface Student {

    public static final int bachelor = 0;
    public static final int master = 1;

    /*@ invariant getCredits() >= 0;
       @ invariant getStatus() == bachelor ||
       @           getStatus() == master;
       @ invariant getStatus() == master ==>
       @           getCredits() >= 180;
       @
       @ initially getCredits() == 0;
       @ initially getStatus() == bachelor;
       @
       @ constraint getCredits() >= \old(getCredits());
       @ constraint \old(getStatus()) == master ==>
       @           getStatus() == master;
       @ constraint \old(getName()) == getName();
    @*/

    public /*@ pure @*/ String getName();

    public /*@ pure @*/ int getStatus();

    public /*@ pure @*/ int getCredits();

    /*@ ensures getName().equals(n);
       public void setName(String n);

    /*@ requires c >= 0;
       @ ensures getCredits() == \old(getCredits()) + c;
    @*/
    public void addCredits(int c);

    /*@ requires getCredits() >= 180;
       @ requires getStatus() == bachelor;
       @ ensures getCredits() == \old(getCredits());
       @ ensures getStatus() == master;
    @*/
    public void changeStatus();

}

```

Listing 4: Interface Student with class level specifications

Thus any method that overrides a method from a superclass still has to respect the invariants. And any method that one adds in the subclass also has to respect the invariants from the superclass. This leads to a very nice separation of concerns.

An important point to realize is that invariants have to hold only in all visible object states, i.e., in all states in which a method is called or terminates. Thus, inside the method, the invariant may be temporarily broken.

Note that the kind of termination of a method does not matter. Regardless of terminating normally, exceptionally or erroneously, a method has to meet the invariant in every visible state.

Example 12. The following possible implementation of `addCredits` is correct, even though it breaks the invariant that a student can only be studying for a Master if he or she has earned more than 180 points inside the method: if `credits + c` is sufficiently high, the status is changed to Master. After this assignment the invariant does not hold, but because of the next assignment, the invariant is re-established before the method terminates.

```
/*@ requires c >= 0;
   @ ensures getCredits() == \old(getCredits()) + c;
   */
public void addCredits(int c) {
    if (credits + c >= 180) {status = master;} // invariant broken!
    credits = credits + c;
}
```

However, if a method calls another method on the same object, it has to ensure that the invariant holds before this callback. Why this is necessary, is best explained with an example.

Example 13. Consider the interface `CallBack` in Listing 5.

```
interface CallBack {

    //@ invariant getX() > 0;
    //@ invariant getY() > 0;

    /*@ pure @*/ public int getX();
    /*@ pure @*/ public int getY();

    //@ ensures getX() == x;
    public void setX(int x);

    //@ ensures getY() == y;
    public void setY(int y);

    //@ ensures \result == getX() % getY();
    public int remainder();

    public int longComputation();
}
```

Listing 5: Interface `CallBack`

Typically, correctness of the method `remainder` crucially depends on the value of `getY` being greater than 0. Suppose we have an implementation of the `CallBack` interface, where the method `longComputation` is sketched as follows.

```
public int longComputation(){
    ...
    if (getY() ...) {
        setY(0); // invariant broken
    }
    ...
    int r = remainder(); // callback
    ...
    setY(r + 1) // invariant re-established
    ...
    return ...
}
```

Naively, one could think that the fact that the invariant about `getY` is broken inside this method, is harmless, because the invariant is re-established by the `setY(r + 1)` statement. However, the call to the method `remainder` is a callback, and the invariant should hold at this point. In fact, correct functioning of this method call depends on the invariant holding. The invariant implicitly is part of `remainder`'s precondition. If the invariant does not hold at the point of the callback, this means that `remainder` is called outside its precondition, and no assumption can be made about its result as well.

Although invariants are always specified within a class or interface, their effective scope is global. E.g., a method m in a class C is obliged to respect invariants of class D . There is a way to avoid the requirement that the invariant has to hold upon callback: this is by specifying that a method is a **helper** method. Such methods cannot depend on the invariant to hold, and they do not guarantee that the invariant will hold afterwards. Typically, only private methods should be specified as helper methods, because one does not want that any other object can directly invoke a helper method.

Where do invariants come from? Sometimes they are imposed by some kind of 'reality' that the code is a model of. The interface `Student` in Listing 4 is such an example. Students can only have a positive number of credits, they must be either Master or Bachelor students, and so forth. Another common source of invariants is efficiency. Efficient computations often require to organise the data in a specific way. One way is introducing redundancy, like for instance in an index of a book, mapping words to pages where they occur. Such an index is redundant (we can always search through the whole book to find the occurrences of a word), but it enables efficient lookup. On the downside, redundancy opens up for inconsistencies. The countermeasure are invariants indeed, formalising the consistency conditions (like each word in an index appearing in the text as well, at the page given by the index). Other ways to increase efficiency limit the organization of data to comply to certain restrictions. A classic example of that is sortedness, that allows for quicker look-up. To extend, for instance, the example of `LimitedIntergerSet` (Listing 3) by sortedness, we add the invariant

```
/*@ public invariant (\forall int i;
    @           0 < i && i < size;
    @           arr[i-1] <= arr[i]) ;
@*/
```

to that class. With that, the implementer of each method can *rely* on sortedness, and the implementor of each impure method has to *guarantee* sortedness.

Defining a precise semantics for invariants is still an active area of research, see e.g. [Poetzsch-Heffter, 1997, Leino and Müller, 2004, Barnett et al., 2004, Müller et al., 2006, Bruns, 2009]. A complication is that, although invariants are declared in a particular class, not only instances of that class have to respect it, but all objects in the system. An alternative approach, that is used in the Spec# framework, is to explicitly add specification statements `unpack` and `pack` for invariants. An invariant may only be broken if it has been explicitly unpacked. When the invariant is reestablished, it has to be explicitly be packed again, and this only succeeds if the invariant indeed holds at this point. Every method can then specify explicitly whether it assumes invariants to hold, i.e., to be packed, or not. This approach is sometimes referred to as the Boogie methodology [Barnett et al., 2005a].

Similar to the Boogie methodology, in the KeY system, invariants are not implicitly added to specifications, but it is left to the specification to include specific invariants. This specification may be more verbose, but it is clear from the given specification that invariants are assumed or established. The invariant for an object `o` can be referred to through `\invariant_for(o)`. This allows that both visible and observed state semantics of invariants can be simulated. Unlike in Boogie, explicit packing/unpacking instructions in the code are not necessary. Instead, the specifier has to specify a set of locations on that the invariant depends at most (`accessible` clause). Usually, methods rely at least on the invariant of the current receiver. For convenience, this invariant is implicitly included for non `helper` methods.

Finally, it is important to realize that the notion of object invariant that we discussed here only makes sense in a sequential setting. In a multithreaded setting, there always may be another thread accessing the object simultaneously, and one cannot talk about visible state semantics anymore. Instead, in a multithreaded setting, one sometimes specifies strong invariants that may never be broken. See, e.g., [Zaharieva-Stojanovski and Huisman, 2014] for a modular specification and verification technique for class invariants in a concurrent setting.

5.3 Initially Clauses

Sometimes, one explicitly wishes to specify the conditions that are satisfied by an object upon creation. Each (non-helper) constructor⁸ of the object has to establish the predicate specified by the initially clause. Another advantage of initially clauses is that they are inherited; that means that also constructors of subclasses have to fulfil them. Constructors in Java itself are not inherited. As a consequence, a constructor can rely on the guarantees provided by a called super constructor but does not have to maintain them.

Example 14. Listing 4 shows some possible initially clauses for the `Student` interface.

Again, it would be possible to specify this property as a postcondition of all constructors, instead of as a single initially clause. But in this way, any additional constructor has to respect the initially clause, and we ensure that also subclasses respect it.

5.4 History Constraints

Invariants as we discussed above define a predicate that every (visible) state of the object should respect. However, sometimes one also wishes to specify how an object may

⁸Again, typically only private constructors would be annotated as a helper constructor.

evolve over time, i.e., the relationship that exists between the pre-state and the post-state of a method call. This could be seen as a sort of general postcondition that has to be respected by every method, however the definition is actually more fine grained than that. For this, *history constraints* (usually *constraints* for short) have been introduced [Liskov and Wing, 1993]. Constraints can be seen as implicit postconditions, but just as invariants and initially clauses, they have the advantage that they are inherited, and immediately are required to hold for any additional methods.

History constraints are in a way similar to invariants as they constrain the state that an object may be in. But while invariants must hold for every visible state, history constraints describe the relation of two consecutive visible states in a program execution. Constraints may rely on syntactical features that are used to measure changes between states such as the `\old` operator as well as frame expressions. Similar to invariants, there may be several constraint definitions and non-private constraints are being inherited. Assigning suitable semantics to history constraints is non-trivial; a possibility would be to see them as special two-state model methods. This is not yet implemented in KeY at the time of writing (KeY version 2.2).

Example 15. Listing 4 defines several constraints for the `Student` interface. The first constraint specifies that the amount of credits can never decrease. The second constraint specifies that if a student has obtained the Master status, he will remain a Master student, and cannot be downgraded to a Bachelor student again. Finally, the third constraint specifies that a student's name can never change.

When specifying constraints, it is important that they should denote a nonstrict relation, i.e., it should be possible to respect a constraint without actually changing the state. This is sensible in practice, since it is nontrivial for an observer to deduce that states are visible. In particular, any pure method should be able to respect the constraint. Therefore, one should not specify the following strict constraint:

```
constraint \old(getCredits()) < getCredits();
```

as it is impossible to respect this constraint with a pure method. Typically, constraints will also be transitive, so that when you consecutively call two methods from the same object, you also know the relationship that holds between the pre-state of the first method, and the post-state of the second method.

```
//@ constraint \old(getCredits()) <= getCredits();

/*@ requires c >= 0;
   @ ensures getCredits() == \old(getCredits()) + c;
   @*/
// pre-state
public void addCredits(int c) {
    credits = credits + c;
    if (credits >= 180) {
        // call-state changeStatus
        changeStatus();
        // return-state changeStatus
    }
} // post-state
```

Listing 6: Checking history constraints

Example 16. Consider the possible implementation of `addCredits` in Listing 6. To show that the constraint is respected, it has to hold for the following visible state pairs:

- (pre-state, call-state `changeStatus`)
- (call-state `changeStatus`, return-state `changeStatus`)
- (return-state `changeStatus`, post-state)

Notice that if the constraint is transitive, the relationship also holds for the pair of prestate and poststate, which is indeed what we want.

Again, in a multithreaded setting, the semantics of constraints would become less clear. Because any interleaving is possible, all intermediate states must be assumed to be visible states, because of the possible thread interleavings. However, a constraint such as that `getName` returns a constant value could still be meaningful also in a multithreaded setting (except that the number of possible visible state pairs that have to be considered might grow exponentially). Therefore, in a concurrent setting one could imagine a notion of *strong* constraints, i.e., a relationship that has to hold for any pair of consecutive states.

5.5 Static Class Specifications

For all class level specification constructs, static variants exist. For example, an invariant might restrict the value of a static variable, or a constraint might restrict the evolution of a static variable. All static specifications have to be preceded by the keyword `static`. Since instance methods might change static variables, static invariants and constraints have to be respected by instance methods. In contrast, invariants and constraints that only restrict the instance variables of a method cannot be invalidated by a static method—and thus this does not have to be checked explicitly.

5.6 Inheritance of Specifications

Design by Contract allows to impose the concept of *behavioral subtyping* [Liskov, 1988], that is usually defined through the *Liskov substitution principle*, or Liskov principle for short [Liskov and Wing, 1994].⁹ A type T' is a behavioral subtype of type T if every observable behavior of T is also observable on T' . In an object oriented program, this means that any subclass may be used wherever a superclass is expected. Behavioural subtyping expresses the idea that a subclass thus should behave as the superclass (at least, when it is used in a superclass context). Subclasses in Java do not always define behavioral subtypes. They can be used simply for the purpose of code reuse.

But what exactly is the *expected* behavior of a linked list? Surely, given concrete implementations in Java, there is no indeterminism that can be refined. This means there cannot be strict behavioral subtypes regarding all behaviors; the substitution principle as originally stated by Liskov [1988] is too strong in practice (cf. [Leavens, 1988]). Instead, we focus on the client perspective again and define behavior subtypes regarding contracts (and invariants). This means that a class C' is a behavioral subtype of a super class C , if for every method m implemented in both C and C' (i.e., the implementation in C' is overriding), every specification case for $C :: m$ is also a specification case for $C' :: m$, and that the contract of $C :: m$ is refined by the contract of $C' :: m$. A

⁹It is named the “constraint rule” there.

full formalization of this definition of behavioral subtyping can be found in [Leavens and Naumann, 2006].

To ensure that a subclass indeed defines a behavioral subtype, specification inheritance can be used [Dhara and Leavens, 1995, Leavens and Dhara, 2000]: In JML, every (nonprivate) method in the subclass inherits the overridden method's specification cases defined in the superclass. And in addition, all invariants of the superclass are inherited by the subclass. Notice that this same approach applies for interfaces and implementing classes. An interface can be specified with its desired behavior. Every class that implements this interface should be a behavioural subtype of the interface, i.e., it should satisfy all the specifications of the interface. Concretely, this means the following:

- every method that overrides a method from a superclass, or implements from an interface, has to respect the method specification from the superclass;
- every class that implements an interface has to respect the specifications of the interface; and
- every class that extends another class has to respect the specifications of that class.

Still, it is possible to refine specifications in subclasses (or implementing classes), in addition to what is inherited. Any additional specification of an inherited method (whether or not the implementation is overridden) is added to the inherited specifications from the superclass, using the **also** keyword.

```
/*@ also
   @
   @ <subclass-specific-spec-cases>
   @*/
public void method () { ...
```

Note that the JML comment starts with **also**, not preceded by anything. This is because the inherited specification cases are still there, even if implicit, to be extended here by whatever is written after the **also**.

Invariants are also fully, and implicitly, inherited. Extending the set of inherited invariants by additional invariants specific for a subclass is easy, by simply writing them in the subclass, using the normal syntax for invariants. The same applies also to initially clauses and constraints.

The idea of behavioral subtypes is crucial for the correctness of object-oriented programs. We can specify the behavior of a class in an abstract way. For example, in class *Average* in Listing 7, we have an array of *Student* instances; the concrete instances that are stored in the array may have different implementations, but we know that they all implement the methods specified in the interface *Student* in Listing 1. This means that we can rely on the specification case of *Student#getCredits()* in Line ?? of *Average#averageCredits()*.

Respecting inherited specifications is a good practice, but it does not guarantee behavioral subtyping per se. JML allows to make program elements more visible in the specification than they are in the implementation (through the **spec_public** modifier, see Sect. 4.1). In this way, specifications may expose implementation details. While it is also a good practice to declare those specifications private, in many cases, this would disable us from giving *any* meaningful specification. A solution to this dilemma is abstraction, that will be covered in Sect. 8.1 below.

6 Nonnull versus Nullable Object References

In Java, the set of values of reference type include the null reference. (Note that the same is true for the values of array type, because each array type is also a subtype of `Object`.) But even if the type system always allows `null`, the specifier may want to exclude the null reference in many cases. Whether or not null is allowed can be expressed by means of simple (in)equations, like, for instance, `o != null`, in pre/post-conditions or invariants. However, this issue is of so dominant importance that JML offers two special modifiers just for that, `non_null` and `nullable`. Class members (i.e., fields), method parameters, and method return values can be declared as `non_null` (meaning null is forbidden), or `nullable` (in which case null is allowed, but not enforced).

Here are some examples for forbidding null values.

```
private /*@ non_null @*/ String name;
```

adds the *implicit* invariant `invariant name != null;` to the class at hand.

```
public void setName(/*@ non_null @*/ String n) {...
```

adds the *implicit* precondition `requires n != null;` to each specification case of `setName`.

```
public /*@ non_null @*/ String getName() {...
```

adds the *implicit* postcondition `ensures \result != null;` to each specification case of `getName`.

The reader can imagine that `non_null` modifiers can easily bloat the specification. Therefore, JML has built-in `non_null` as the *default* for all fields, method parameters, and return types, such that the `non_null` in the above examples is actually redundant. By only writing

```
private String name;
```

```
public void setName(String n) {...
```

```
public String getName() {...
```

without any explicit `non_null`, we get *exactly* the same implicit invariants, preconditions, and postconditions as mentioned above.

But how can we allow null anyway? We can avoid the nonnull default by the aforementioned modifier `nullable`. In the above examples, we could allow null (and thereby avoid the implicit conditions), by writing

```
private /*@ nullable @*/ String name;
```

```
public void setName(/*@ nullable @*/ String n) {...
```

```
public /*@ nullable @*/ String getName() {...
```

Notice that the nonnull by default also can have some unwanted effects, as illustrated by the following example.

Example 17. Consider the following declaration of a `LinkedList`.

```
public class LinkedList {
    private Object elem;
    private LinkedList next;
```

```

    ....
}

```

Because of the nonnull by default behavior of JML, this means that all elements in the list are nonnull. Thus the list must be cyclic, or infinite.¹⁰ This is usually not the intended behavior, and thus the `next` reference should be explicitly annotated as `nullable`.

```

public class LinkedList {
    private Object elem;
    private /*@ nullable @*/ LinkedList next;
    ....
}

```

In short, it is important to remember that for all class fields, method parameters, and method results, the `null` reference is *forbidden* wherever we do not state otherwise with the JML modifier `nullable`.

In the context of allowing vs. forbidding the null reference, handling of arrays deserves special mentioning. The additional question here is whether, or not, the prohibition of null holds for the *elements* of the array. Without loss of generality, we consider the following array typed field declaration: `String[] arr;`. Because of nonnull being the default, this is equivalent to writing `/*@ non_null @*/String[] arr;`. Now, in both cases, the prohibition of null references extends, in JML, to the elements of the array! In other words, both the above forms have the same meaning as if the following invariants were added:

```

/*@ invariant arr != null;
/*@ invariant (\forall int i; i >= 0 && i < arr.length; arr[i] != null);

```

Again, no such invariant is needed for disallowing null; writing `String[] arr;` is enough. We can, however, allow null for *both*, the whole array and its elements (at first), by writing `/*@ nullable @*/String[] arr;`. To that, we can add further restrictions. For instance, if only the elements may be null, but not the whole array, we can write:

```

/*@ invariant arr != null;
/*@ nullable @*/ String[] arr;

```

7 Exceptional Behavior

So far, we have only considered normal termination of methods. But in some cases, exceptions cannot be avoided. Therefore JML also allows one to specify explicitly under what conditions an exception may occur.

The `signals` and `signals_only` clauses are introduced to specify *exceptional postconditions*. In addition, one can give an `exceptional_behavior` method specification that expresses that a method must terminate with an exception. Exceptional postconditions have the form `signals (E e) P`, where E is a subtype of `Throwable`, and the following meaning: if the method terminates because of an exception that is an instance of type E , then the predicate P has to hold. The variable name `e` can be used to refer to the exception in the predicate. The `signals_only` clause is optional in a

¹⁰A linked data structure having infinite length is indeed a contradiction. At runtime, there are only finitely many created objects on the heap.

method specification. Its syntax is `signals_only E1, E2, ..., En`, meaning that if the method terminates because of an exception, the dynamic type of the exception has to be a subclass of `E1, E2, ..., or En`. If `signals_only` is left out, only *unchecked exceptions*, i.e., instances of `Error` and `RuntimeException`, and the exception types declared in the method's `throws` clause are permitted.

```

1 class Average {
2
3     /*@ spec_public @*/ private Student[] s1;
4
5     /*@ signals_only ArithmeticException;
6         @ signals (ArithmeticException e) s1.length == 0;
7         @*/
8     public int averageCredits() {
9         int sum = 0;
10        for (int i = 0; i < s1.length; i++) {
11            sum = sum + s1[i].getCredits();
12        };
13        return sum/s1.length;
14    }
15 }

```

Listing 7: Class Average

Example 18. Consider for example class `Average` in Listing 7. The specification of method `averageCredits` states that the method may only terminate normally, or with an `ArithmeticException`—and thus, it will not throw an `ArrayIndexOutOfBoundsException`. Moreover, if an `ArithmeticException` occurs, then in this exceptional state the length of `s1` is 0.

Notice that it is incorrect to use an `ensures` clause, instead of a `signals` clause: an `ensures` clause specifies a normal postcondition, that only holds upon normal termination of the method.

Above, in Sect. 2 we discussed `normal_behavior` specifications. Implicitly, these state that the method has to terminate normally. Similarly, JML also has an `exceptional_behavior` method specification. This specifies that the method has to terminate, because of an exception. As mentioned above, a `behavior` specification only enforces that a method terminates, but it does not exclude exceptional termination. Thus a `behavior` specification may well contain a `signals` or `signals_only` clause, whereas a normal behavior specification may not contain these, and an exceptional behavior specification may not contain an `ensures` clause.

As mentioned above, a single method can be specified with several method specifications, joined with `also`. Exceptional behavior specifications are typically used in this case.

Example 19. Consider the more detailed specification for `averageCredits` in Listing 8.

This states that if `s1.length > 0`, i.e., there are students in the list, then the method terminates and the result is the average value of the credits obtained by these students. If `s1.length == 0` then the method will terminate exceptionally, with a `ArithmeticException`.

In this example, the two preconditions together cover the complete state space for the value of `s1.length`. If `s1.length` could be less than 0, the method's behavior would not be specified.

```

class Average2 {

    /*@ spec_public @*/ private Student[] sl;

    /*@ normal_behavior
    @ requires sl.length > 0;
    @ ensures \result ==
    @         (\sum int i; 0 <= i && i < sl.length;
    @             sl[i].getCredits())/sl.length;
    @ also
    @ exceptional_behavior
    @ requires sl.length == 0;
    @ signals_only ArithmeticException;
    @ signals (ArithmeticException e) true;
    @*/
    public int averageCredits() {
        int sum = 0;
        for (int i = 0; i < sl.length; i++) {
            sum = sum + sl[i].getCredits();
        };
        return sum/sl.length;
    }
}

```

Listing 8: Class Average

Finally, it is important to realize that invariants and constraints also must hold when a method terminates exceptionally. This might seem strange at first: something goes wrong during the execution, so why would it be necessary that the object stays in a good state. But in many cases, the object can recover from the exception, and normal execution can be resumed. But this means that it is necessary that also when an exception occurs, the object stays in a ‘welldefined’ state, i.e., a state in which the invariants hold, and that evolves according to the constraints.

A Note on False In exceptional behavior specifications, one often sees specifications like:

```
signals (Throwable e) false;
```

This is a way to state that an exception should not occur: if the exception occurs, the property `false` should hold. And as this is this never the case, the exception may not occur.

Similarly, if one specifies a postcondition `ensures false`; this states that a method should not terminate normally. Thus a method specification:

```
ensures false;
signals (Throwable e) false;
diverges true;
```

implicitly says that a method should never terminate (either normally, or exceptionally).

Finally, a method can also be specified with a precondition `requires false`; This means that the method may not be invoked, as no caller can fulfill the precondition of the method.

8 Specification Only Class Members

The previous sections shows how the behavior of code members is specified in JML. But sometimes it is easier or even required to introduce new members only for specification.

Model fields, as discussed in Sect. 8.1, allow one to provide abstraction from the concrete program state. For each abstract state, a relationship to the concrete program state can be defined. In addition to model fields, sometimes it is also useful to define model methods, i.e., methods that are used in specifications only.

This section also introduces ghost variables (Sect. 8.2). These can be used to extend the state space with specification only information. They do not provide abstraction, but can record extra information. The use of model and ghost fields is often confused, and therefore Sect. 8.3 compares both approaches, and highlights their differences.

8.1 Model Fields and Model Methods

An important feature of specifications is that they provide abstraction over the concrete implementations. *Model fields* serve as an abstraction feature in a familiar guise. They are declared like regular fields, but within JML specifications and with the modifier keyword `model`. Model fields can be read from like regular fields, but there are no assignments to them since they do not have a state of their own. Instead, to make sure that the concrete implementation corresponds to the abstract specification, a link between the two has to be made. For this purpose, the *represents clause* defines how the value of the abstract variable is defined in terms of the values of the concrete entities. In the so called functional form, the represents clause, that is a class member, appears similar to an assignment, as can be seen in the following example taken from [Breunese et al., 2005].

Example 20. Class `Decimal` implements decimal variables using an `intPart` and `decPart` variable, but the specification is given in terms of a single model field that represents the value of the composed decimal number.

```
class Decimal {
    public static final short PRECISION = (short) 1000;
    /*@ spec_public @*/ private short intPart = (short) 0;
    /*@ spec_public @*/ private short decPart = (short) 0;

    /*@ model int value;
    /*@ represents value = intPart * PRECISION + decPart;
}
```

Sometimes, a represents clause cannot be defined directly as a translation into concrete variables; sometimes a (nonfunctional) relation between the abstract and the concrete state can be expressed, sometimes only a dependency relation. JML provides a way to define non-functional represents clauses. Instead of the assignment operator, they consist of the keyword `\such_that` followed by a boolean expression. It means that the model fields points to some value such that this condition is satisfied.

Example 21. Consider class `MatrixImplem` in Listing 9. It implements a matrix as a single array (on some platforms, as `JavaCard`, only one-dimensional arrays are allowed). A model variable `matrix` is declared, that specifies the abstract representation of the matrix.

Unfortunately, no functional `represents` clause can be specified for this. Instead, the `such_that` keyword is used to define a relational `represents` clause, that enables to write the specifications of the matrix methods in terms of the abstract matrix variable.

Model fields are useful in many cases. Typical examples are specifications of interfaces. The behavior of an interface is specified in terms of model variables, and the classes implementing the interface define `represents` clauses for these model variables, relating them to their own concrete implementation. Because of the flexible connection between concrete and abstract state using the `represents` clause, this does not impose any restriction on the internal state of a class implementing the interface. Note that in interfaces, model field declarations are `static` by default, nonstatic model field declarations must use the modifier `instance`.

Example 22. Listing 10 gives an alternative specification for interface `Student` using model fields. It shows the specification for an implementing class `CCStudent`. Note that it does not declare the model variables, but only defines the `represents` clause.

Sometimes, to complete a specification, one needs a *method* that only is intended for specification. To support this, JML provides *model methods*. A model method is defined as part of the specification. It can be implemented, but it may also be abstract. And the behavior of a model method is typically defined in terms of its pre- and postconditions again. Typical usages for model methods are:

- if the specification needs a method that is not related to the code, for example to sum all the elements in an array;
- if the specification needs a method that cannot be implemented easily, but that can be specified without any problem.

8.2 Ghost Variables

Sometimes the information needed in specifications is not provided by the source code itself. Typical examples are specifications that express something about the control flow, e.g., how often or in which order methods are called, or about the used resources, e.g., to limit the number of objects. This additional knowledge can be modelled with ghost variables. These are specification-only variables that can be updated in the specification, via a special set annotation, placed inside a method body.

A ghost variable in JML can be defined as a class/instance member or as a local variable. In both cases, it is declared as a normal Java variable, but inside a JML comment preceded by the keyword `ghost`. It is important to know that the used type might be a specification-only type such as `\bigint` (see Sect. 9). The initial value of a ghost variable can be directly assigned during its declaration. Its value can be updated during method execution by a `set` statement. This is a JML comment placed like a normal statement within a method body, consisting of a keyword `set` followed by an assignment. The left side of the assignment has to be a ghost variable and the right side can be any side effect free JML expression.

Example 23. Consider class `Tries` in Listing 11. This class supports operations that might occasionally fail, and then should be retried, instead of returning an error. A typical example of such an operation would be modifying a file, that might be temporarily locked by the operating system. Method `tryMaximalNTimes` does a maximum of `n` tries of the `doTry` operation, before returning a failure. To describe the behavior of this method, ghost variables are very convenient, because the number of successful and failed tries is only part of the specification, and not of the program itself.

```

public class MatrixImplem {

    /*@ public model int[][] matrix;
    private int x;
    private int y;
    private int[] matrix_implem;
    /*@ represents matrix \such_that
       @ (\forall int i; i >= 0 && i < x;
       @   (\forall int j; j >= 0 && j < y;
       @     matrix[i][j] == matrix_implem[x * j + i]));
    @*/

    /*@ ensures
       @ (\forall int i; i >= 0 && i < x;
       @   (\forall int j; j >= 0 && j < y;
       @     matrix[i][j] == 0));
    @*/
    public MatrixImplem(int x, int y) {
        this.x = x;
        this.y = y;
        matrix_implem = new int [x * y];
    }

    /*@ ensures \result == matrix[i][j];
    public /*@ pure @*/ int get (int i, int j) {
        return matrix_implem[x * j + i];
    }

    /*@ ensures \result >= 0
       @   ==> matrix[\result][coordY(elem)] == elem;
    @*/
    public /*@ pure @*/ int coordX (int elem) {
        for (int i = 0; i < matrix_implem.length; i++)
            if (matrix_implem[i] == elem)
                return i % x;
        return -1;
    }

    /*@ ensures \result >= 0
       @   ==> matrix[coordX(elem)][\result] == elem;
    @*/
    public /*@ pure @*/ int coordY (int elem) {
        for (int i = 0; i < matrix_implem.length; i++)
            if (matrix_implem[i] == elem)
                return i / x;
        return -1;
    }
}

```

Listing 9: Represents clause using \such_that

```

public interface Student {

    /*@ public instance model int status;
       @ public instance model int credits;
       @ represents status = (credits < 180 ? bachelor : master);
       @*/

    /*@ instance invariant status == bachelor || status == master;
       @ instance invariant credits >= 0;
       @*/

    public static final int bachelor = 0;
    public static final int master = 1;

    /*@ pure @*/ public String getName();

    /*@ ensures \result == status;
       @ pure @*/ public int getStatus();

    /*@ ensures \result == credits;
       @ pure @*/ public int getCredits();

    /*@ ensures getName().equals(n);
       public void setName(String n);

    /*@ requires c >= 0;
       @ ensures credits == \old(credits) + c;
       @*/
    public void addCredits(int c);

    /*@ requires credits >= 180;
       @ requires status == bachelor;
       @ ensures credits == \old(credits);
       @ ensures status == master;
       @*/
    public void changeStatus();
}

class CCStudent implements Student {

    private int[] creditList;

    /*@ represents credits =
       @ (\sum int i; 0 <= i && i < creditList.length; creditList[i]);
       @*/

    // rest of class continued...
}

```

Listing 10: Interface Student with model fields and an implementation.

Two ghost variables are declared: `successful` and `failed`. The specification of `tryMaximalNTimes` ensures that if it returns `true` method `doTry` was exactly once successful and failed at most $n - 1$ times and if it returns `false` otherwise method `doTry` was never successful and failed exactly n times. The ghost variables are updated in each loop iteration.

```

1 public abstract class Tries {
2     //@ ghost private int successfulTries;
3     //@ ghost private int unsuccessfulTries;
4
5     public abstract boolean doTry();
6
7     /*@ ensures \result ==> successfulTries == \old(successfulTries) + 1 &&
8         @                unsuccessfulTries >= \old(unsuccessfulTries) &&
9         @                unsuccessfulTries <= \old(unsuccessfulTries) + 2;
10        @ ensures !\result ==> successfulTries == \old(successfulTries) &&
11        @                unsuccessfulTries == \old(unsuccessfulTries) + 3;
12        @*/
13    public boolean tryMaximalThreeTimes() {
14        int i = 0;
15        boolean done = false;
16        while (i < 3 && !done) {
17            if (doTry()) {
18                //@ set successfulTries = successfulTries + 1;
19                done = true;
20            }
21            else {
22                //@ set unsuccessfulTries = unsuccessfulTries + 1;
23                done = false;
24            }
25            i++;
26        }
27        return done;
28    }
29 }

```

Listing 11: Usage of ghost variables

8.3 Ghost Variables vs Model Fields

It is important to understand the difference between model and ghost variables. Both are variables that are used for specification purposes only, and they do not occur during the execution of the program.

However, model variables provide an abstract representation of the state. If the underlying state changes, implicitly the model variable also changes. Often it is possible to define this relationship explicitly as a translation, but sometimes it can only be given in a nonconstructive manner (or even as a dependency relation).

In contrast, ghost variables extend the state. They provide some additional information that cannot be directly related to the object state. Ghost variables are often used to keep track of the events that have happened on an object, e.g., which methods have been called, how often have these methods been called etc. There also exists

work where ghost variables have been used to keep track of the resources used by the program: every time a new object is created, there is an associated set-annotation that increases a resource counter, modelled as a ghost variable [Barthe et al., 2005]. In this way, the specification can state something about the number of objects that have been created by the program. This information allows then to define a resource analysis over the application.

9 Integer Semantics

Since JML incorporates Java expressions, specifications also adhere to the semantics of the Java numerical data types. This means in particular that always special care has to be taken regarding overflows in integer operations¹¹. Undoubtedly, dealing with finite numerical data types is a very common source of programming errors. The most infamous example from the real world is the maiden flight of Ariane 5, where conversion of 64-bit floating point data to 16-bit integers finally caused the spacecraft to be destroyed just seconds after lift off [Nuseibeh, 1997]. It is thus desirable to detect such errors and to not repeat them in the specification. We will show how to avoid this problem through the use of JML’s `\bigint` data type, that represents the mathematical integers. This section does not discuss semantics of integral data types in general; that can be found in [Beckert et al., 2007, Chap. 12] or (more elaborate) in [Schlager, 2002].

Example 24. Regard the short method `mult()` below; it returns `a*b`, but this is not multiplication in the mathematical sense, since an overflow may occur.

```
public int mult (int a, int b) { return a*b; }
```

The naive specification `ensures \result == a*b`; would be trivially true since it uses the same overflow semantics.

This example shows a feature of Java that may be a large source of confusion. Integer operators in Java are often misunderstood to equal their mathematical counterparts, see, e.g., the survey by Chalin [2003]. But the actual functionality¹² represented by, e.g., `a*b` (where both are `int` expressions) is $((a + 2^{31}) \cdot (b + 2^{31}) \bmod 2^{32}) - 2^{31}$. In addition, these operators are overloaded – the `*` operator has different semantics if one operand is of type `long` (64-bit integers). This means that, in many situations, naive specifications are just incorrect due to the presence of overflows. For instance, in Listing 4, the invariant that credits are non-negative can be broken by method `addCredits()`, that does not check for overflows.

Example 25. To display even more obscure characteristics of overflow semantics, the following boolean JML expression is trivially true. We leave it to the reader to find out with which element the quantifier would be instantiated.

```
(\exists int x; x-1 > x
  && (\forall int y; x <= y))
  && x == -x
  && x != 0 && x * 2 == 0);
```

¹¹Similar issues arise with rounding in floating point operations, which however will not be covered here.

¹²More mathematically speaking, the `int` data type with operators `+` and `*` forms a finite Abelian ring that is isomorphic to $\mathbb{Z}/\mathbb{Z}_{2^{32}}$. This means that addition and multiplication are commutative, associative, and distributive; but there are zero-dividers—as shown in Example 25.

Besides Java’s bounded integer types (also known as *bit vector* types), JML offers the specification only primitive type `\bigint`, that represents the mathematical integers \mathbb{Z} . ‘Specification only’ means that, besides variables bound by a quantifier, only ghost variables and ghost and model fields can be declared with type `\bigint`. The Java standard library also provides a type called `BigInteger`, that represents arbitrary precision integers. While `\bigint` is a primitive type with an infinite number of elements, `BigInteger` is just a regular Java object type. This means, in particular, that instances of `BigInteger` must be created through constructors and that quantification makes little sense since it only ranges over the (finitely many) created instances. It is therefore inadequate for specification purposes.

Let us come back to Example 24. How can we specify that there is no overflow? In Java, all arithmetic operations are *unchecked*, i.e., an overflow is not indicated in any way, e.g., by exceptions. A precondition like `a*b <= Integer.MAX_VALUE` is trivially true. Instead, we can apply numerical conversion to `\bigint` to expressions of type `int`. Note that this kind of conversion, a widening, has no effect on the *values* of `a` and `b`, but on the semantics of the `*` operator. Under the preconditions that the (mathematical) product of `a` and `b` is within the bounds of `int`, we can ensure that the result is indeed the mathematical product:

```
//@ requires Integer.MIN_VALUE <= (\bigint) a * (\bigint) b;
//@ requires Integer.MAX_VALUE >= (\bigint) a * (\bigint) b;
//@ ensures \result == (\bigint) a * (\bigint) b;
public int mult (int a, int b) { return a*b; }
```

Because this specification is tedious to write and even more horrible to read, classes and methods can be annotated in JML with *math modifiers* [Chalin, 2004]. The default integer semantics in specifications can be changed by declaring the method `spec_bigint_math`, that achieves the above while saving to write down casts explicitly.

```
//@ requires Integer.MIN_VALUE <= a * b;
//@ requires Integer.MAX_VALUE >= a * b;
//@ ensures \result == a * b;
public /*@ spec_bigint_math @*/ int mult (int a, int b) { return a*b; }
```

An even simpler way to express the absence of overflows is to change the semantics of the Java implementation through the `code_safe_math` modifier. It causes the program to be interpreted as if operations were checked, leading to an exception in case of overflow. The only thing left to show is that there are no exceptions:

```
//@ signals_only \nothing;
public /*@ code_safe_math @*/ int mult (int a, int b) { return a*b; }
```

There are six math modifiers in total, declaring integer expressions in specifications or code to be interpreted as either Java integers with default operations, mathematical integers, or Java integers with checked operations. While these modifiers are currently not directly supported, the KeY prover offers to select different integer semantics.

10 Specification for Verification

The previously discussed specification constructs are essential to the Design by Contract philosophy and relevant to all analysis techniques. However, for static verification of Java programs it is typically required to provide some additional information, like the locations a method might access (Sect. 10.1); guidance for the verification tool in

the presence of loops via loop invariants (Sect. 10.2); or in general via assert statements (Sect. 10.3).

10.1 Framing

An important aspect of verification is modularity. Each method is verified in isolation, and any method call inside a body is abstracted by its method specification.

To achieve this, it is not enough to specify what a method does; it is also required to specify what a method does not do. This is known as the *frame problem* [Borgida et al., 1995, Müller et al., 2002]. Basically, for modular verification one needs to know what is the *frame* of a method, i.e., what are the variables that may be changed by the method, and what is the *anti frame*, i.e., which variables may not be changed by the method.

To specify this, JML uses the *assignable clause*. This provides a set of variable locations that may be modified by a method (thus, it may be an over-approximation of the actual set of locations that is modified by the method). Location sets can be given through comma separated lists of single variables or one of the special keywords `\nothing` (only locations of newly allocated objects may be changed, corresponds to weak purity, see Sect. 4.5), `\everything` (any location may be changed), `this.*` (all locations provided by the current object), and `array[*]` or `array[i..j]` (all elements in the array or between indices *i* and *j*). While assignable clauses are attached to single specification cases, pure methods are defined to have an empty frame under *any* precondition. The JML extension used in KeY, called JML*, provides additional constructs to specify frames, offering more flexibility. Most important, in JML*, the keyword `\strictly_nothing` denotes strictly the empty set of locations; strictly pure methods are annotated with `strictly_pure`, see Sect. 4.5.

Example 26. Listing 12 contains the specification of Listing 1, but with assignable clauses added.

Method `addCredits` increases the achieved credits, which means that it may have to update the master flag to maintain the invariant. Therefore, the assignable clause of this method lists the instance variables `credits` and `master`. Even though the variables are not modified directly by the method, it is required to list them in the assignable clause, because they may be modified during the method execution.

Methods `updateCredits`, `changeToMaster` and `setName` modify only one instance variable, that is listed in the assignable clause of their method specifications.

Finally, method `getName` is specified as a pure method, that automatically implies that the assignable clause is `\nothing` by default.

Of course it would be possible to add the information in the assignable clause to the postcondition, explicitly specifying that the variables not mentioned in the assignable clause are not changed. But this is not a satisfactory solution: a class might have many variables and only a few are typically changed by a method. Moreover, when a new variable is added, for every method that does not change it, an additional postcondition about this variable not being changed would have to be added. As one can imagine, this is error-prone, and leads to overly verbose specifications.

10.2 Loop Invariants

A verification tool typically needs some guidance in presence of loops to verify that a method implementation complies to its specification. This is due to the general impos-

```

1 public class Student {
2     private String name;
3
4     /*@ invariant credits >= 0;
5        */
6     private int credits;
7
8     /*@ invariant credits < 180 ==> !master &&
9        @           credits >= 180 ==> master;
10        */
11    private boolean master;
12
13    /*@ requires c >= 0;
14       @ ensures credits == \old(credits) + c;
15       @ assignable credits, master;
16       */
17    public void addCredits(int c) {
18        updateCredits(c);
19        if (credits >= 180) {
20            changeToMaster();
21        }
22    }
23
24    /*@ requires c >= 0;
25       @ ensures credits == \old(credits) + c;
26       @ assignable credits;
27       */
28    private void updateCredits(int c) {
29        credits += c;
30    }
31
32    /*@ requires credits >= 180;
33       @ ensures master;
34       @ assignable master;
35       */
36    private void changeToMaster() {
37        master = true;
38    }
39
40    /*@ ensures this.name == name;
41       @ assignable this.name;
42       */
43    public void setName(String name) {
44        this.name = name;
45    }
46
47    /*@ ensures \result == name;
48       */
49    public /*@ pure */ String getName() {
50        return name;
51    }
52 }

```

Listing 12: Full specification of Student with assignable clauses

sibility to statically evaluate the loop body repeatedly until the loop condition evaluates to false. The number of iterations is not static but depends on dynamic input parameters and initial states. In program verification, the dominating solution to this problem is the usage of a *loop invariant* [Floyd, 1967, Hoare, 1969]. This is a formula whose validity is preserved by the loop body (given the loop condition was true before). From this we can conclude that, if the entire loop starts in a state where the loop invariant holds, then it will still hold once the loop terminates, in addition the negated loop condition¹³.

There exists approaches to automated invariant generation [German and Wegbreit, 1975, Karr, 1976], and the recent years saw a very dynamic development in this area. Yet, much more needs to be done to automatically find good invariants, and to integrate that into verification tools. (The bottleneck is currently not to generate formulas that are invariant over the loop body, but to identify those that contribute to the overall correctness proof.) For the time being, finding loop invariants that allow to verify some code unit is still a largely manual task. Guidance on how to write loop invariants is beyond our scope here.

In the first place, loop invariants are proof artifacts, comparable to induction hypotheses in inductive proofs. But JML offers the possibility to annotate loops, in the source code, with invariants, to be used by verification tools during the proof process. The corresponding keyword is `loop_invariant`, followed by a boolean JML expression. The JML comment that contains this must be placed directly in front of the loop. Notice that a loop invariant may contain an `\old(E)` expression. This refers to the value of the expression *E* before the method started, *not* to the value of *E* at the previous iteration of the loop.

As long as no `diverges` clause (see Sect. 4.3) is defined, it is required to prove that a method terminates. In the presence of a loop this is only possible if a `decreasing` clause (also named *variant*) is provided together with the loop invariant. The decreasing term must be wellfounded, which means that it cannot decrease forever. For the decreasing clause, it has to be shown that it is strictly decreasing for each loop iteration and that it evaluates to a nonnegative value in any state satisfying the invariant. Therefore, this is sufficient to conclude that the loop terminates. In JML the decreasing term is specified via keyword `decreasing`, followed by an expression of type integer.

Example 27. The loop invariant in method `search` in Listing 13 shows a very common loop invariant pattern for methods iterating over an array. All the elements that have been examined so far respect a certain property, and the loop terminates at least when all the elements in the array have been examined. Variable `found` indicates in this example whether the element to search is contained in the already examined elements or not. A loop invariant restricting the range of loop variables is typically always needed, but not sufficient alone. In this example, the range of loop variable `i` is limited to valid array indices (`0 <= i && i <= a.length`). Finally, a well-founded decreasing clause is provided, that allows to prove termination.

Loop invariants are sensitive to the frame problem as discussed for method calls in Sect. 10.1. Basically, it is necessary to specify which variable locations might be changed by a loop and which not. In KeY this is done with the assignable clause. Only non-local locations have to be specified since local variables changed by the loop are automatically added to the set of assignable locations by KeY.

Example 28. Method `sum` of Listing 14 computes the sum of the values provided by an array using a for-each loop. The assignable clause is explicitly set to `\strictly_nothing`

¹³In fact, to reason about Java, it is required to also support abrupt loop termination, caused by an exception or programmatically by a `return`, `break` or `continue` statement.

```

1 /*@ normal_behavior
2   @ requires a != null;
3   @ ensures \result == (\exists int i;
4     @                               0 <= i && i < a.length; a[i] == val);
5   @*/
6 public boolean search(int[] a, int val) {
7     int i = 0;
8     /*@ maintaining !(\exists int j;
9       0 <= j && j < i; a[j] == val);
10      @ maintaining 0 <= i && i <= a.length;
11      @ decreasing a.length - i;
12      @*/
13     while (i < a.length) {
14         if (a[i] == val)
15             return true;
16         i++;
17     }
18     return false;
19 }

```

Listing 13: Loop invariant example to search an element in an array

to make sure that no objects are created during loop execution. Local variables are not listed in the assignable clause since they are automatically added by KeY.

Java 1.5 introduced so called *enhanced for* loops (also called *foreach* loops, see [Gosling et al., 2013, Sect. 14.14]) that iterate over elements of an array or a collection. Here, the index variable is only implicit. As proposed by Cok [2008], the keyword `\index` refers to this value. An example is also shown in Listing 14.

10.3 Assertions and Block Contracts

Sometimes, the program verifier needs some additional guidance in proving a contract. This can be given as an intermediate *assertion*: `assert P`;¹⁴ We have to prove that P is true in this intermediate state. Afterwards, we can use this additional knowledge to prove the overall proof obligation. In this way, assertions in the code are similar to cuts in proofs. JML also provides a dual `assume` statement. It is supposed to be assumed to be true without verifying it.

While the intuition behind these constructs is clear, they perturb the concept of design by contract. In particular, the statement `assume false`; would make any contract trivially satisfied. For this reason, in KeY `assert` and `assume` are replaced by *block contracts* [Wacker, 2012]. The behavior of any Java block can be specified in the same way as a method is specified (see Sect. 2) by placing the specification directly in front of the Java block. It can contain any clause that is available for method contracts. The only differences are: 1. that `\old` represents the value before executing the block, and not the one before executing the method, and 2. that the `\signals_only` definition must be explicitly specified, because a block has no `throws` definition from which it can be

¹⁴JML `assert` statements are not to be confused with Java `assert` statements. The former are only present in specifications and meant to guide the prover. The latter is an actual program statement to be checked at runtime, that raises an exception upon failure.

```

1  /*@ requires array != null;
2  @ ensures \result == (\sum int i;
3  @
4  @*/
5  public static int sum(int[] array) {
6      int sum = 0;
7      /*@ loop_invariant sum == (\sum int j;
8      @
9      @ loop_invariant \index >= 0 && \index <= array.length;
10     @ decreasing array.length - \index;
11     @ assignable \strictly_nothing;
12     @*/
13     for (int value : array) {
14         sum += value;
15     }
16     return sum;
17 }

```

Listing 14: Loop invariant example to compute the sum of an array

computed. Listing 15 shows the usage of a block contract within a longer method. The block itself swaps the value of the two variables x and y .

```

1  public void complicated() {
2      // Some code
3
4      /*@ ensures x == \old(y);
5      @ ensures y == \old(x);
6      @ assignable x, y;
7      @ signals_only \nothing;
8      @*/
9      {
10         y = x + y;
11         x = y - x;
12         y = y - x;
13     }
14
15     // Some code
16 }

```

Listing 15: Usage of a block contract to swap two values

11 Conclusion

This chapter has provided a short overview of the Java Modeling Language (JML), its main features and how it can be used to describe intended program behavior. More information about JML, including people involved in the community effort, the reference manual, tools supporting JML, teaching material, and relevant papers are available from the JML webpage <http://jmlspecs.org/>.

To conclude, we briefly discuss other related program annotation languages, and the wide range of tool support that exists for JML.

11.1 Tool Support for JML

One of the strong points of JML is that many different kinds of tool support exist for it, covering the whole spectrum of formal methods. For an—unfortunately outdated—overview of JML tools, the reader may refer to [Burdy et al., 2003]. We briefly describe a few, more information is available from the JML webpage. It should be noted that most recent tool development, including KeY, aims at combining different kinds of tool support within a single environment. In particular both the JMLEclipse [Chalin et al., 2010] and OpenJML [Cok, 2011] tool suites include their own runtime checker, static analysis tool, and test case generator.

The original developers of JML started the work on JML with runtime checking in mind, i.e., JML should provide support to check pre- and postconditions *during* program execution. Many different tools exist that support this, for different subsets of JML, e.g., JMLRac [Cheon, 2003], ajml [Rebêlo et al., 2008], and as mentioned subtools of JMLEclipse and OpenJML. The runtime checking approach has also been the basis for model checking of JML annotated programs in Bogor: every program annotation is translated into an assertion, that is validated during the software model checking procedure [Robby et al., 2006].

JML is also used for test case generation. JMLunitNG [Zimmerman and Nagmoti, 2010] extends standard unit testing with knowledge derived from the program annotations. KeyTestGen uses information from the KeY prover to improve test case generation. As mentioned, also JMLEclipse and OpenJML provide support for test case generation.

There are also several tools that support static checking of JML annotations, i.e., at compile time, without executing the program. These tools differ in the level of automation and the support they provide for manually constructing a proof. In general, the more user intervention is possible, the more complex properties can be verified. KeY is a typical example of a tool that can verify complex properties, but requires manual intervention. Other tools in this category are Krakatoa [Marché et al., 2004] and KIV [Balsler et al., 2000]. ESC/Java [Leino et al., 2000] and its followup ESC/Java2 [Cok and Kiniry, 2005] intended to provide automatic support for proving program correctness (if necessary, compromising soundness or completeness). Another tool that has been developed with automation in mind is JACK [Barthe et al., 2006], however it also provided support to fall back on interactive proving using Coq. Also the static verification subtools of JMLEclipse and OpenJML are developed with automation in mind. Finally, the VerCors tool set [Amighi et al., 2012] combines separation logic support for concurrent programs with JML annotations.

Last, it should be mentioned that there are also very different tools that support JML. There is a JMLdoc facility that allows one to generate webpages for JML annotations (similar to Javadoc). There also exist tools that generate JML annotations. These range from generating arbitrary JML specifications such as Daikon [Ernst et al., 2007], and Houdini [Flanagan and Leino, 2000] to tools that can generate one specific class of annotations, such as Chase [Cataño and Huisman, 2003].

11.2 Other Program Annotation Languages

The JML language has been a pioneer in the area of *annotation based specification languages* dedicated to a single programming language. As explained above, in Sect. 1, the intention of the developers was to provide a language to write assertions for Java programs. Its design has been inspired by earlier experiences of some of the developers on annotating Modula 3 [Leino and Nelson, 1998], and C++ (the Larch project) [Cheon and Leavens, 1994].

As a major difference to more abstract specification languages, such as UML [Rumbaugh et al., 2010], Z [Spivey, 1992], VDM [Fitzgerald et al., 2008], Alloy [Jackson, 2003], and the B method [Abrial, 1996], JML focuses solely on the phases of software development in which source code is written. It is also primarily intended to specify existing code, rather than to implement programs according to a preexisting specification. However, it should be noted that some work has been done on translating specifications in these high level languages into JML, e.g., for B [Cataño et al., 2012].

JML also has a number of similarities to OCL [Warmer and Kleppe, 1999], a language for annotating UML class diagrams with constraints on object states. It is used for both meta modeling and application modeling. In the latter case, annotations are added to the fine design of the implementation, much like class and method specifications in JML. But unlike JML, OCL does not subscribe to any programming language, and therefore does not address language-specific concerns (like, e.g., exceptions). Earlier versions of KeY supported OCL as well [Beckert et al., 2007], but this has been discontinued.

JML has been an inspiration for many other program annotation languages that have emerged over the last years, such as the ANSI/ISO C Specification Language (ACSL) [Baudin et al., 2010], and the language of the VCC tool (formerly “Verifying C Compiler”) [Cohen et al., 2009], Spec# for C# [Barnett et al., 2005b], and Dafny [Leino, 2010], that is an integrated annotation *and* programming language.

Recently, separation logic [O’Hearn et al., 2001, 2004] has become a popular alternative to Hoare logic to specify program behavior. Separation logic allows explicit reasoning about the heap, that makes it suitable for reasoning about pointer programs, and for concurrent programs. Several approaches exist that combine separation logic with JML (or JML like languages), to enable reasoning about pointers and/or concurrent programs, while maintaining the expressiveness of JML [Tuerk, 2009, Jacobs and Piessens, 2011, Amighi et al., 2012].

References

- Jean-Raymond Abrial. *The B Book: Assigning Programs to Meanings*. Cambridge University Press, August 1996.
- Afshin Amighi, Stefan Blom, Marieke Huisman, and Marina Zaharieva-Stojanovski. The VerCors project: Setting up basecamp. In *Programming Languages meets Program Verification (PLPV 2012)*, pages 71–82. ACM, 2012. doi: 10.1145/2103776.2103785.
- Michael Balser, Wolfgang Reif, Gerhard Schellhorn, Kurt Stenzel, and Andreas Thums. Formal system development with KIV. In T. S. E. Maibaum, editor, *FASE*, volume 1783 of *Lecture Notes in Computer Science*, pages 363–366. Springer, 2000. ISBN 3-540-67261-3.

- M. Barnett, R. DeLine, M. Fähndrich, K. R. M. Leino, and W. Schulte. Verification of object-oriented programs with invariants. *Journal of Object Technology*, 3(6):27–56, 2004.
- M. Barnett, B.-Y. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *Formal Methods for Components and Objects*, volume 4111 of *Lecture Notes in Computer Science*. Springer-Verlag, 2005a.
- Mike Barnett, K. Rustan M. Leino, and Wolfram Schulte. The Spec# programming system: an overview. In Gilles Barthe, Lilian Burdy, Marieke Huisman, Jean-Louis Lanet, and Traian Muntean, editors, *Post Conference Proceedings of CASSIS: Construction and Analysis of Safe, Secure and Interoperable Smart devices, Marseille*, volume 3362 of *LNCS*, pages 49–69. Springer, 2005b. URL <http://www.springerlink.com/openurl.asp?genre=article&issn=0302-9743&volume=3362&spage=151>.
- Mike Barnett, David A. Naumann, Wolfram Schulte, and Qi Sun. 99.44% pure: Useful abstractions in specification. In *ECOOP Workshop FTJP'2004 Formal Techniques for Java-like Programs*, pages 51–60, January 2005c.
- G. Barthe, M. Pavlova, and G. Schneider. Precise analysis of memory consumption using program logics. In *Software Engineering and Formal Methods*, pages 86–95. IEEE Press, 2005.
- G. Barthe, L. Burdy, J. Charles, B. Grégoire, M. Huisman, J.-L. Lanet, M. Pavlova, and A. Requet. JACK: A tool for validation of security and behaviour of Java applications. pages 152–174, 2006. doi: 10.1007/978-3-540-74792-5_7.
- Patrick Baudin, Pascal Cuoq, Jean-Christophe Filliâtre, Claude Marché, Benjamin Monate, Yannick Moy, and Virgile Prevosto. *ACSL: ANSI/ISO C Specification Language*. CEA LIST and INRIA, 2010. URL <http://frama-c.cea.fr/acsl.html>. Version 1.5.
- Bernhard Beckert, Reiner Hähnle, and Peter H. Schmitt, editors. *Verification of Object-Oriented Software: The KeY Approach*. Number 4334 in *Lecture Notes in Computer Science*. Springer-Verlag, 2007.
- A. Borgida, J. Mylopoulos, and R. Reiter. On the frame problem in procedure specifications. *IEEE Transactions on Software Engineering*, 21(10):785–798, 1995.
- C. Breunesse, N. Cataño, M. Huisman, and B. Jacobs. Formal methods for smart cards: an experience report. *Science of Computer Programming*, 55:53–80, 2005.
- Daniel Bruns. Formal semantics for the Java Modeling Language. Diploma thesis, Universität Karlsruhe, 2009. URL <http://digbib.ubka.uni-karlsruhe.de/volltexte/1000012399>.
- 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. In Thomas Arts and Wan Fokkink, editors, *Proc. Eighth International Workshop on Formal Methods for Industrial Critical Systems (FMICS 03)*, volume 80 of *Electronic Notes in Theoretical Computer Science*, pages 73–89. Elsevier, 2003. URL <http://www1.elsevier.com/gej-ng/31/29/23/137/23/27/80006.pdf>.

- Néstor Cataño, Tim Wahls, Camilo Rueda, Víctor Rivera, and Danni Yu. Translating b machines to jml specifications. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing, SAC '12*, pages 1271–1277, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-0857-1. doi: 10.1145/2245276.2231978. URL <http://doi.acm.org/10.1145/2245276.2231978>.
- N. Cataño and M. Huisman. Chase: A static checker for JML’s assignable clause. volume 2575 of *Lecture Notes in Computer Science*, pages 26–40. Springer-Verlag, 2003.
- Patrice Chalin. Improving JML: For a safer and more effective language. In Keijiro Araki, Stefania Gnesi, and Dino Mandrioli, editors, *Formal Methods (FME)*, volume 2805 of *Lecture Notes in Computer Science*, pages 440–461. Springer-Verlag, 2003. URL <http://www.springerlink.com/content/26cpmd9b3vbgd2et>.
- Patrice Chalin. JML support for primitive arbitrary precision numeric types: Definition and semantics. *Journal of Object Technology*, 3(6):57–79, June 2004. URL http://www.jot.fm/issues/issue_2004_06/article3. Special issue: ECOOP 2003 Workshop on FTfJP.
- Patrice Chalin, Robby, Perry R. James, Jooyong Lee, and George Karabotsos. Towards an industrial grade IVE for java and next generation research platform for JML. *STTT*, 12(6):429–446, 2010.
- Y. Cheon. *A Runtime Assertion Checker for the Java Modeling Language*. PhD thesis, 2003. Technical Report 03-09.
- Yoonsik Cheon and Gary T. Leavens. A quick overview of Larch/C++. *Journal of Object-oriented Programming*, 7(6):39–49, 1994.
- Ernie Cohen, Markus Dahlweid, Mark Hillebrand, Dirk Leinenbach, Michał Moskal, Thomas Santen, Wolfram Schulte, and Stephan Tobies. VCC: A practical system for verifying concurrent C. In Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel, editors, *Theorem Proving in Higher Order Logics, 22nd International Conference, TPHOLs 2009*, volume 5674 of *Lecture Notes in Computer Science*, pages 23–42, Berlin, August 2009. Springer-Verlag.
- David Cok. Adapting JML to generic types and Java 1.6. In *Seventh International Workshop on Specification and Verification of Component-Based Systems (SAVCBS 2008)*, number CS-TR-08-07 in Technical Report, pages 27–34, 2008.
- David Cok. OpenJML: JML for Java 7 by extending OpenJDK. In Mihaela Bobaru, Klaus Havelund, Gerard Holzmann, and Rajeev Joshi, editors, *NASA Formal Methods*, volume 6617 of *Lecture Notes in Computer Science*, pages 472–479. Springer-Verlag, Berlin, 2011. doi: http://dx.doi.org/10.1007/978-3-642-20398-5_35.
- David Cok and Gary T. Leavens. Extensions of the theory of observational purity and a practical design for JML. In *Seventh International Workshop on Specification and Verification of Component-Based Systems (SAVCBS 2008)*, number CS-TR-08-07 in Technical Report, pages 43–50, 4000 Central Florida Blvd., Orlando, Florida, 32816-2362, 2008. School of EECS, UCF. URL <http://www.eecs.ucf.edu/SAVCBS/2008/papers/Cok-Leavens.pdf>.

- David R. Cok and Joseph Kiniry. ESC/Java2: Uniting ESC/Java and JML. In Gilles Barthe, Lilian Burdy, Marieke Huisman, Jean-Louis Lanet, and Traian Muntean, editors, *Post Conference Proceedings of CASSIS: Construction and Analysis of Safe, Secure and Interoperable Smart devices, Marseille*, volume 3362 of *LNCS*, pages 108–128. Springer-Verlag, 2005. URL <http://www.springerlink.com/openurl.asp?genre=article&issn=0302-9743&volume=3362&spage=151>.
- Á. Darvas and R. Leino. Practical reasoning about invocations and implementations of pure methods. In *FASE*, volume 4422 of *Lecture Notes in Computer Science*, pages 336–351. Springer-Verlag, 2007.
- Á. Darvas and P. Müller. Reasoning about method calls in JML specifications. In *Workshop on Formal Techniques for Java Programs*, 2005.
- Ádám Darvas and Peter Müller. Formal encoding of JML Level 0 specifications in JIVE. Technical Report 559, ETH Zurich, 2007.
- Krishna Kishore Dhara and Gary T. Leavens. Forcing behavioral subtyping through specification inheritance. Technical Report 95-20c, Department of Computer Science, Iowa State University, Ames, Iowa, 50011, December 1995. URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR95-20/TR.ps.gz>.
- Christian Engel. A translation from JML to JavaDL. Studienarbeit, Fakultät für Informatik, Universität Karlsruhe, February 2005.
- Michael D. Ernst, Jeff H. Perkins, Philip J. Guo, Stephen McCamant, Carlos Pacheco, Matthew S. Tschantz, and Chen Xiao. The Daikon system for dynamic detection of likely invariants. *Science of Computer Programming*, 69(1–3):35–45, December 2007. ISSN 0167-6423.
- John S. Fitzgerald, Peter Gorm Larsen, and Marcel Verhoef. Vienna development method. In Benjamin W. Wah, editor, *Wiley Encyclopedia of Computer Science and Engineering*. John Wiley & Sons, Inc., 2008.
- C. Flanagan and K.R.M. Leino. Houdini, an annotation assistant for ESC/Java. Technical Report 2000-003, DEC-SRC, December 2000.
- Robert W. Floyd. Assigning meanings to programs. In J. T. Schwartz, editor, *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages 19–32, Providence, Rhode Island, 1967. American Mathematical Society.
- S.M. German and Ben Wegbreit. A synthesizer of inductive assertions. *IEEE Transactions on Software Engineering*, SE-1(1):68–75, March 1975. ISSN 0098-5589. doi: 10.1109/TSE.1975.6312821.
- James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. *The Java Language Specification, Java SE 7 Edition*. The Java Series. Addison-Wesley, Boston, Mass., 2013. ISBN 978-0133260229.
- John V. Guttag and James J. Horning. *Larch: Languages and Tools for Formal Specification*. Springer-Verlag, 1993.
- C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 583, October 1969.

- Daniel Jackson. Alloy: A logical modelling language. In Didier Bert, Jonathan P. Bowen, Steve King, and Marina A. Waldén, editors, *ZB 2003: Formal Specification and Development in Z and B, Third International Conference of B and Z Users, Turku, Finland, June 4-6, 2003, Proceedings*, volume 2651 of *Lecture Notes in Computer Science*, page 1. Springer, 2003. ISBN 3-540-40253-5.
- B. Jacobs and F. Piessens. Expressive modular fine-grained concurrency specification. In *POPL*, pages 271–282, 2011.
- Bart Jacobs and Erik Poll. A logic for the java modeling language JML, January 24 2001. URL <http://citeseer.ist.psu.edu/501115.html>; <http://www.cs.kun.nl/~erikpoll/publications/jmllogic.ps.gz>.
- Michael Karr. Affine relationships among variables of a program. *Acta Informatica*, 6(2):133–151, 1976. ISSN 0001-5903. doi: 10.1007/BF00268497. URL <http://dx.doi.org/10.1007/BF00268497>.
- Gary T. Leavens and Yoonsik Cheon. Preliminary design of Larch/C++. In U. Martin and J. Wing, editors, *Proceedings of the First International Workshop on Larch, July, 1992*, Workshops in Computing, pages 159–184. Springer-Verlag, New York, NY, 1993.
- 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, editors, *Foundations of Component-Based Systems*, pages 113–135. Cambridge University Press, 2000. URL citeseer.ist.psu.edu/leavens00concepts.html.
- Gary T. Leavens and David A. Naumann. Behavioral subtyping is equivalent to modular reasoning for object-oriented programs. Technical Report 06-36, Department of Computer Science, Iowa State University, Ames, Iowa, 50011, December 2006. URL <ftp://ftp.cs.iastate.edu/pub/techreports/TR06-36/TR.pdf>.
- Gary T. Leavens, Albert L. Baker, and Clyde Ruby. Preliminary design of JML: a behavioral interface specification language for Java. *ACM SIGSOFT Software Engineering Notes*, 31(3):1–38, 2006. URL <http://doi.acm.org/10.1145/1127878.1127884>.
- Gary T. Leavens, Erik Poll, Curtis Clifton, Yoonsik Cheon, Clyde Ruby, David Cok, Peter Müller, Joseph Kiniry, Patrice Chalin, Daniel M. Zimmerman, and Werner Dietl. *JML Reference Manual*, May 31, 2013. URL <http://www.eecs.ucf.edu/~leavens/JML/documentation.shtml>. Draft Revision 2344.
- Gary Todd Leavens. *Verifying Object-Oriented Programs that use Subtypes*. PhD thesis, Massachusetts Institute of Technology, December 1988. Published as MIT/LCS/TR-439 in February 1989.
- K. Rustan M. Leino. Dafny: An automatic program verifier for functional correctness. In Edmund M. Clarke and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 6355 of *Lecture Notes in Computer Science*, pages 348–370. Springer Berlin Heidelberg, 2010. ISBN 978-3-642-17510-7. doi: 10.1007/978-3-642-17511-4_20.

- K. Rustan M. Leino and Peter Müller. Object invariants in dynamic contexts. In M. Odersky, editor, *European Conference on Object-Oriented Programming*, volume 3086 of *Lecture Notes in Computer Science*, pages 491–516. Springer-Verlag, 2004. Available from www.sct.inf.ethz.ch/publications/index.html.
- K. Rustan M. Leino and Greg Nelson. An extended static checker for Modula-3. In Kai Koskimies, editor, *Compiler Construction, 7th International Conference, CC'98, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, volume 1383 of *Lecture Notes in Computer Science*, pages 302–305. Springer, 1998. ISBN 3-540-64304-4.
- K. Rustan M. Leino, Greg Nelson, and J.B. Saxe. ESC/Java user's manual. Technical Report SRC 2000-002, Compaq System Research Center, 2000.
- Barbara Liskov. Data abstraction and hierarchy. *SIGPLAN Notices*, pages 17–34, May 1988.
- Barbara Liskov and Jeanette M. Wing. Specifications and their use in defining subtypes. In Andreas Paepcke, editor, *Proceedings of the 8th Annual Conference on Object-Oriented Programming Systems, Languages and Applications*, pages 16–28, Washington DC, USA, 1993. ACM Press. ISBN 0-201-58895-1.
- Barbara Liskov and Jeanette M. Wing. A behavioral notion of subtyping. *ACM Transactions on Programming Languages and Systems*, 16(6):1811–1841, November 1994.
- C. Marché, C. Paulin-Mohring, and X. Urbain. The Krakatoa tool for certification of Java/JavaCard programs annotated with JML annotations. *JLAP*, 58:89–106, 2004.
- Bertrand Meyer. Applying “design by contract”. *IEEE Computer*, 25(10):40–51, October 1992.
- Bertrand Meyer. *Object-Oriented Software Construction*. Prentice-Hall, 1997. ISBN 0-13-629155-4.
- Peter Müller, Arndt Poetzsch-Heffter, and Gary T. Leavens. Modular specification of frame properties in JML. *Concurrency Computation Practice and Experience*, 2002. URL [ftp://ftp.cs.iastate.edu/pub/techreports/TR02-02/TR.pdf](http://ftp.cs.iastate.edu/pub/techreports/TR02-02/TR.pdf).
- Peter Müller, Arnd Poetzsch-Heffter, and Gary T. Leavens. Modular invariants for layered object structures. *Science of Computer Programming*, 62(3):253–286, October 2006.
- David A. Naumann. Observational purity and encapsulation. *Theoretical Computer Science*, 376(3):205–224, 2007.
- Bashar Nuseibeh. Ariane 5: Who dunnit? *IEEE Software*, 14(3):15–16, May / June 1997. ISSN 0740-7459.
- P. W. O'Hearn, J. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In L. Fribourg, editor, *CSL*, volume 2142 of *Lecture Notes in Computer Science*, pages 1–19. Springer-Verlag, 2001. Invited paper.

- P. W. O’Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. In *POPL*, pages 268–280. ACM, January 2004.
- Arndt Poetzsch-Heffter. *Specification and Verification of Object-Oriented Programs*. PhD thesis, Technical University of Munich, 1997. Habilitation thesis.
- A. D. Raghavan and G. T. Leavens. Desugaring JML method specifications. Technical Report TR #00-03e, Department of Computer Science, Iowa State University, 2000. Current revision from May 2005.
- Henrique Rebêlo, Ricardo Massa Ferreira Lima, Márcio Cornélio, and Sérgio Soares. A jml compiler based on aspectj. In *First International Conference on Software Testing, Verification, and Validation, ICST 2008, Lillehammer, Norway, April 9-11, 2008*, pages 541–544. IEEE Computer Society, 2008.
- Robby, Edwin Rodríguez, Matthew B. Dwyer, and John Hatcliff. Checking jml specifications using an extensible software model checking framework. *STTT*, 8(3):280–299, 2006.
- James Rumbaugh, Ivar Jacobson, and Grady Booch. *The Unified Modeling Language Reference Manual*. Object Technology Series. Addison-Wesley, Reading/MA, 2nd edition, 2010.
- Steffen Schlager. Handling of integer arithmetic in the verification of Java programs. Master’s thesis, July 10 2002. URL <http://citeseer.ist.psu.edu/530777.html>; <http://i12www.ira.uka.de/~key/doc/2002/DA-Schlager.ps.gz>.
- J. Michael Spivey. *Z Notation - a reference manual (2. ed.)*. Prentice Hall International Series in Computer Science. Prentice Hall, 1992. ISBN 978-0-13-978529-0.
- Thomas Tuerk. A formalisation of smallfoot in HOL. In Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel, editors, *TPHOL*, volume 5674 of *Lecture Notes in Computer Science*, pages 469–484. Springer-Verlag, 2009.
- J. van den Berg and B. Jacobs. The LOOP compiler for Java and JML. In T. Margaria and W. Yi, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2001)*, volume 2031 of *Lecture Notes in Computer Science*, pages 299–312. Springer-Verlag, 2001.
- Simon Wacker. Blockverträge. Studienarbeit, Karlsruhe Institute of Technology, 2012.
- Jos Warmer and Anneke Kleppe. *The Object Constraint Language: Precise Modelling with UML*. Object Technology Series. Addison-Wesley, Reading/MA, 1999.
- Marina Zaharieva-Stojanovski and Marieke Huisman. Verifying class invariants in concurrent programs. In Stefania Gnesi and Arend Rensink, editors, *Fundamental Approaches to Software Engineering, (FASE)*. Springer-Verlag, 2014.
- D.M. Zimmerman and R. Nagmoti. JMLUnit: The Next Generation. In B. Beckert and C. Marché, editors, *1st International Conference on Formal Verification of Object-Oriented Software (FoVeOOS 2010)*, volume 6528 of *Lecture Notes in Computer Science*. Springer-Verlag, 2010.