

# Verifying Class Invariants in Concurrent Programs

Marina Zaharieva-Stojanovski and Marieke Huisman

University of Twente, the Netherlands

**Abstract.** Class invariants are a highly useful feature for the verification of object-oriented programs, because they can be used to capture all valid object states. In a sequential program setting, the validity of class invariants is typically described in terms of a *visible state semantics*, i.e., invariants only have to hold whenever a method begins or ends execution, and they may be broken inside a method body. However, in a concurrent setting, this restriction is no longer usable, because due to thread interleavings, any program state is potentially a visible state.

In this paper we present a new approach for reasoning about class invariants in multithreaded programs. We allow a thread to explicitly break an invariant at specific program locations, while ensuring that no other thread can observe the broken invariant. We develop our technique in a permission-based separation logic environment. However, we deviate from separation logic's standard rules and allow a class invariant to express properties over shared memory locations (the *invariant footprint*), independently of the permissions on these locations. In this way, a thread may break or reestablish an invariant without holding permissions to all locations in its footprint. To enable modular verification, we adopt the restrictions of Müller's ownership-based type system.

## 1 Introduction

In object-oriented programs, class invariants are typically used to express properties about the object's state that should hold throughout the object's life cycle. However, in practice it is often impossible to maintain the invariant continuously. For example, for an invariant that expresses a relation between fields  $x$  and  $y$ ,  $x == y$ , when  $x$  is updated,  $y$  must also be updated, and both updates can not be done atomically. Therefore, invariant theory should provide for the possibility that a class invariant is temporarily *broken* at specific program parts.

In the sequential setting, the theory about invariant validity is well-developed; in essence, class invariants only have to hold in the program's *visible states*, i.e., in pre- and poststates of public methods [17]. In particular, if a class invariant  $I$  holds in a method's prestate, the method must end in a state satisfying  $I$ .

However, in the setting of multithreading programs, this approach can not be carried over directly. Due to possible interference between parallel threads, any program state may be *visible*. For example, when the field  $x$  in the invariant above is updated, any other thread might observe this change and the broken invariant. This problem is sometimes called a *high-level data race* [2].

Therefore, this paper defines an approach to define validity of class invariants in a multithreaded setting. Our approach supports explicit *breaking of invariants*, under the condition that other threads can not see that the invariant is broken. We build our technique on *permission-based separation logic* [4], using a Java-like language. However, in contrast to standard separation logic, we explicitly make a distinction between *state formulas*, which describe a property about the shared state, and *resource formulas*, which describe when a thread holds a permission to access a certain location. We ensure modular verification using the restrictions from *ownership-based type systems* [7].

Our approach works as follows. A class invariant is specified as a condition on the shared memory. For each class invariant, we maintain a token that indicates whether the class invariant can be inspected. This token can be split and combined: if a thread has the complete token, it can *break* the invariant; otherwise it can only *use* it. Breaking the invariant is done by executing a (specification-only) **unpack** statement. When a thread reestablishes the invariant, the token to inspect the invariant becomes available again for other threads to break or inspect the invariant. This behaviour is modeled by a (specification-only) **pack** statement. Thus, within the unpacked segment, a thread is free to do whatever it wants with the class invariant, as our verification approach ensures that no other thread can observe the invariant in parallel.

To guarantee that class invariants can be verified in a modular way, when a class invariant is broken, a thread is not allowed to obtain any new permissions anymore. In particular, if a thread requires a lock to change any of the fields associated to the invariant, it should obtain this lock before breaking the invariant. This requirement shows that there is close connection between the locking strategy and the functional invariant properties that can be maintained in an application. Further, it is important that with our approach, a thread does not need to have all access permissions that are associated with the invariant, but only the access permissions needed to break the invariant; all other variables are implicitly assumed to be unchanged. Moreover, our technique does allow creating new (helper) threads when an invariant is broken; however, these threads need to be finished and joined before the invariant is reestablished again.

The main contribution of this paper is a sound modular technique for verification of class invariants in multithreaded programs, which:

- is flexible and permissive, because it allows a thread to break an invariant without holding all permissions associated to the invariant property; and
- reveals the connection between locking policy and invariant properties that can be maintained.

The motivation and applicability of our approach is illustrated on several examples. Its implementation as part of the VerCors tool set is under development.

**Outline.** We begin by introducing a short overview of permissions in separation logic, Sec. 2. Next, in Sec. 3 we present the main concepts of our approach, which is further formalised in Sec. 4. Sec. 5 reviews others approaches that tie in with our work. Finally, in Sec. 6 we summarise our work and discuss our future plans.

## 2 Background

This paper builds on Parkinson’s work on separation logic for Java-like programs [21], and its extension by Haack *et al.* [11] for concurrency.

Separation logic [23] is an extension of Hoare Logic [12] for reasoning about separate parts of the heap. The base of this logic is the binary *separating conjunction* operation:  $P * Q$  describes that  $P$  and  $Q$  hold for disjoint parts of the heap. O’Hearn shows that separation logic is also convenient for reasoning about multithreaded programs [19]. To allow parallel reads of the same data, basic separation logic is extended with *fractional permissions* [4]. Permission  $\pi$  is a value in the domain  $(0, 1]$ . At any point in time, a thread holds a number of permissions on locations. If a thread has a write permission for a certain location, i.e., the value 1, it is allowed to change this location. If a thread has a fractional permission, i.e., a fraction less than 1, then it may only read this location. Permissions can be split and combined, to change between read and write permissions. The soundness of this logic ensures that the sum of all threads’ permissions for a certain location never exceeds 1, which guarantees data-race freedom. The predicate  $\text{Perm}(x.f, \pi)$  indicates that  $x.f$  points to a location for which the actual thread has a permission  $\pi$ . Permission expressions are combined with the *separating conjunction* operation.

Parkinson adapts separation logic for object-oriented concepts in a Java-like language [21]. He proposes *abstract predicates* [20] to provide abstraction. Later, Haack *et al.* extended this logic to show how to reason about multithreaded Java-like programs [11] that include reentrant locks and dynamic thread creation. For each lock, a *resource invariant* is specified, i.e., an abstract predicate describing which permissions are stored in the lock. A newly created lock is still **fresh** and not ready to be acquired. The thread must first execute the **commit** command on the lock, which transfers the permissions from the thread to the lock and changes the lock’s state to **initialized**. Any thread then may acquire the **initialized** lock to get the resource invariant (except for reentrant acquiring). Upon final release of the lock, the thread returns the resource invariant back to the lock.

## 3 Verification Methodology for Class Invariants

This section gives a conceptual understanding of our methodology, presented from two different aspects. First, we discuss how we model the *invariant protocol*, i.e., when an invariant may be assumed, and how it can be broken and reestablished. Then, we describe how our method supports modular verification.

### 3.1 Class Invariant Protocol

We assume that class invariants express properties over non-static class fields. Thus, a class invariant  $I$  defined in a class  $C$  is always associated with a particular object  $v$  of class  $C$ , we write  $v.I$ . We call the set of locations referred to by an invariant  $v.I$  the *footprint* of  $v.I$ , denoted  $\text{fp}(v.I)$  (formally defined in Sec. 4).

*Assuming a Class Invariant.* Our technique should guarantee absence of high-level data races; therefore, it should control access to the invariant's footprint. To provide this control, to every invariant  $v.I$ , we associate a special abstract predicate  $\text{holds}(v.I, 1)$ , distributed as a token among the threads. The intuitive meaning of this predicate is the following: when a thread holds a predicate  $\text{holds}(v.I, \pi)$ ,  $\pi > 0$ , it may assume that the invariant  $v.I$  holds; if  $\pi = 1$ , the running thread may additionally break the invariant. The predicate might be divided among different threads by using the following equivalence:

$$\text{holds}(v.I, \pi) * - * \text{holds}(v.I, \pi/2) * \text{holds}(v.I, \pi/2)$$

This approach guarantees that: 1) a class invariant  $v.I$  is stable and all threads that hold a token  $\text{holds}(v.I, \pi)$  may rely on  $v.I$ 's correctness; or 2) *at most one* thread has the token  $\text{holds}(v.I, 1)$  and no other thread may assume  $v.I$ .

*Breaking a Class Invariant.* Inspired by the work of Leino *et al.* [14], we explicitly specify the segment in the program where an invariant property might be violated: for an invariant  $v.I$ , specification command  $\text{unpack}(v.I)$  must be executed at the beginning of such a segment, and  $\text{pack}(v.I)$  at its end. The segment between both commands is called an *unpacked segment of  $v.I$* . A special case is object initialisation: the program segment between the end of  $v$ 's construction and the first execution of the  $\text{pack}(v.I)$  command is also  $v.I$ 's unpacked segment.

The  $\text{unpack}(v.I)$  command consumes the token  $\text{holds}(v.I, 1)$ , and issues a predicate  $\text{unpacked}(v.I, 1)$  (*breaking token*). This token serves as a license for the thread to break the invariant  $v.I$ . Once all updates are done, the running thread must reestablish the validity of  $v.I$  and call the  $\text{pack}(v.I)$  command, which trades the  $\text{unpacked}(v.I, 1)$  token for the  $\text{holds}(v.I, 1)$  token. The  $\text{unpack}(v.I)$  command is always followed by  $\text{pack}(v.I)$  within the same method and executed by the same thread. This thread is called a *holder* of the unpacked segment.

Lst. 1 illustrates the use of an unpacked segment: a class `Point`, represents a point lying on or above the line  $y = -x$ . Since method `move()` updates the fields  $x$  and  $y$  to which invariant  $I$  refers, these updates must happen within an unpacked segment of  $I$ . (The annotation `safe` at line 7 is discussed next.)

*Restrictions to Unpacked Segments.* We showed how a thread obtains permission to modify an invariant footprint location  $p.f$ . Once  $p.f$  is assigned, we say that  $p.f$  is in a *critical state* until the end of the unpacked segment. More precisely:

**Definition 1.** (*Critical state of a location*) Let  $v.I$  be an invariant,  $p.f$  a location, such that  $p.f \in \text{fp}(v.I)$ , and let  $p.f$  be assigned inside an unpacked segment of  $v.I$ . Then, any program execution state between the assignment and the end of the unpacked segment is a critical state for  $p.f$ .

To prevent a thread to observe a broken invariant, a location in a critical state must not be publicly exposed. Therefore, within an unpacked segment we forbid the running thread to release permissions and make them accessible to other threads. Concretely, within an unpacked segment, we allow only `safe`

```

class Point {
2   int x; int y;
   // @ invariant I : this.x + this.y >= 0;
4   // ...constructors
   // @ requires holds(this.I, 1) * Perm(this.x, 1) * Perm(this.y, 1);
6   // @ ensures holds(this.I, 1) * Perm(this.x, 1) * Perm(this.y, 1);
   /* safe @ */ void move() {
8       // the invariant I may now be assumed because of the holds token
       {holds(this.I, 1) * Perm(this.x, 1) * Perm(this.y, 1) * this.I}
10      // @ unpack(this.I);      // trades holds token for unpacked token
       {unpacked(this.I, 1) * Perm(this.x, 1) * Perm(this.y, 1) * this.I}
12      this.x = this.x - 1;      // the invariant I is broken
       this.y = this.y + 1;      // the invariant I can now be reestablished
14      {unpacked(this.I, 1) * Perm(this.x, 1) * Perm(this.y, 1) * this.I}
       // @ pack(this.I);        // trades unpacked token for holds token
16      {holds(this.I, 1) * Perm(this.x, 1) * Perm(this.y, 1)}
       }}

```

**Lst. 1.** Unpacked segment of a class invariant

commands, i.e., commands that exclude any lock-related operation (acquiring, releasing or committing a lock). This means that all permissions used in the unpacked segment must be obtained before the segment begins. A **safe** command may call only **safe** methods, i.e., methods composed of **safe** commands only. These methods are specified with the optional modifier **safe** (see Lst. 1, line 7).

We allow forking a **safe** thread, i.e., threads with a **safe** *run()* method, under the condition that the thread must be joined within the unpacked segment. We call these threads *local to the segment*. A **safe** thread may further fork other **safe** threads. The breaking token might be shared among all local threads of the unpacked segment, and thus, they might all update different locations of the invariant footprint in parallel. For this purpose, we define the following axiom:

$$\text{unpacked}(v.I, \pi) * - * \text{unpacked}(v.I, \pi/2) * \text{unpacked}(v.I, \pi/2)$$

Lst. 2 shows a modified version of the **move** method (from Lst. 1) that can not be verified since acquiring/releasing a lock is used within the unpacked segment.

*Object Initialisation.* In our language, object initialisation (the object constructor) is divided into two steps: 1) *object construction* creates an empty object  $v$  (all  $v$ 's fields get a default value), and gives the running thread write permission for each of  $v$ 's fields and a token  $\text{unpacked}(v.I, 1)$  for each invariant  $v.I$ . 2) the **init** method follows obligatorily after object construction, where object fields are initialised. Additionally, for every invariant  $v.I$ , the **pack**( $v.I$ ) is called by default at the end of the **init** method. Hence, at the end of  $v$ 's initialisation, all  $v$ 's invariants hold, and therefore,  $v$  is a valid object.

```

Lock lock; // resource invariant:  $\text{Perm}(x, 1) * \text{Perm}(y, 1)$ ;
2 // @ requires holds(this.l, 1);
  // @ ensures holds(this.l, 1);
4 void move(){
  // @ unpack(this.l); // trades holds token for unpacked (breaking) token
6  lock.lock(); // invalid call (permissions to x and y must be gained before unpacking)
  t.fork(); // another thread t may get half of the breaking token to modify x
8  updateY(); // for updating y another method is called, which must be safe
  lock.unlock(); // invalid call, must happen after packing
10 t.join(); // t is a safe thread, thus joining must be before packing
  // @ pack(this.l);
12 }

```

**Lst. 2.** Restrictions to unpacked segments

A verified program with our approach is free of *high-level data races*. This is expressed by the following theorem:

**Theorem 1.** (*High-level data race freedom*) *If a value  $p.f$  is in a critical state  $s$  of an unpacked segment  $S$  of an invariant  $v.I$ , then any thread that is neither holder nor a local thread of  $S$  can not access  $p.f$ .*

*Proof.* See [25].

As discussed initially, a thread that holds a token  $\text{holds}(v.I, \pi)$ ,  $\pi > 0$  may use the invariant  $v.I$ . This is justified by the following theorem:

**Theorem 2.** (*Use of a class invariant*) *An invariant  $v.I$  holds in a program state in which the running thread  $t$  holds the predicate  $\text{holds}(v.I, \pi)$ ,  $\pi > 0$ .*

*Proof.* See [25].

Lst. 3 extends the program with the **Point** class (see Lst. 1) to show how a class invariant may be used for verifying a client class. The main thread creates initially a valid **Point** object  $s$  for which the invariant  $s.I$  holds ( $s.x + s.y \geq 0$ ) and obtains the token  $\text{holds}(s.I, 1)$  (lines 3,4). The thread then forks a set of new threads (lines 5-9), passing each of them a reference to  $s$  and part of the  $\text{holds}$  token. Each forked thread has a task to create a sequence of new points at specific locations calculated from the location of  $s$  (line 21). To prove that each new **Point**  $p$  is a valid object ( $p.x + p.y \geq 0$ ) (line 24), each thread uses the class invariant  $s.I$ , which is guaranteed by the token  $\text{holds}(s.I, \pi)$ .

To conclude, we summarise the rules that define the invariant protocol:

**R1** (*Assuming*). A thread  $t$  may assume (use) a class invariant  $v.I$  if  $t$  holds the predicate  $\text{holds}(v.I, \pi)$ ,  $\pi > 0$ .

**R2** (*Breaking*). A thread  $t$  may write on a location  $p.f$  if apart from holding a write permission to  $p.f$ , it holds a breaking token  $\text{unpacked}(v.I, \pi)$ ,  $\pi > 0$  for each invariant  $v.I$  that refers to  $p.f$ , i.e.,  $p.f \in \text{fp}(v.I)$ .

```

class DrawPoints {
2  void create(){
    Point s = new Point (0, 0);
4  //holds(s.l,1) is produced
    for (int k = 1; k<=10; k++){
6      Task t = new Task(s, k);
        //each t gets part of holds token
8      t.fork();
    }
10 //join Task threads
    } }
12

14 class Task {
    Point s; int k;
16 // ... constructors
    //@ requires holds(s.l.  $\pi$ ) * ... ;
18 //@ ensures holds(s.l.  $\pi$ ) * ... ;
    void run(){
20     for (int i = 1; i < 10; i++) {
        int x = s.x+i; int y = s.y+ki;
22 //s.l holds(because of the holds token)
        //use s.l to validate p.l
24     Point p = new Point(x, y);
        draw(p);
26     } } }

```

**Lst. 3.** Using a class invariant for verifying a client class

**R3** (*Reestablishing*). An invariant  $v.I$  must have been reestablished when  $\text{pack}(v.I)$  is executed.

**R4** (*Exchanging tokens*). The token  $\text{unpacked}(v.I, 1)$  is produced at  $v$ 's construction; commands  $\text{unpack}(v.I)$  and  $\text{pack}(v.I)$  exchange the  $\text{holds}(v.I, 1)$  token for the  $\text{unpacked}(v.I, 1)$  token, and vice versa.

### 3.2 Modular Verification

As a second step, we discuss the additional properties needed to support modular verification. In the prestate of the assignment to a location  $p.f$ , rule **R2** requires a breaking token for all invariants that refer to  $p.f$ . However, in the context (class) where the assignment happens, not all invariants in the program are known. To support modularity, the breaking token is only explicitly checked for the invariants of the object  $p$ . Additionally, it is guaranteed that this token is implicitly held for all other invariants. We use Müller's *ownership type system* [7], which is strongly connected to modular verification of invariants [18,3,16,8].

*Ownership-Based Types*. The ownership type system organises the objects in the heap in an *ownership tree*, where each object has one *owner* (either the root of the tree, or another object in the heap). We say that each ancestor of an object  $p$  in the tree is  $p$ 's *transitive owner*. The position of the object  $p$  in the tree is determined on  $p$ 's creation, with an attached required modifier from the set  $\{\text{rep}, \text{peer}, \text{rd}\}$  where: **peer** indicates that  $r$  has the same owner as the object **this**; **rep** specifies that  $r$  is owned by **this**, and **rd**(readonly) is any other relation. Additionally, the **self** modifier is used for references that point to the **this** object. An array  $a$  of object references has an additional modifier to define the relation of each element  $a[i]$  with the **this** reference (see Lst. 4, line 2). When an object changes its context, for example, via transfer as a method parameter, the type of the new reference is determined by applying the *viewpoint adaptation* function  $\triangleright$  (see [25]). For example, if the **this** reference owns  $r$ , while  $r$  owns  $x$ , the type of the reference  $r.x$  in the context of **this** is  $\text{rep} \triangleright \text{rep} = \text{rd}$ .

Additionally, the following discipline is imposed in the program: writing to a field  $p.f$  or a call to a *non-pure* method (i.e. with side-effects) with a receiver  $p$  is forbidden when  $p$  has a modifier **rd**. In this way, each object controls all updates that happen in its transitively owned objects. This guarantees the following:

**RO.** If a field  $p.f$  is modified in a method  $m$ , for each transitive owner  $o$  of  $p$ , the call stack contains a method invocation where  $o$  is a receiver.

We require that all class invariants in the program are *ownership admissible*:

**Definition 2.** A class invariant  $v.I$  is *ownership admissible* if it expresses properties over fields  $p_1.p_2 \dots p_n.f$ , where  $n \geq 1$ ,  $v == p_1$  and  $p_i$  is a *rep field* in the class of  $p_{i-1}$  ( $i = 2..n$ ).

*Verification Technique via Ownership Types.* Based on Def. 2, we observe the following: for a location  $p.f$ , an invariant  $v.I$  may refer to  $p.f$  only if  $v == p$  or  $v$  is a transitive owner of  $p$ . Our verification technique suggests that before assigning to a location  $p.f$ , it is enough to require a breaking token only for the invariants of the object  $p$  ( $p.I$ ) that refer to  $p.f$ . If an invariant  $v.I$ , where  $v$  is a transitive owner of  $p$ , refers to  $p.f$ , then the rule **RO** ensures that assignment of  $p.f$  is preceded by a method call where  $v$  is a receiver. To support modular verification, the check that the actual thread holds a breaking token for  $v.I$  should therefore be a requirement of the method call where object  $v$  is a receiver. More precisely, we replace the rule **R2** listed above with the following two rules:

**R2'** A precondition for assigning a field  $p.f$  requires a token  $\text{unpacked}(p.I, \pi)$  ( $\pi > 0$ ) for each invariant  $I$  of the object  $p$  that refers to  $p.f$ .

**R2''** A precondition for invoking a method  $m$  that assigns a field  $p.f$  requires the token  $\text{unpacked}(\text{this}.I, \pi)$  ( $\pi > 0$ ) for each invariant  $I$  of the **this** object that refers to  $p.f$ .

To establish **R2''**, the contract of the called method  $m$  should provide information to the caller about the locations it assigns to. In permission-based separation logic, assigning to a location  $p.f$  in  $m$  requires a write permission  $\pi = 1$  for  $p.f$ . The caller can identify the locations assignable by  $m$  from the precondition formula  $\text{Pre}_m$ : this is the set of locations for which  $\text{Pre}_m$  requires a write permission, denoted  $\text{wrt}(\text{Pre}_m)$  (see [25]). However,  $\pi$  might also be obtained by acquiring a lock during the execution of  $m$ . We ensure that this scenario is not possible. In particular, if a location  $p.f$  is in the footprint of an invariant  $v.I$ ,  $p.f$  should not be protected by a lock object that is transitively owned by  $v$ , because this would mean that other threads might observe a broken invariant (see the example below). This restriction is imposed by the following rule (the used functions are defined in [25]):

**RL**  $\forall I \in \text{inv}(C); \forall f \in \text{relFld}(C); \text{fld}(I) \cap \text{fldResInv}(\text{classOf}(f)) = \emptyset$

The rule is translated as: for any invariant  $I$  defined in a class  $C$ , and a field  $f$  *relevant* to  $C$ , the set of fields that appear in  $I$  is disjoint from the set of fields



```

class PointsSet {
2   rep rep Point[] points = new rep rep Point[100];
   //@ Invariant  $I_1$ :  $(\forall \text{int } i: 0 \leq i < 100) (\text{points}[i].x \leq 10) * (\text{points}[i].y \leq 10)$ ;
4   //@ requires holds(this. $I_1$ , 1) * Perm(points[i].x, 1) * Perm(points[i].y, 1)
   //@ ensures holds(this. $I_1$ , 1) * Perm(points[i].x, 1) * Perm(points[i].y, 1)
6   void moveAt(int i) {
   //@ unpack(this. $I_1$ ); // trades the holds token for unpacked token
8   if (points[i].y <= 9) {
   //required unpacked token for  $I_1$  (as  $\text{points}[i].x, \text{points}[i].y \in \text{wrt}(\text{Pre}_{\text{move}}) \cap \text{fp}(I_1)$ )
10  points[i].move(); }
   //@ pack(this. $I_1$ ); // trades the unpacked token for holds token
12 } }

```

Lst. 4. Modular verification

that appear in the resource invariant definition in the class of  $f$ . A field  $f$  is *relevant* to a class  $C$  if it may be expressed as a  $p_1.p_2..p_n.f$ , where  $p_1$  is a **rep** field defined in  $C$ , and  $p_i$  is a **rep** or **peer** field in the class of  $p_{i-1}$ ,  $i = 2..n$ ,  $n \geq 1$ .

In Lst. 4, we extend our program (from Lst. 1) to illustrate modular verification. Class `PointsSet` represents a set of points that lie within a predefined area. When calling the method `move()` (line 10), the caller provides a breaking token for its own invariants that `move()` might break (in this case invariant  $I_1$ ). After the call to `move()`, invariant  $I_1$  is reestablished (line 11), even though the actual thread has permissions to the  $i^{\text{th}}$  array element only; our approach ensures that the other locations in  $\text{fp}(I_1)$  are stable until the end of the unpacked segment.

Fields  $x$  and  $y$  from class `Point` are relevant to the `PointsSet` class and used in  $I_1$ ; hence, Rule **RL** forbids a lock that protects  $x$  and/or  $y$  to be transitively owned by a `PointsSet` object. This is necessary: if permissions to  $x$  and  $y$  could be obtained by a lock in `Point`, other threads might observe that  $I_1$  is broken. To avoid this, the lock would have to be already acquired before the unpacked segments for  $I_1$ , but this would violate modularity. The example shows that the invariants that can be maintained strongly depend on the locking strategy used.

## 4 Formalisation

We formalise our approach using a Java-like concurrent language. The formalisation is mainly inspired by Haack *et al.* [11]. We concentrate on those points that are relevant for class invariants. For other concepts, e.g., those associated to locks, we only provide some basic intuition to make the paper self-contained.

### 4.1 Language

Fig. 1 shows the grammar of our language. With  $\bar{x}$  we define sequences of  $x$ , while  $x?$  represents an optional  $x$ . A class is composed of fields, methods, predicates, and class invariants. The special predicate *res\_inv* is associated to a lock

$cl \in \text{Class}$	$::= \text{class } C \{fd * md * inv * pd*\}$
$fd \in \text{Field}$	$::= Tf$
$md \in \text{Method}$	$::= \text{spec } T \ m(\overline{V} \ \overline{x})\{c\}$
$\text{spec} \in \text{MethSpec}$	$::= \text{requires } F \text{ ensures } F \text{ pure? safe?}$
$pd \in \text{Predicate}$	$::= \text{pred } P = F_{\text{res}}(P \neq \text{res\_inv}) \mid \text{pred } \text{res\_inv} = F_{\text{res}}$
$inv \in \text{Invariant}$	$::= \text{Invariant } I : F_{\text{inv}}$
$c \in \text{Command}$	$::= v \text{ (return value or null in case of type void)}$ $\mid T \ x; c \mid x = v; c \mid x = \text{op}(\overline{v}); c \mid x = v.f; c$ $\mid x = \text{new rtype } C; c \mid (x = v.m(\overline{v}); c \mid \text{if } v \text{ then } c \text{ else } c; c$ $\mid v.f = v; c \mid v.\text{lock}(); c \mid v.\text{commit}(); c \mid v.\text{unlock}(); c$ $\mid v.\text{fork}(); c \mid v.\text{join}(); c \mid \text{unpack}(v.I); c \mid \text{pack}(v.I); c$
$F \in \text{Formula}$	$::= e \mid \text{Perm}(v.f, \pi) \mid \pi.P \mid F \oplus F \mid (qt \ T \ \alpha)F$ $\mid \text{holds}(v.I, \pi) \mid \text{unpacked}(v.I, \pi) \mid e.\text{fresh}() \mid e.\text{initialized}()$
$F_{\text{res}} \in \text{Formula}_{\text{res}}$	$::= e \mid \text{Perm}(v.f, \pi) \mid \pi.P \mid F_{\text{res}} \oplus F_{\text{res}} \mid (qt \ T \ \alpha)(F_{\text{res}}) \mid \text{holds}(v.I, \pi)$
$F_{\text{inv}} \in \text{Formula}_{\text{inv}}$	$::= e_{\text{inv}} \mid (qt \ T \ \alpha)(F_{\text{inv}}) \mid F_{\text{inv}} \oplus F_{\text{inv}}$
$e \in \text{Exp}$	$::= \pi \mid v.f \mid v \mid \text{op}(\overline{e})$
$e_{\text{inv}} \in \text{Exp}_{\text{inv}}$	$::= v_1.v_2..v_n.f \mid \text{op}(\overline{e_{\text{inv}}})$
$T, U, V \in \text{Type}$	$::= \text{void} \mid \text{int} \mid \text{bool} \mid \text{perm} \mid (\text{rtype}, C)$
$\text{rtype} \in \text{RefType}$	$::= \text{rep} \mid \text{peer} \mid \text{self} \mid \text{rd}$
$\pi \in \text{SpecVal}$	$::= \alpha \mid v \mid 1 \mid \text{split}(\pi) \text{ (1/2 of a fractional permission } \pi)$
$u, v, w \in \text{Val}$	$::= \text{null} \mid n \mid b \mid o \mid x$

$\oplus \in \{*, \wedge, \vee\} \quad \text{op} \in \text{Op} \supseteq \{=, !, \wedge, \vee, \Rightarrow\} \quad qt \in \{\exists, \forall\}$   
 $n \in \text{int} \quad b \in \{\text{true}, \text{false}\} \quad x, y, z \in \text{Variables} \quad o, p \in \text{ObjectId}$

Fig. 1. Language Syntax

object, and is used to describe the resources that the lock protects. Methods may be declared as **pure** and/or **safe**, as explained below. The set of commands is extended with the specification commands **pack**( $v.I$ ) and **unpack**( $v.I$ ).

*Specification Formulas.* We distinguish three types of specification formulas: i) *Standard formulas*  $F$ , expressed in permission-based separation logic and used to specify methods. Predicates **holds** and **unpacked**, and **fresh** and **initialized** are special tokens that describe the state of a class invariant or a lock, respectively. ii) *Resource invariant formulas*  $F_{\text{res}}$ , used to express the **res\_inv** predicate. They are more restrictive than  $F$ :  $F_{\text{res}}$  must not use the special tokens **unpacked**, **fresh** and **initialized**.

iii) *State formulas*  $F_{\text{inv}}$ , first-order logic formulas, used to specify class invariants and describe properties over shared memory locations only. Thus, their syntax does not include the predicate  $\text{Perm}(v.f, \pi)$  or any of the special tokens.

We define the invariant footprint  $\text{fp}(v.I)$  by induction of the structure of  $v.I$ :

$$\begin{aligned} \text{fp}(v_1.v_2..v_n.f) &= \{v_1, v_1.v_2, \dots, v_1..v_n.f\} & \text{fp}(\text{op}(\overline{e_{\text{inv}}})) &= \bigcup_{e \in \overline{e_{\text{inv}}}} \text{fp}(e) \\ \text{fp}(F_{\text{inv}_1} \oplus F_{\text{inv}_2}) &= \text{fp}(F_{\text{inv}_1}) \cup \text{fp}(F_{\text{inv}_2}) & \text{fp}((qt \ \alpha \ T)(F_{\text{inv}})) &= \bigcup_{v \in T \setminus \{\alpha\}} \text{fp}(F_{\text{inv}}[v/\alpha]) \end{aligned}$$

*Types.* A type of an object reference in our language is represented as a tuple  $T = (\text{rtype}, C)$ . The first component,  $T^1$ , is a type modifier from the set

$\text{RefType} = \{\text{rep}, \text{peer}, \text{self}, \text{rd}\}$ , while the second,  $T^2$ , represents the object's class. Consequently, two references pointing to the same object might have different reference types if they are in a different context. In this paper we do not present the typing rules of the language; rules that represent constraints imposed by the ownership type system are listed in [25].

*Safe and Pure Commands.* Above, we introduced the notion of **safe** commands. For a **safe** command  $c$  the predicate  $\text{safe}(c, V)$  holds, where  $V$  is a set that keeps track of all identifiers of threads that are forked and expected to be joined. The  $V$  parameter is used to capture that threads forked within a **safe** command  $c$ , must also be joined within  $c$ . For a method  $m$  defined as  $\text{safe } T \ m(\overline{V} \ \bar{i}) \ \{c\}$ , the relation  $\text{safe}(m)$  holds iff  $\text{safe}(c, \emptyset)$  holds. A **safe** method is annotated with the optional modifier **safe**. We define inductively the set of **safe** commands.

$\text{safe}(v, V)$	$\Leftrightarrow \text{true}$
$\text{safe}(c, V)$	$\Leftrightarrow \text{false}, \text{ if } c \in \{v.\text{lock}(), v.\text{unlock}(), v.\text{commit}()\}$
$\text{safe}(c; c_1, V)$	$\Leftrightarrow \text{safe}(c_1, V), \text{ if } c \in \{T \ x, x = v, x = v.f, v.f = v, x = \text{op}(\overline{v}), \text{new rtype } C, \text{unpack}(v.I), \text{pack}(v.I)\}$
$\text{safe}(x = v.m(\overline{v}); c, V)$	$\Leftrightarrow \text{safe}(m) \wedge \text{safe}(c, V)$
$\text{safe}(v.\text{fork}(); c, V)$	$\Leftrightarrow \text{safe}(c, V \cup \{v\})$
$\text{safe}(v.\text{join}(); c, V)$	$\Leftrightarrow \text{safe}(c, V \setminus \{v\})$
$\text{safe}(\text{if } v \text{ then } c_1 \text{ else } c_2; c, V)$	$\Leftrightarrow \text{safe}(c_1, \emptyset) \wedge \text{safe}(c_2, \emptyset) \wedge \text{safe}(c, V)$

Our method uses also the notion of **pure** commands, i.e., commands that do not make any changes to the shared state (defined in [25]). **Pure** methods are composed of pure commands and specified with the optional modifier **pure**.

## 4.2 Hoare Triples

Fig. 2 shows the Hoare triples relevant to our approach (for the complete list of rules see [11]). We use:  $\otimes_i F_i$  to abbreviate a separation conjunction of all formulas  $F_i$ ;  $\text{PointsTo}(v.f, \pi, w)$  to abbreviate  $\text{Perm}(v.f, \pi) \wedge v.f == w$ ; functions  $\text{fld}(C)$  and  $\text{inv}(C)$  to represent respectively the set of fields and invariants in the class  $C$ ;  $\text{df}(T)$  for the default value of type  $T$ ;  $\text{wrt}(F)$  for the set of locations for which  $F$  expresses a write permission (all defined formally in [25]).

The rule (New) shows that construction of object  $v$  produces an **unpacked** token for each invariant of  $v$ , and a write permission for each field of  $v$ . Rules (Set) and (MethCall) encode **R2'** and **R2''** (see Sec. 3.2); they ensure that the breaking token is a condition for breaking the invariant  $v.I$ . Rules (Pack) and (Unpack) describe the invariant protocol and encode **R3** and **R4** (see Sec. 3.1). Finally, the rule (RuleInv) shows that the token  $\text{holds}(v.I, \pi)$  provides the actual thread the right to use the invariant  $v.I$  (as justified by Theorem 2 in Sec 3.1).

## 4.3 Semantics

We define a program state as:  $st \in \text{State} = \text{Heap} \times \text{ThreadPool} \times \text{LockTable}$ . A **Heap** models the shared memory:  $h \in \text{Heap} = \text{ObjId} \mapsto \text{Type} \times (\text{FieldId} \mapsto \text{Value})$ .

(New)	$\frac{\{true\}}{v = \text{new rtype } C}$ $\{\otimes_{T.f \in \text{fld}(C)} \text{PointsTo}(v.f, 1, \text{df}(T^1)) * \otimes_{I \in \text{inv}(C)} \text{unpacked}(v.I, 1)\}$
(Set)	$\frac{v : V}{\{v \neq \text{null} * \text{PointsTo}(v.f, 1, u) * \otimes_{I \in \text{inv}(V^2), v.f \in \text{fp}(v.I)} \text{unpacked}(v.I, \pi)\}}$ $v.f = w;$ $\{\text{PointsTo}(v.f, 1, w) * \otimes_{I \in \text{inv}(V^2), v.f \in \text{fp}(v.I)} \text{unpacked}(v.I, \pi)\}$
(MethCall)	$\frac{\text{md} ::= \text{requires } F \text{ ensures } F' \text{ safe? pure? } T \ m(\overline{U} \ \overline{u})\{c\} \ \text{this} : V}{\{u \neq \text{null} * F * \otimes_{I \in \text{inv}(V^2), \text{wrt}(F) \cap \text{fp}(\text{this}.I) \neq \emptyset} \text{unpacked}(\text{this}.I, \pi)\}}$ $x = u.m(\overline{i})$ $\{\exists T \alpha (\alpha == x * F') * \otimes_{I \in \text{inv}(V^2), \text{wrt}(F) \cap \text{fp}(\text{this}.I) \neq \emptyset} \text{unpacked}(\text{this}.I, \pi)\}$
(Unpack)	$\{\text{holds}(v.I, 1)\} \text{unpack}(v.I) \{\text{unpacked}(v.I, 1) * v.I\}$
(Pack)	$\{\text{unpacked}(v.I, 1) * v.I\} \text{pack}(v.I) \{\text{holds}(v.I, 1)\}$
(RuleInv)	$\frac{\{\text{holds}(v.I, \pi) * v.I\} \ c \ \{F\}}{\{\text{holds}(v.I, \pi)\} \ c \ \{F\}}$

Fig. 2. Hoare triples

The **ThreadPool** component describes all threads that operate on the heap:  $ts \in \text{ThreadPool} = \text{ObjId} \mapsto \text{Thread}$ , where each thread contains its own local memory and a command to execute,  $t \in \text{Thread} = \text{Stack} \times \text{Cmd}$ . The **LockTable** expresses for every lock whether it is free, or it is acquired by a thread a certain number of times:  $l \in \text{LockTable} = \text{ObjId} \mapsto \text{free} \uplus (\text{ObjId} \times \mathbb{N})$ . Operationally, the two specification commands **unpack**( $v.I$ ) and **pack**( $v.I$ ) are no operations. The small-step operational semantics of the other commands is standard, see [11].

*Semantics of Formulas.* The specification formulas are interpreted using the semantics relation  $\Gamma \vdash \mathcal{E}, \mathcal{R}, s \models F$ , which expresses validity of the formula  $F$  in a type environment  $\Gamma$ , a predicate environment  $\mathcal{E}$  and a stack  $s$ , given a resource  $\mathcal{R}$ . Type environment  $\Gamma$  is a partial function of type  $\text{ObjId} \cup \text{Var} \mapsto \text{Type}$  that maps each object or variable to its type, while  $\mathcal{E}$  maps each predicate symbol to an appropriate relation that represents its definition. For details see [11].

The resource  $\mathcal{R}$  is an abstraction of a program state represented by an 8-tuple,  $\mathcal{R} = (h, \mathcal{P}, \mathcal{J}, \mathcal{L}, \mathcal{F}, \mathcal{I}, \mathcal{U}, \mathcal{T})$ , where each component describes part of the state: i)  $h$  represents the heap:  $\text{ObjId} \mapsto \text{Type} \times (\text{FieldId} \mapsto \text{Val})$  ii)  $\mathcal{P}$  is a permission table that stores permissions to object fields from the heap ( $\text{ObjId} \times \text{FieldId} \mapsto [0, 1]$ ); iii)  $\mathcal{J}$  is a join table ( $\text{ObjId} \mapsto [0, 1]$ ), where  $\mathcal{J}(t)$  represents how much of the postcondition of a thread  $t$  is given to other forked threads; iv)  $\mathcal{L}$  is an abstraction of the lock table, which maps each thread to the set of locks that it holds; v)  $\mathcal{F}$  keeps a set of **fresh** locks; vi)  $\mathcal{I}$  keeps a set of **initialized** locks; vii)  $\mathcal{U}$  keeps the parts of the **unpacked** tokens for each invariant; and analogously viii)  $\mathcal{T}$  keeps the **holds** tokens. Both components  $\mathcal{U}$  and  $\mathcal{T}$  are defined as functions  $\text{ObjId} \times \text{InvId} \mapsto [0, 1]$ .

We define a *compatibility* binary relation ( $\#$ ) and a *resource joining operation* ( $*$ ) over resources. Compatibility ensures that two different threads always observe the abstract state as two compatible resources,  $\mathcal{R}\#\mathcal{R}'$ : the object fields that are common for the heaps in  $\mathcal{R}$  and  $\mathcal{R}'$  are mapped to the same value; the sum of permissions for a location in  $\mathcal{R}$  and  $\mathcal{R}'$ , or the sum of the parts of the special tokens (**holds** and **unpacked**) for an invariant in both resources never exceeds 1; etc. The intuitive meaning of the operation  $\mathcal{R} * \mathcal{R}'$  is joining (summing) both resources. For example,  $\mathcal{R} * \mathcal{R}'$  contains all permissions from both resources or all tokens from both resources. The definition of the  $\#$  and  $*$  is component-wise. We give the formal definitions for the structure  $(\#, *)$  for the components  $\mathcal{U}$  and  $\mathcal{T}$ , while for the others we refer to [11].

$$\begin{aligned} \mathcal{U}\#\mathcal{U}' &\Leftrightarrow \forall i \in \text{dom}(\mathcal{U}) \cap \text{dom}(\mathcal{U}'). \mathcal{U}(i) + \mathcal{U}'(i) \leq 1 & (\mathcal{U} * \mathcal{U}')(i) &= \mathcal{U}(i) + \mathcal{U}'(i) \\ \mathcal{T}\#\mathcal{T}' &\Leftrightarrow \forall i \in \text{dom}(\mathcal{T}) \cap \text{dom}(\mathcal{T}'). \mathcal{T}(i) + \mathcal{T}'(i) \leq 1 & (\mathcal{T} * \mathcal{T}')(i) &= \mathcal{T}(i) + \mathcal{T}'(i) \end{aligned}$$

Below we define that the specification formula  $\text{holds}(v.I, \pi)$  holds for a resource  $\mathcal{R}$  if the part of the **holds** token for the invariant  $v.I$  in  $\mathcal{R}$  is at least  $\pi$ . The validity of the  $\text{unpacked}(v.I, \pi)$  formula is defined analogously. The semantics of a class invariant  $v.I$  is expressed as a validity of the representation formula of  $v.I$ , i.e.,  $F_{\text{inv}}$ .

$$\begin{aligned} \Gamma \vdash \mathcal{E}, (h, \mathcal{P}, \mathcal{J}, \mathcal{L}, \mathcal{F}, \mathcal{I}, \mathcal{U}, \mathcal{T}), s &\models \text{holds}(v.I, \pi) \Leftrightarrow \mathcal{T}(v.I) \geq \pi \\ \Gamma \vdash \mathcal{E}, (h, \mathcal{P}, \mathcal{J}, \mathcal{L}, \mathcal{F}, \mathcal{I}, \mathcal{U}, \mathcal{T}), s &\models \text{unpacked}(v.I, \pi) \Leftrightarrow \mathcal{U}(v.I) \geq \pi \\ \Gamma \vdash \mathcal{R} = \mathcal{E}, (h, \mathcal{P}, \mathcal{J}, \mathcal{L}, \mathcal{F}, \mathcal{I}, \mathcal{U}, \mathcal{T}), s &\models v.I(I = F_{\text{inv}}) \Leftrightarrow \Gamma \vdash \mathcal{E}, \mathcal{R}, s \models F_{\text{inv}} \end{aligned}$$

As our language contains *state formulas*, not all locations in the partial heap must be 'framed' by a positive permission (unlike in standard permission-based separation logic). For a sound resource  $\mathcal{R} = (h, \mathcal{P}, \mathcal{J}, \mathcal{L}, \mathcal{F}, \mathcal{I}, \mathcal{U}, \mathcal{T})$  we require:

$$\begin{aligned} \forall p \in \text{dom}(h), f \in \text{dom}(h(p)_2), \mathcal{P}(p, f) > 0 \vee \\ (\exists v.I \in \text{dom}(\mathcal{T}) \ p.f \in \text{fp}(v.I) \wedge (\mathcal{T}(v.I) > 0 \vee \mathcal{U}(v.I) > 0)) \end{aligned}$$

The rule states that if a location  $p.f$  is not protected by a read permission ( $\mathcal{P}(p, f) = 0$ ), then it must be protected by (a part of) the **holds** or **unpacked** token ( $\mathcal{T}(v.I) > 0 \vee \mathcal{U}(v.I) > 0$ ), for an invariant  $v.I$  that refers to  $p.f$ . This ensures that the location  $p.f$  is stable and might not be modified by other threads.

## 5 Related Work

The early work on verification of class invariants in sequential programs [17,15] is unsound for more complex data structure, for example if an invariant captures properties over different objects. Later, Poetzsch-Heffter [22] and Huizing *et al.* [13] presented sound techniques that do not restrict the invariant definition or the program itself; however, both approaches are not modular.

Müller *et al.* [18] propose two sound techniques for modular reasoning: the *ownership technique* and the less restrictive *visibility technique*. Both concepts, as well as Lu *et al.*'s modular technique [16], are designed for ownership-based

type systems. These techniques are captured in Drossopoulou *et al.*'s abstract unified framework [9]. Although it is stated that this abstract framework should be suitable to model class invariants in a concurrent setting, the framework has never been applied on a concrete verification technique for concurrent programs.

Weiß models class invariants with a boolean model field *inv* [24]. Their validity is checked only on demand. Specifications use *inv* explicitly where needed, while *this.inv* is implicitly generated in each method pre- and postcondition.

We are not aware of much work done on verification of class invariants for multithreaded programs. Comparable to our approach is Jacobs *et al.*'s technique [14] for verifying multithreaded programs with class invariants, using the *Boogie methodology* [3] for sequential programs. However, this technique allows a thread to break an invariant of an object only if it completely owns this object. Instead, with our technique, breaking a class invariant is independent of permissions on heap memory. This ensures a broader applicability of our technique.

A different approach for modular verification of object invariants in concurrent programs is proposed by Cohen [6], implemented in VCC [5]. Each object is assigned a two-state invariant expressing the required relation between any two consecutive states of execution that has to be respected by every state update in the program. Modular verification of multithreaded programs with class invariants is also supported by the static checker Calvin [10]. However, both methodologies do not allow breaking of a class invariant in the program.

## 6 Conclusion and Future Work

We introduced a sound and modular approach for verifying class invariants in multithreaded Java-like programs in a permission-based separation logic setting. We do, however, deviate from the standard rules in separation logic: we impose that class invariants may express properties only over state and thus, their definition is free of permission expressions. We allow a thread to explicitly break an invariant, and we ensure that no other thread can observe the invalidated object's state. Moreover, breaking and reestablishing an invariant is allowed without holding all permissions associated to the invariant. This makes our technique broadly applicable. To achieve modularity, we restrict our technique to ownership-based type systems only. The method requires simple specifications support.

For future work, we plan to integrate our technique in the VerCors tool [1], and to use it to verify data structures from the *java.util.concurrent* package. We plan to extend the concept to support class inheritance, to allow more permissive invariants with model methods and/or abstract predicates, to allow more fine-grained permission handling, as well as to support *history constraints*.

**Acknowledgments.** We thank Christian Haack and Stefan Blom for their useful feedback. This work was supported by ERC grant 258405 for the VerCors project.

## References

1. Amighi, A., Blom, S., Huisman, M., Zaharieva-Stojanovski, M.: The VerCors project: setting up basecamp. In: PLPV, pp. 71–82 (2012)
2. Artho, C., Havelund, K., Biere, A.: High-level data races. *Softw. Test., Verif. Reliab.* 13(4), 207–227 (2003)
3. Barnett, M., DeLine, R., Fähndrich, M., Leino, K.R.M., Schulte, W.: Verification of object-oriented programs with invariants. *Journal of Object Technology* 3(6), 27–56 (2004)
4. Bornat, R., Calcagno, C., O’Hearn, P., Parkinson, M.: Permission accounting in separation logic. In: Palsberg, J., Abadi, M. (eds.) POPL, pp. 259–270. ACM (2005)
5. Cohen, E., Dahlweid, M., Hillebrand, M., Leinenbach, D., Moskal, M., Santen, T., Schulte, W., Tobies, S.: VCC: A practical system for verifying concurrent C. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 23–42. Springer, Heidelberg (2009)
6. Cohen, E., Moskal, M., Schulte, W., Tobies, S.: Local verification of global invariants in concurrent programs. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 480–494. Springer, Heidelberg (2010)
7. Dietl, W., Müller, P.: Universes: Lightweight ownership for JML. *Journal of Object Technology* 4(8), 5–32 (2005)
8. Dietl, W., Müller, P.: Object ownership in program verification. In: Clarke, D., Noble, J., Wrigstad, T. (eds.) *Aliasing in Object-Oriented Programming*. LNCS, vol. 7850, pp. 289–318. Springer, Heidelberg (2013)
9. Drossopoulou, S., Francalanza, A., Müller, P., Summers, A.J.: A unified framework for verification techniques for object invariants. In: Vitek, J. (ed.) ECOOP 2008. LNCS, vol. 5142, pp. 412–437. Springer, Heidelberg (2008)
10. Flanagan, C., Freund, S.N., Qadeer, S., Seshia, S.A.: Modular verification of multithreaded programs. *Theor. Comput. Sci.* 338(1–3), 153–183 (2005)
11. Haack, C., Huisman, M., Hurlin, C., Amighi, A.: Permission-based separation logic for Java, 201x. Conditionally accepted for LMCS
12. Hoare, C.A.R.: An axiomatic basis for computer programming. *Commun. ACM* 12(10), 576–580 (1969)
13. Huizing, K., Kuiper, R.: Verification of object oriented programs using class invariants. In: Maibaum, T. (ed.) FASE 2000. LNCS, vol. 1783, pp. 208–221. Springer, Heidelberg (2000)
14. Jacobs, B., Piessens, F., Leino, K.R.M., Schulte, W.: Safe concurrency for aggregate objects with invariants. In: SEFM, pp. 137–147 (2005)
15. Liskov, B., Guttag, J.: Abstraction and specification in program development. MIT Press, Cambridge (1986)
16. Lu, Y., Xue, J.: Validity invariants and effects. In: Ernst, E. (ed.) ECOOP 2007. LNCS, vol. 4609, pp. 202–226. Springer, Heidelberg (2007)
17. Meyer, B.: *Object-Oriented Software Construction*, 2nd edn. Prentice-Hall (1997)
18. Müller, P., Poetzsch-Heffter, A., Leavens, G.T.: Modular invariants for layered object structures. *Sci. Comput. Program.* 62(3), 253–286 (2006)
19. O’Hearn, P.W.: Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* 375(1–3), 271–307 (2007)
20. Parkinson, M., Bierman, G.: Separation logic, abstraction and inheritance. In: *Principles of programming languages (POPL 2008)*, pp. 75–86. ACM (2008)
21. Parkinson, M.J.: Local reasoning for Java. Technical Report UCAM-CL-TR-654, University of Cambridge, Computer Laboratory (November 2005)

22. Poetzsch-Heffter, A.: Specification and Verification of Object-Oriented Programs. PhD thesis, Habilitation thesis, Technical University of Munich (1997)
23. Reynolds, J.: Separation logic: A logic for shared mutable data structures. In: 17th IEEE Symposium on LICS 2002, pp. 55–74. IEEE Computer Society (2002)
24. Weiß, B.: Deductive Verification of Object-Oriented Software: Dynamic Frames, Dynamic Logic and Predicate Abstraction. PhD thesis, Karlsruhe Institute of Technology (2011)
25. Zaharieva-Stojanovski, M., Huisman, M.: Verifying class invariants in concurrent programs. Technical Report TR-CTIT-13-10, Centre for Telematics and Information Technology, University of Twente (2014)