Separation Logic Contracts for a Java-like Language with Fork/Join

Christian Haack^{1*} and Clément Hurlin^{2**}

¹ Radboud Universiteit Nijmegen, The Netherlands

² INRIA Sophia Antipolis - Méditerranée, France

Abstract. We adapt a variant of permission-accounting separation logic to a concurrent Java-like language with fork/join. To support both concurrent reads and information hiding, we combine fractional permissions with abstract predicates. As an example, we present a separation logic contract for iterators that prevents data races and concurrent modifications. Our program logic is presented in an algorithmic style: we avoid structural rules for Hoare triples and formalize logical reasoning about typed heaps by natural deduction rules and a set of sound axioms. We show that verified programs satisfy the following properties: data race freedom, absence of null-dereferences and partial correctness.

1 Introduction

1.1 Context

Over the past ten years or so, substructural logics and type systems have proven to be very valuable formalisms for reasoning about pointer-manipulating programs. Examples include static capabilities [10,11], alias types [29] and separation logic [18,28]. In these systems, the underlying specification language contains linear formulas for specifying memory access policies. Whereas traditional program logics control memory access via frame conditions, separation logic tightly integrates access policy specifications into the formula language itself. Formulas represent access tickets to heap space, and possession of access tickets gets verified statically. Access policies are tightly coupled with assertions about memory content, so that separation logic's Hoare rules make it impossible to maintain assertions that can be invalidated by thread interference or memory updates through unknown aliases. This is achieved without annoying side conditions like non-interference tests or frame conditions.

While initially separation logic mostly focused on low level programs, researchers have more recently started to adapt it to object-oriented features for use in contract languages for OO [25,26], and very recently [9,27].

1.2 Contributions

We present the careful design of a small Java-like model language with separation logic contracts, including the definition of a program logic and its soundness proof. Our language has simple threads, with fork/join as concurrency primitives. In order to facilitate concurrent reads we employ fractional permissions [5]. Our rules allow multiple

* Supported in part by IST-FET-2005-015905 Mobius project.

** Supported in part by IST-FET-2005-015905 Mobius and ANR-06-SETIN-010 ParSec project.

threads to join on the same thread, in order to read-share the dead thread’s resources. This is not possible with a lexically scoped parallel composition operator or with Posix threads, and is thus not supported by recent work that adapts separation logic to Posix threads [14]. To support data abstraction and recursive data types, we use abstract predicates [26]. *Class axioms* complement abstract predicates to export relations between predicates without revealing their full definitions. Abstract predicates satisfying a split/merge axiom generalize datagroups [21], which are common in specification languages for OO. In order to support concurrent read access to whole datagroups (rather than single fields), access permission to datagroups can be split by splitting their permission parameters. In order to allow fine-grained permission-splitting for overlapping datagroups, we support datagroups with multiple permission parameters. To achieve modular soundness in the presence of subclassing, we axiomatize the “stack of class frames” [12,1] in separation logic. We support *value-parametrized classes*, where class parameters have the same purpose as `final` ghost fields in specification languages like JML [20]. In particular, class parameters can represent static ownership relations.

1.3 Background on Separation Logic and Fractional Permissions

Separation logic combines the usual logical operators with the points-to predicate $x.f \mapsto v$, the resource conjunction $F * G$, and the resource implication $F \multimap G$.

The predicate $x.f \mapsto v$ has a *dual purpose*: firstly, it asserts that the object field $x.f$ contains data value v and, secondly, it represents a *ticket* that grants permission to access the field $x.f$. This is formalized by separation logic’s Hoare rules for reading and writing fields:

$$\{x.f \mapsto _ * F\} x.f = v \{x.f \mapsto v * F\} \quad \{x.f \mapsto v * F\} y = x.f \{x.f \mapsto v * v == y * F\}$$

The crucial difference to standard Hoare logic is that both these rules have a precondition of the form $x.f \mapsto _$.³ This formula functions as an *access ticket* for $x.f$.

It is important that tickets are not forgeable. One ticket is not the same as two tickets! For this reason, the resource conjunction $*$ is not idempotent: F is not equivalent to $F * F$. The resource implication $\multimap$ matches the resource conjunction $*$, in the sense that the modus ponens law is satisfied: $F * (F \multimap G)$ implies G . However, $F * (F \multimap G)$ *does not* imply $F * G$. In English, $F \multimap G$ is pronounced as “consume F yielding G ”. In terms of tickets, $F \multimap G$ permits to trade ticket F and receive ticket G in return.

Separation logic is particularly useful for concurrent programs: two concurrent threads simply split the resources that they may access, as formalized by the rule for the parallel composition $t \mid t'$ of threads t and t' [22].

$$\frac{\{F\}t\{G\} \quad \{F'\}t'\{G'\}}{\{F * F'\}t \mid t'\{G * G'\}}$$

With this concurrency rule, separation logic prevents data races. There is a caveat, though. The rule does not allow concurrent reads. Boyland [5] solved this problem with a very intuitive idea, which was later adapted to separation logic [4]. The idea is that (1) *access tickets are splittable*, (2) *a split of an access ticket still grants read access* and (3) *only a whole access ticket grants write access*. To account for multiple

³ $x.f \mapsto _$ is short for $(\exists v)(x.f \mapsto v)$.

Fig. 1. Usage Protocol for Iterators

splits, Boyland uses fractions, hence the name *fractional permissions*. In permission-accounting separation logic [4], access tickets $x.f \mapsto v$ are superscripted by fractions π . $x.f \mapsto^\pi v$ is equivalent to $x.f \mapsto^{\pi/2} v * x.f \mapsto^{\pi/2} v$. In the Hoare rules, writing requires the full fraction 1, whereas reading just requires *some* fraction π :

$$\{x.f \mapsto^1 _ * F\} x.f = v \{x.f \mapsto^1 v * F\} \quad \{x.f \mapsto^\pi v * F\} y = x.f \{x.f \mapsto^\pi v * v == y * F\}$$

Permission-accounting separation logic maintains the global invariant that the sum of all fractional permissions to the same cell is always at most 1. This prevents read-write and write-write conflicts, but permits concurrent reads.

In our Java-like language, we use ASCII and write $\text{Perm}(x.f, \pi)$ for $x.f \mapsto^\pi _$, and $\text{PointsTo}(x.f, \pi, v)$ for $x.f \mapsto^\pi v$.

1.4 Example: A Usage Protocol for Iterators

Often one wants to constrain object clients to adhere to certain usage protocols. Object usage protocols can, for instance, be specified in tpestate systems [12] or, using ghost fields, in general purpose specification languages. A limitation of these techniques is that state transitions must always be associated with method calls. This is sometimes not sufficient. Consider for instance a variant of Java’s `Iterator` interface (enriched with an `init` method to avoid constructor contracts):

```

interface Iterator {
    void init(Collection c);
    boolean hasNext();
    Object next();
    void remove();
}

```

If iterators are used in an undisciplined way, there is the danger of unwanted concurrent modification of the underlying collection (both of the collection elements and the collection itself). Moreover, in concurrent programs bad iterator usage can result in data races. It is therefore important that `Iterator` clients adhere to a usage discipline.

Figure 1 shows a state machine that defines a safe iterator usage discipline. Unfortunately, the dashed transitions are not supported by existing tpestate systems, because they are not associated with method calls. Specifying

this protocol with classical program logics would be clumsy. In [13] (Section 1.1.4), Girard explains how *linear implications can be used to logically represent state transitions*. Applying this idea to the iterator protocol, we obtain the following formalization where the dashed transitions are represented by resource implications:

```
interface Iterator<perm p, Collection iteratee> {
  pred ready; // prestate for iteration cycle
  pred readyForNext; // prestate for next()
  pred readyForRemove<Object element>; // prestate for remove()

  axiom ready -* iteratee.state<p>; // stop iterating
  axiom readyForRemove<e> * e.state<p> -* ready; // back to ready

  req init * c.state<p> * c==iteratee; ens ready;
  void init(Collection c);

  req ready; ens ready & (result -* readyForNext);
  boolean hasNext();

  req readyForNext; ens result.state<p> * readyForRemove<result>;
  Object next();

  req readyForRemove<_> * p==1; ens ready;
  void remove();
}
```

The interface has two parameters: firstly, a permission p and, secondly, the *iteratee*. If the permission parameter is instantiated by a fraction $p < 1$, one obtains a read-only iterator, otherwise a read-write iterator. The states of our state diagram are represented by three abstract predicates: *ready*, *readyForNext* and *readyForRemove*.

Class axioms express relations between abstract predicates, without revealing the complete predicate definitions. Implementations must define the abstract predicates by separation logic formulas such that the class axioms are tautologically true. In the example, the two class axioms represent the dashed transitions of the state machine. We represent the heap space associated with an object by a generic datagroup *state*, which has a default definition in the *Object* class and needs to be overridden by each class. The definition of this *state* predicate should describe the heap space associated with the object. Often this heap space will consist of the object’s fields only, but sometimes it will also include other objects and change dynamically, as in the case of *Collection* objects. The *state* predicate is parametrized by a fraction so that it can be read-shared.

The precondition of *init()* consumes a fraction p of the access right for the *iteratee* and puts the iterator in the *ready* state. The crux is that, by linearity, the iterator client temporarily loses a p -fraction of the access right on the collection, which he can only gain back by “invoking” the first class axiom. The *init* predicate in *init()*’s precondition is a special abstract predicate that every object enters right after object creation and that grants access to all of the object’s fields.

The postcondition of *hasNext()* uses a resource implication whose antecedent is a boolean expression. We treat boolean expressions as *copyable resources* that satisfy $e \multimap (e * e)$.⁴ Furthermore, *hasNext()*’s postcondition uses *additive conjunction* $\&$. A

⁴ We could equivalently use classical implication: $\text{result} \Rightarrow \text{readyForNext}$. Even the two class axioms could equivalently use classical implication, because they are tautologies, and in

resource satisfies $F \& G$, if it satisfies both F and G .⁵ Operationally, $\&$ represents *choice*. If $F \& G$ holds, then F and G are available, but are interdependent: using either one of them destroys the other one, too. Additive conjunction can conveniently represent non-deterministic state transitions, as exhibited in `hasNext()`'s postcondition. Note that this postcondition allows clients to stay in the ready-state, even if `hasNext()==true`. This can, for instance, be useful for removing the 10th element of an ordered collection.

In our companion report [16], we have implemented the `Iterator` interface for a doubly linked list implementation of the `Collection` interface. In [15], we refine the protocol to support unrestricted access to immutable collection elements, and to support shallow collections that do not govern access to their elements.

1.5 Example: Representing Datagroups

We represent *datagroups* [21] as abstract predicates satisfying a datagroup axiom that says *split/merging datagroup parameters split/merges datagroups*:

$$\text{group } P\langle \bar{T} \bar{x} \rangle; \triangleq \text{pred } P\langle \bar{T} \bar{x} \rangle; \text{ axiom } P\langle \bar{x} \rangle *-* (P\langle \bar{e} \rangle * P\langle \bar{e} \rangle);$$

where $e_i \triangleq x_i/2$, if $T_i = \text{perm}$, and $e_i \triangleq x_i$, otherwise

The formula $F *-* G$ is short for $(F -* G) \& (G -* F)$. Here are simple examples of a legal and an illegal datagroup definition (where $|$ is disjunction):

```
group P<perm p> = Perm(this.f,p) * Perm(this.g,p);
  legal because the datagroup axiom holds
group P<perm p> = Perm(this.f,p) | Perm(this.g,p);
  illegal because the datagroup axiom's right-to-left direction does not hold
```

On the right, you see a fractional permission version of Leino's running example [21]. The datagroups `position` and `color` are nested in `state`, as expressed by the two class axioms. The formula " $F \text{ ispartof } G$ " is a derived form for $G -* (F * (F -* G))$. Intuitively, this formula says that F is a physical part of G : one can take G apart into F and its complement $F -* G$, and one can put the two parts back together to obtain G back.

```
interface Sprite {
  group position<perm p>;
  group color<perm p>;
  axiom position<p> ispartof state<p>;
  axiom color<p> ispartof state<p>;
  req position<1>; ens position<1>;
  void updatePosition();
  req color<1>; ens color<1>;
  void updateColor();
  req state<1>; ens state<1>;
  void update();
  req state<p>; ens state<p>;
  void display();
}
```

intuitionistic separation logic $F -* G$ is a tautology if and only if $F \Rightarrow G$ is. However, in our implementations of the `Iterator` interface, we use true resource implications that cannot be replaced by classical implications. In this paper, we avoid classical implication because having just one implication simplifies the natural deduction rules for reasoning about resources.

⁵ In contrast, a resource satisfies $F * G$ if it can be split into separate resources, one of which satisfies F and the other satisfies G .

1.6 Example: Recursive and Overlapping Datagroups

Our next example illustrates that multiple threads can concurrently access overlapping datagroups, as long as they only read-access their intersection. Consider a linked list that implements a simple class roster. Each node stores a student identifier and a grade. We design the roster interface so that multiple threads can concurrently read the roster. Moreover, when a thread updates the grades we allow other threads to concurrently read the student identifiers. To this end, the interface defines two datagroups `ids_and_links<p,q>` and `grades_and_links<p,q>` that overlap in the links of the list. The permission parameter `p` is associated with the student id fields and grade fields, respectively. The permission parameter `q` is associated with the links.

```
interface Roster {
  group ids_and_links<perm p, perm q>;
  group grades_and_links<perm p, perm q>;
  axiom state<p> == (ids_and_links<p,p/2> * grades_and_links<p,p/2>);
  req grades_and_links<1,p> * ids_and_links<q,r>;
  ens grades_and_links<1,p> * ids_and_links<q,r>;
  void updateGrade(int id, int grade);
  req ids_and_links<p,q>; ens ids_and_links<p,q>;
  bool contains(int id);
}
```

The `updateGrade()` method requires write access (permission 1) for the grades and read access for the links and ids. The `contains()` method requires read permission for the ids and the links. The axiom exposes that the `state` datagroup is the union of the datagroups `ids_and_links` and `grades_and_links` and that these datagroups overlap on the links. In our companion report [16], we have implemented this interface.

2 A Model Language with Separation Logic Contacts

2.1 Syntax

We distinguish between read-only variables ι , read-write variables ℓ , and logic variables α . Method parameters (including `this`) are read-only. Logic variables can only occur in specifications and types. They range over both fractional permissions and values (like integers, object identifiers and `null`).

$$\begin{aligned} C, D \in \text{ClassId} \quad I \in \text{InterId} \quad s, t \in \text{TyId} &= \text{ClassId} \cup \text{InterId} \quad o, p, q \in \text{ObjId} \quad f \in \text{FieldId} \\ m \in \text{MethId} \quad P \in \text{PredId} \quad \iota \in \text{RdVar} \quad \ell \in \text{RdWrVar} \quad \alpha \in \text{LogVar} \\ x, y, z \in \text{Var} &= \text{RdVar} \cup \text{RdWrVar} \cup \text{LogVar} \end{aligned}$$

We include read-only variables (but not read-write variables) in the syntax domain of *values*. This is convenient for our substitution-based operational semantics. *Fractional permissions* are represented symbolically: `splitn(1)` represents the concrete fraction $\frac{1}{2^n}$. In examples, we sometimes write $\frac{1}{2^n}$ as syntax sugar for `splitn(1)`. *Specification values* union values and fractional permissions. Interfaces and classes are parametrized by specification values. Correspondingly, object types $t < \bar{\pi} >$ instantiate the parameters.

$n \in \text{Int}$	integers	$b \in \text{Bool} = \{\text{true}, \text{false}\}$	booleans
$u, v, w \in \text{Val} ::= \text{null} \mid n \mid b \mid o \mid \iota$			values
$\pi \in \text{SpecVal} ::= v \mid 1 \mid \text{split}(\pi) \mid \alpha$			specification values
$T, U, V, W \in \text{Ty} ::= \text{void} \mid \text{int} \mid \text{bool} \mid \iota < \bar{\pi} > \mid \text{perm}$			types

Interface Declarations:

$F \in \text{Formula} ::= \dots$	specification formulas (defined in Section 2.3)
$\text{spec} ::= \text{req } F; \text{ens } F;$	pre- and postconditions
$\text{mt} ::= < \bar{T} \bar{\alpha} > \text{spec } U \text{ m } (\bar{V} \bar{\iota})$	method types (scope of $\bar{\alpha}, \bar{\iota}$ is $\bar{T}, \text{spec}, U, \bar{V}$)
$\text{pt} ::= \text{pred } P < \bar{T} \bar{\alpha} >$	predicate types
$\text{ax} ::= \text{axiom } F$	class axioms
$\text{int} \in \text{Interface} ::= \text{interface } I < \bar{T} \bar{\alpha} > \text{ext } \bar{U} \{ \text{pt}^* \text{ ax}^* \text{ mt}^* \}$	interfaces (scope of $\bar{\alpha}$ is $\bar{T}, \bar{U}, \text{pt}^*, \text{ax}^*, \text{mt}^*$)

Syntactic restriction: The type “perm” may only occur inside angle brackets or formulas.

Method types include pre- and postconditions and are parametrized by logic variables. In examples, we often leave these quantifiers over logic variables implicit. Interfaces may declare *abstract predicates* and classes must implement them by providing concrete definitions as separation logic formulas. Like [26], we allow abstract predicates to have parameters in addition to the implicit self-parameter (as listed in the typed formal parameter lists $\bar{T} \bar{\alpha}$). The types $\bar{T}$ for predicate parameters range over all Java types and the distinguished type *perm* for fractional permissions.

We assume that the *Object* class declares a distinguished datagroup called *state*:

```
class Object { group state<perm p> = true; }
```

This datagroup represents the access permissions for the *object state*. Every class must extend it and thereby define what the object states of its instances are. Our syntax for predicate extensions is as follows:

```
class C ext D { ... ext pred P< $\bar{T} \bar{\alpha}$ > by F; ... }
```

Semantically, the extension F of abstract predicate P gets **-conjoined* with P ’s definition in C ’s superclass D . We do not allow arbitrary predicate redefinitions in subclasses in order to facilitate modular verification, avoiding re-verification of inherited methods.

Class Declarations:

$\text{fd} ::= T f$	field declarations
$\text{pd} ::=$	predicate definitions
$\text{final? pred } P < \bar{T} \bar{\alpha} > = F$	root definition (scope of $\bar{\alpha}$ is F)
$\text{final? ext pred } P < \bar{T} \bar{\alpha} > \text{ by } F$	extension (scope of $\bar{\alpha}$ is F)
$\text{md} ::= \text{final? } < \bar{T} \bar{\alpha} > \text{spec } U \text{ m } (\bar{V} \bar{\iota}) \{ c \}$	method (scope of $\bar{\alpha}, \bar{\iota}$ is $\bar{T}, \text{spec}, U, \bar{V}, c$)
$\text{cl} \in \text{Class} ::= \text{final? class } C < \bar{T} \bar{\alpha} > \text{ext } U \text{ impl } \bar{V} \{ \text{fd}^* \text{ pd}^* \text{ ax}^* \text{ md}^* \}$	class (scope of $\bar{\alpha}$ is $\bar{T}, U, \bar{V}, \text{fd}^*, \text{pd}^*, \text{ax}^*, \text{md}^*$)
$\text{ct} \subseteq \text{Interface} \cup \text{Class}$	class tables

Syntactic restrictions:

- The type “perm” may only occur inside angle brackets or specification formulas.
- Cyclic predicate definitions in ct must be positive.

The *first syntactic restriction* ensures that fractional permissions do not spill into the executable part of the language. The *second syntactic restriction* ensures that predicate implementations (which can be recursive) are well-founded. We allow negative dependencies of predicate P on predicate Q as long as Q does not also depend on P .

We use the symbol $\preceq_{ct}$ for the partial order on type identifiers induced by class table ct , usually leaving the subscript ct implicit. We write $s \prec_1 t$ when s and t are neighbours with respect to $\preceq$. *Subtyping* is inductively defined by the following rules:

$$\begin{aligned} T <: T \quad T <: U, U <: V \Rightarrow T <: V \quad s < \bar{T} \bar{\alpha} > \text{ext } t < \bar{\pi}' > \Rightarrow s < \bar{\pi} > <: t < \bar{\pi}' [\bar{\pi} / \bar{\alpha}] > \\ t < \bar{\pi} > <: \text{Object} \quad t < \bar{T} \bar{\alpha} > \text{impl } I < \bar{\pi}' > \Rightarrow t < \bar{\pi} > <: I < \bar{\pi}' [\bar{\pi} / \bar{\alpha}] > \end{aligned}$$

We assume that class tables always contain the following class declaration:

```
class Thread ext Object {
  final void fork(); final void join();
  req false; ens true; void run() { null }
}
```

The `run`-method is meant to be overridden. The contracts for `fork` and `join` are omitted, because our verification system ignores them anyway. Instead, it uses the precondition for `run` as the precondition for `fork` and the postcondition for `run` as the postcondition for `join`. The methods `fork` and `join` do not have implementations, but the operational semantics treats them in a special way⁶: `o.fork()` creates a new thread, whose thread identifier is o , and executes `o.run()` in this thread. The `o.fork`-method should not be called more than once (on the same receiver o). A second call results in blocking. `o.join()` blocks until thread o has terminated.

Commands:

$op \in \text{Op} \supseteq \{=, !, \&, \} \cup \{C \text{ isclassof} \mid C \in \text{ClassId}\}$	
$c \in \text{Cmd} ::=$	commands
v	return value (or null in case of type void)
$T \ell; c$	local variable declaration (scope of ℓ is c)
$\text{final } T \iota = \ell; c$	local read-only variable declaration (scope of ι is c)
$\text{unpack (ex } T \alpha) (F); c$	unpacking an existential (scope of α is F, c)
$hc; c$	first do hc , then do c
$hc \in \text{HeadCmd} ::= \ell = v \mid \ell = op(\bar{v}) \mid \ell = v.f \mid v.f = v \mid \ell = (T)v \mid \ell = \text{new } C < \bar{\pi} > \mid$ $\text{if } (v) \{c\} \text{else} \{c'\} \mid \ell = v.m < \bar{\pi} > (\bar{v}) \mid \text{assert}(F)$	
<i>Synt. Restr.:</i> Logic variables that occur in $\ell = \text{new } C < \bar{\pi} >$ must be bound by class parameters.	

Our command language assumes that Java-like commands have been transformed so that intermediate values are always assigned to local variables. Following [17], we assume that methods only return at the end of their body. We omit the `return`-command. Values are included in the syntax domain of commands, so that a terminating, non-blocking execution of a command results in the return value. Methods of type void return null, which is the only member of type void. We usually omit terminating occurrences of null. The operator for *existential unpacking* has no effect at runtime. It makes the existential variable α available in the continuation c for instantiation of logic method parameters. In examples, we often omit explicit existential unpacking and instantiation of logic method parameters. Making these explicit helps with the theory.

⁶ In reality, they would be implemented natively.

2.2 Operational Semantics

Runtime Structures:

$\text{CIVal} = \text{Val} \setminus \text{RdVar}$	closed values
$s \in \text{Stack} = \text{RdWrVar} \rightarrow \text{CIVal}$	stacks
$t \in \text{Thread} = \text{Stack} \times \text{Cmd} ::= s \text{ in } c$	threads
$ts \in \text{ThreadPool} = \text{ObjId} \rightarrow \text{Thread} ::= o_1 \text{ is } t_1 \mid \dots \mid o_n \text{ is } t_n$	thread pools
$os \in \text{ObjStore} = \text{FieldId} \rightarrow \text{CIVal}$	object stores
$obj \in \text{Obj} = \text{Ty} \times \text{ObjStore} ::= (T, os)$	objects
$h \in \text{Heap} = \text{ObjId} \rightarrow \text{Obj}$	heaps
$st \in \text{State} = \text{Heap} \times \text{ThreadPool} ::= \langle h, ts \rangle$	states
$prog \in \text{Program} = \text{ClassTable} \times \text{Cmd} ::= (ct, c)$	programs

Each thread “ s in c ” consists of a thread-local stack s and a process continuation c . In thread pools, each thread t is associated with a unique object identifier, which serves as a thread identifier. The dynamic semantics of our language is a small-step operational semantics $st \rightarrow_{ct} st'$ and can be found in [16].

There is one (and only one) reduction rule where our operational semantics depends on class parameters, namely the reduction rule for type casts. Downcasts to parametrized types require a runtime check that looks at the type parameters, which the standard JVM does not keep track of. There are at least three ways how one could deal with that in practice: Firstly (and most pragmatically), one could simply forbid downcasts to reference types that have a non-empty parameter list. Secondly, one could develop an enhanced virtual machine that keeps track of class parameters. Thirdly, one could devise a syntactic translation that erases class parameters such that the target of this translation throws a `ClassCastException` whenever the source does.

2.3 Specification Formulas and Their Semantics

Specification Formulas:

$e \in \text{Exp} ::= \pi \mid \ell \mid op(\bar{e})$	$lop \in \{*, -*, \&, \}$	$qt \in \{\text{ex}, \text{fa}\}$
$\kappa \in \text{Pred} ::=$	predicates	
P	P at receiver's dynamic class	
$P@C$	P at class C	
$E, F, G, H \in \text{Formula} ::=$	specification formulas	
e	boolean expression	
$\text{PointsTo}(e.f, \pi, e')$	$e.f$ points to e' and the access permission for $e.f$ is π	
$\text{Perm}(e.\text{join}, \pi)$	permission to use a split of join 's postcondition	
$\pi.\kappa < \bar{\pi}' >$	predicate $\pi.\kappa$ applied to $\bar{\pi}'$	
$F \text{ lop } G$	binary logical operator	
$(qt \ T \ \alpha) (F)$	quantifier	
Derived forms: $F * - * G \triangleq (F - * G) \& (G - * F)$ $F \text{ assures } G \triangleq F - * (F * G)$		
$F \text{ ispartof } G \triangleq G - * (F * (F - * G))$		

The formula semantics is defined by a Kripke resource interpretation [23] of the form $\Gamma \vdash \mathcal{E}; \mathcal{R}; s \models F$, where Γ is a type environment, $\mathcal{E}$ is a predicate environment that maps predicate names to predicates, $\mathcal{R}$ is a resource, and s is a stack. Resources are triples $\mathcal{R} = (h, \mathcal{P}, \mathcal{Q})$ of heaps h and two permission tables $\mathcal{P}$ and $\mathcal{Q}$. Permission

tables are functions of type $\text{ObjId} \times (\text{FieldId} \times \{\text{join}\}) \rightarrow [0, 1]$ that map fields and join to fractional permissions. The resource components h and $\mathcal{P}$ are local resources, whereas $\mathcal{Q}$ is a global resource. We denote the projections to the resource components by $\mathcal{R}_{\text{hp}}$, $\mathcal{R}_{\text{loc}}$ and $\mathcal{R}_{\text{glo}}$. The definition of the forcing relation $\models$ is pretty standard, and we refer to [16] for details.

2.4 Soundness Theorems

Below, we define a verification system whose top level judgment is $\text{prog} : \diamond$ (read: “*prog* is verified”). We have proven a *preservation theorem* from which we can draw several corollaries, namely, *data race freedom*, *null error freedom* and *partial correctness*.

A pair (hc, hc') of head commands is called a *data race* iff $hc = (o.f = v)$ and either $hc' = (o.f = v')$ or $hc' = (\ell = o.f)$ for some o, f, v, v', ℓ . A head command hc is called a *null error* iff $hc = (\ell = \text{null}.f)$ or $hc = (\text{null}.f = v)$ or $hc = (\ell = \text{null}.m < \bar{\pi} > (\bar{v}))$. We define *initial states*: $\text{init}(c) = \langle \{\text{main} \mapsto (\text{Thread}, \emptyset)\}, \text{main is } (\emptyset \text{ in } c) \rangle$, where *main* is some distinguished object id for the main thread. The main thread has an empty set of fields (hence the first $\emptyset$), and its stack is initially empty (hence the second $\emptyset$).

Theorem 1 (Verified Programs are Data Race Free). *If $(ct, c) : \diamond$ and $\text{init}(c) \rightarrow_{ct}^* \langle h, ts \mid o \text{ is } (s_1 \text{ in } hc_1; c_1) \mid o_2 \text{ is } (s_2 \text{ in } hc_2; c_2) \rangle$, then (hc_1, hc_2) is not a data race.*

Theorem 2 (Verified Programs are Null Error Free). *If $(ct, c) : \diamond$ and $\text{init}(c) \rightarrow_{ct}^* \langle h, ts \mid o \text{ is } (s \text{ in } hc; c) \rangle$, then hc is not a null error.*

Theorem 3 (Partial Correctness).

If $(ct, c) : \diamond$ and $\text{init}(c) \rightarrow_{ct}^ \langle h, ts \mid o \text{ is } (s \text{ in } \text{assert}(F); c) \rangle$, then $(\Gamma \vdash \mathcal{E}; \mathcal{R}; s \models F[\sigma])$ for some $\Gamma, \mathcal{E}, \mathcal{R}$ such that $\mathcal{R}_{\text{hp}} = h$, and $\sigma \in \text{LogVar} \rightarrow \text{SpecVal}$.*

3 The Verification System

3.1 Proof Theory

Many presentations of separation logic are based on a model-theoretic logical consequence. We, instead, define logical consequence proof-theoretically. This gives our system an algorithmic flavour, similar to recent static assertion checkers for fragments of separation logic [2, 9] that are built upon proof-theoretic decision procedures⁷.

$$\begin{array}{ll} \Gamma; v; \bar{F} \vdash G & \text{from } v\text{'s point of view, } G \text{ is a logical consequence of the } * \text{-conjunction of } \bar{F} \\ \Gamma; v \vdash F & \text{from } v\text{'s point of view, } F \text{ is an axiom} \end{array}$$

In the former judgment, $\bar{F}$ is a *multiset* of formulas. The parameter v represents the *current receiver*, which is needed to determine the scope of predicate definitions.

The logical consequence judgment is driven by standard natural deduction rules that are common to the logic of bunched implications [23] and linear logic [30]. These rules are detailed in [16]. We admit weakening, because Java is a garbage-collected language. The link between $\Gamma; v; \bar{F} \vdash G$ and the axiom judgment $\Gamma; v \vdash F$ is established by the following rule. (We omit the definitions of typing judgments $\Gamma \vdash v : T$ and $\Gamma \vdash F : \diamond$.)

⁷ Unfortunately, these fragments do not include $\neg*$, as needed for our iterator implementation.

$$\frac{\Gamma; v \vdash G \quad \Gamma \vdash v : \text{Object} \quad \Gamma \vdash \bar{F}, G : \diamond}{\Gamma; v; \bar{F} \vdash G}$$

We now define the complete set of axioms. First, we repeat the split/merge law:

$$\begin{aligned} \Gamma; v \vdash \text{PointsTo}(e.f, \pi, e') & \text{--} * (\text{PointsTo}(e.f, \frac{\pi}{2}, e') * \text{PointsTo}(e.f, \frac{\pi}{2}, e')) \\ \Gamma; v \vdash \text{Perm}(e.\text{join}, \pi) & \text{--} * (\text{Perm}(e.\text{join}, \frac{\pi}{2}) * \text{Perm}(e.\text{join}, \frac{\pi}{2})) \end{aligned}$$

For the following axioms, recall that “ F assures G ” abbreviates “ $F \text{--} * (F * G)$ ”.

$$\begin{aligned} \Gamma; v \vdash \text{true} \quad \Gamma; v \vdash \text{false} & \text{--} * F \quad \Gamma; v \vdash (e \& F) \text{--} * (e * F) \\ \Gamma; v \vdash (\text{PointsTo}(e.f, \pi, e') \& \text{PointsTo}(e.f, \pi', e'')) & \text{assures } e' == e'' \\ (\Gamma \vdash e, e' : T \wedge \Gamma, x : T \vdash F : \diamond) \Rightarrow \Gamma; v \vdash & (F[e/x] * e == e') \text{--} * F[e/x] \end{aligned}$$

The third of these axioms implies that boolean expressions are copyable: $e \text{--} * (e * e)$.

The following axiom lifts semantic validity of boolean expressions (which we do not axiomatize) to our proof theory:

$$(\Gamma \models !e_1 \mid !e_2 \mid e') \Rightarrow \Gamma; v \vdash (e_1 * e_2) \text{--} * e'$$

The next axiom allows to apply class axioms. Here, $\text{axiom}(t < \bar{\pi}' >)$ is the $*$ -conjunction of all class axioms in $t < \bar{\pi}' >$ and its supertypes.

$$(\Gamma \vdash \pi : t < \bar{\pi}' > \wedge \text{axiom}(t < \bar{\pi}' >) = F) \Rightarrow \Gamma; v \vdash F[\pi/\text{this}]$$

The *open/close axiom* allows predicate receivers to replace abstract predicates by their definitions. It uses a function $\text{pbody}(v.P < \bar{\pi} >, C < \bar{\pi}' >)$ that returns the extension F of predicate $v.P < \bar{\pi} >$ in class $C < \bar{\pi}' >$.

$$\begin{aligned} (\Gamma \vdash v : C < \bar{\pi}'' > \wedge \text{pbody}(v.P < \bar{\pi}, \bar{\pi}' >, C < \bar{\pi}'' >) = F \wedge C <_1 D) \\ \Rightarrow \Gamma; v \vdash v.P @ C < \bar{\pi}, \bar{\pi}' > \text{--} * (F * v.P @ D < \bar{\pi} >) \end{aligned}$$

Note that the current receiver, as represented on the left of $\vdash$, has to match the predicate receiver on the right. This rule is the only reason why our logical consequence judgment tracks the current receiver. Note also that $P @ C$ may have a higher arity than $P @ D$: following [26] we allow subclasses to extend predicate arities.

The following axiom deals with unqualified predicates with missing parameters:

$$\Gamma; v \vdash \pi.P < \bar{\pi} > \text{--} * (\text{ex } \bar{T} \bar{\alpha}) (\pi.P < \bar{\pi}, \bar{\alpha} >)$$

The following axioms capture additional facts about abstract predicates. Recall that “ F ispartof G ” is defined as $G \text{--} * (F * (F \text{--} * G))$.

$$\begin{aligned} \Gamma; v \vdash \text{null.k} < \bar{\pi} > \quad \Gamma; v \vdash \pi.P @ \text{Object} \quad \Gamma; v \vdash \pi.P @ C < \bar{\pi} > \text{ispartof } \pi.P < \bar{\pi} > \\ C \preceq D \Rightarrow \Gamma; v \vdash \pi.P @ D < \bar{\pi} > \text{ispartof } \pi.P @ C < \bar{\pi}, \bar{\pi}' > \end{aligned}$$

The next axioms allow to drop the class modifier C from $\pi.P @ C$, if we know that C is π ’s dynamic class:

$$\begin{aligned} \Gamma; v \vdash (\pi.P @ C < \bar{\pi} > * C \text{isclassof } \pi) \text{--} * \pi.P < \bar{\pi} > \\ (C \text{ is final or } P \text{ is final in } C) \Rightarrow \Gamma; v \vdash \pi.P @ C < \bar{\pi} > \text{--} * \pi.P < \bar{\pi} > \end{aligned}$$

Here, the expression “ $\pi \text{isclassof } C$ ” evaluates to `true` whenever C is π ’s dynamic class. “ $C \text{isclassof } \pi$ ” surely holds right after object π of class C has been created. Consequently, our Hoare rules introduce it as a postcondition of object creation commands. The second axiom makes use of `final` classes (resp. predicates), which are classes (resp. predicates) that are prohibited to be extended.

3.2 Method Subtyping

Method types are of the following form:

$$\langle \bar{T} \bar{\alpha} \rangle \text{req } F; \text{ens } G; U m(V_0 i_0; \bar{V} \bar{i})$$

In method types, we make the self-parameter explicit, separated from the other formal parameters by a semicolon. In the scheme above, i_0 is the self-parameter.

Before presenting the method subtyping rule in full generality, we present its instance for method types without logic parameters:

$$\frac{U, V_0, \bar{V}' <: U', V'_0, \bar{V} \quad \Gamma, i_0 : V_0, \bar{i} : \bar{V}'; i_0; \text{true} \vdash F' \multimap (F * (\text{fa } U \text{ result}) (G \multimap G'))}{\Gamma \vdash \text{req } F; \text{ens } G; U m(V_0 i_0; \bar{V} \bar{i}) <: \text{req } F'; \text{ens } G'; U' m(V'_0 i_0; \bar{V}' \bar{i})}$$

This rule has the following two derived rules (where types are elided):

$$\frac{\vdash F' \multimap F \quad \vdash G' \multimap G}{\vdash \text{req } F; \text{ens } G <: \text{req } F'; \text{ens } G'} \quad \frac{}{\vdash \text{req } F; \text{ens } G <: \text{req } F * H; \text{ens } G * H}$$

The first of these derived rules is standard behavioural subtyping, the second one abstracts separation logic's frame rule. In order to see that these two rules follow from the above rule, note that the following two formulas are tautologies (as can be easily proven by natural deduction):

$$(F' \multimap F) * H \multimap F' \multimap F * H \quad F * H \multimap F * (\text{fa } U x) (G \multimap G * H)$$

The general method subtyping rule also accounts for logic parameters:⁸

$$\frac{m \neq \text{run} \quad \bar{T}', U, V_0, \bar{V}' <: \bar{T}, U', V'_0, \bar{V} \quad \Gamma, i_0 : V_0; i_0; \text{true} \vdash (\text{fa } \bar{T}' \bar{\alpha}) (\text{fa } \bar{V}' \bar{i}) (F' \multimap (\text{ex } \bar{W} \bar{\alpha}') (F * (\text{fa } U \text{ result}) (G \multimap G'))))}{\Gamma \vdash \langle \bar{T} \bar{\alpha}, \bar{W} \bar{\alpha}' \rangle \text{req } F; \text{ens } G; U m(V_0 i_0; \bar{V} \bar{i}) <: \langle \bar{T}' \bar{\alpha} \rangle \text{req } F'; \text{ens } G'; U' m(V'_0 i_0; \bar{V}' \bar{i})}$$

Note that the subtype may have more logic parameters than the supertype. For instance, we obtain the following derived rule:

$$\frac{}{\vdash \langle T \alpha \rangle \text{req } F; \text{ens } G <: \text{req } (\text{ex } T \alpha) (F); \text{ens } (\text{ex } T \alpha) (G)}$$

This derived rule is an abstraction of separation logic's auxiliary variable rule. It follows from the method subtyping rule by the following tautology:

$$(\text{ex } T \alpha) (F) \multimap (\text{ex } T \alpha) (F * (\text{fa } U x) (G \multimap (\text{ex } T \alpha) (G)))$$

3.3 Hoare Triples

Our Hoare rules are syntax-directed, omitting structural rules. Separation logic's frame rule is admissible. Separation logic's auxiliary variable rule is subsumed by our syntax for existential unpacking. We omit the rules of conjunction and disjunction, and did not need them in the examples we considered. We could soundly add the rule of disjunction. To add the rule of conjunction, we would need to assume that preconditions of $\text{run}()$ are *supported* [14].⁹

⁸ The subtyping rule for run is restricted to avoid dependencies between pre- and postcondition.

⁹ Supported formulas are formulas that have the property that, for any resource, the set of sub-resources that satisfy it is either empty or has a least element. They play a similar role for intuitionistic predicates, as *precise* formulas for non-intuitionistic predicates [24]. In our variant of separation logic, all predicates are intuitionistic, as we admit weakening.

Hoare triples have the forms $(\Gamma; v \vdash \{F\}c : T\{G\})$ and $(\Gamma; v \vdash \{F\}hc\{G\})$, where v is the receiver parameter. We present a few selected rules and refer to [16] for the complete rule system.

The rules for reading and writing fields are standard:

$$\frac{\Gamma; v; F \vdash \text{PointsTo}(w.f, \pi, u) \quad \Gamma \vdash w : U \quad T f \in \text{fld}(U) \quad T[w/\text{this}] <: \Gamma(\ell) \quad \ell \notin F}{\Gamma; v \vdash \{F\}\ell = w.f\{F * \ell == u\}}$$

$$\frac{\Gamma \vdash v, F : \text{Object}, \diamond \quad \Gamma \vdash u : U \quad T f \in \text{fld}(U) \quad \Gamma \vdash w : T[u/\text{this}]}{\Gamma; v \vdash \{F * \text{PointsTo}(u.f, 1, T)\}u.f = w\{F * \text{PointsTo}(u.f, 1, w)\}}$$

The rule for forking a thread consumes run's precondition. The postcondition of `fork()` is empty.¹⁰ The rule makes use of the function `mtype(m, T)`, which looks up m 's type in the smallest supertype of T that declares m :

$$\frac{\text{mtype}(\text{run}, T) = \text{req } G; \text{ens } G'; \text{void run}(T \iota_0;) \quad \ell \notin F \quad \Gamma(\ell) = \text{void} \quad \Gamma \vdash u : T <: \text{Thread} \quad \Gamma; v; F \vdash u != \text{null}}{\Gamma; v \vdash \{F * G[u/\iota_0]\}\ell = u.\text{fork}()\{F\}}$$

The most interesting rule is the one for joining threads. It allows the caller to exchange a fraction fr of the join-permission `Perm(u.join, 1)` for a fraction fr of $u.\text{run}$'s postcondition:¹¹

$$\frac{\text{mtype}(\text{run}, T) = \text{req } G; \text{ens } G'; \text{void run}(T \iota_0;) \quad fr = \text{all} \text{ or } G' \text{ is supported} \quad \ell \notin F \quad \Gamma(\ell) = \text{void} \quad \Gamma \vdash u : T <: \text{Thread} \quad \Gamma; v; F \vdash u != \text{null}}{\Gamma; v \vdash \{F * fr \cdot \text{Perm}(u.\text{join}, 1)\}\ell = u.\text{join}()\{F * fr \cdot G'[u/\iota_0]\}}$$

Here, fr ranges over *linear combinations*. These represent numbers of the forms 1 or $\sum_{i=1}^n bit_i \cdot \frac{1}{2^i}$:

$$bit \in \{0, 1\} \quad bits ::= 1 \mid bit, bits \quad fr \in \text{BinFrac} ::= \text{all} \mid fr() \mid fr(bits)$$

To define the scalar multiplication $fr \cdot F$, we first extend the split-operation from permissions to formulas:

$$\begin{aligned} \text{split}(e) &\triangleq e & \text{split}(\pi. \kappa < \bar{\pi}' >) &\triangleq \pi. \kappa < \text{split}(\bar{\pi}') > \\ \text{split}(\text{PointsTo}(e.f, \pi, e')) &\triangleq \text{PointsTo}(e.f, \text{split}(\pi), e') \\ \text{split}(\text{Perm}(e.\text{join}, \pi)) &\triangleq \text{Perm}(e.\text{join}, \text{split}(\pi)) \\ \text{split}(F \text{ lop } G) &\triangleq \text{split}(F) \text{ lop } \text{split}(G) & \text{split}((qt \ T \ \alpha)(F)) &\triangleq (qt \ T \ \alpha)(\text{split}(F)) \end{aligned}$$

Now, the scalar multiplication $fr \cdot F$ is defined as follows: $\text{all} \cdot F = F$, $fr() \cdot F = \text{true}$, $fr(1) \cdot F = \text{split}(F)$, $fr(0, bits) \cdot F = fr(bits) \cdot \text{split}(F)$, and $fr(1, bits) \cdot F = \text{split}(F) * fr(bits) \cdot \text{split}(F)$. For instance, $fr(1, 0, 1) \cdot F ** (\text{split}(F) * \text{split}^3(F))$.

Via the bijection $fr(bits) \mapsto \sum_{i=1}^n bit_i \cdot \frac{1}{2^i}$, we can define an addition on linear combinations that reflects the addition on concrete binary fractions. For proving soundness of the `join()`-rule, it is crucial that `join()`'s postcondition satisfies the following distributivity law, which holds if G' is supported:

$$(fr_1 + fr_2) \cdot G' ** (fr_1 \cdot G' * fr_2 \cdot G')$$

¹⁰ The permission `Perm(u.join, 1)` gets introduced when the thread object u is created.

¹¹ We assume that postconditions of methods with return type `void` do not mention the `result`-variable.

4 Comparison to Related Work and Conclusion

Parkinson/Bierman are the first to adapt separation logic to a Java-like language [25,26]. We build on their work, using abstract predicates, but extend it to a concurrent language and combine abstract predicates with fractional permissions.

Boyland and Retert [7] explain the relation between write-effects, uniqueness and datagroups in terms of a linear type-and-effect system. Their system features a nesting operation and recursive definitions, which serve as an abstraction mechanism similar to abstract predicates, but in addition promote linear formulas to non-linear ones. Recently, Boyland presented a semantics for formulas that combine nesting and fractional permissions [6]. His semantics is quite different from ours. Generally speaking, our semantics is closer to standard semantics of BI [23]. Boyland facilitates permission splitting for datagroups through an operation that scales formulas by fractions, whereas we require datagroups to be fully permission-parametrized and scale the parameters. Because we allow multiple parameters, our approach permits more fine-grained scaling for overlapping datagroups (see Section 1.6 for an example).

Bierhoff and Aldrich [3] combine tpestates and fractional permissions to specify object usage protocols. They use iterators as an example, but they do not allow linear implications in method contracts. As a result, their usage protocol regulates access to the collection itself, but not access to the elements of the collection, and their protocol would not prevent data races in concurrent programs. Krishnaswami [19] (in higher-order separation logic) and Boyland et al [8] (in their linear type-and-effect system) present iterator contracts that use linear implication and are related to ours.

Gotsman et al [14] recently adapted concurrent separation logic to Posix threads, treating storable locks. They do not support read-sharing of join's postcondition like us.

Regarding the interplay between abstract predicates and subclassing, we axiomatize the “stack of class frames” [12,1] to control predicate extensions in subclasses. The stack of class frames supports the use of subclassing for specialization and is well-suited for dealing with extended object state. Furthermore, the stack of class frames facilitates fully modular verification, avoiding the need to re-verify inherited methods, which is required in [25,26] where unrestricted predicate re-definitions in subclasses are allowed. In recent work, Parkinson and Bierman argue that a verification systems should support subclassing for code reuse in addition to subclassing for specialization, and present a system that supports both uses of subclassing while avoiding re-verification of inherited methods [27]. To this end, they associate with each method *two* contracts: a concrete “static” contract, and an abstract “dynamic” contract. Their system checks that predicate re-definitions in subclasses are compatible with concrete static contracts of inherited methods, thereby avoiding re-verification of implementations of inherited methods. The advantage over the stack of class frames is increased flexibility, the disadvantage is heavier specification machinery, although much of this can be hidden behind good defaults. Chin et al [9] make a similar proposal.

Conclusion. We have presented a variant of concurrent separation logic with fractional permissions for a Java-like language with fork/join and proved it sound. Future work includes algorithmic checking and extension to handle lock synchronization.

Acknowledgments. We thank John Boyland, Marieke Huisman, Erik Poll and anonymous reviewers for their very useful comments that helped to improve this paper.

References

1. M. Barnett, R. DeLine, M. Fähndrich, K. R. M. Leino, W. Schulte. Verification of object-oriented programs with invariants. *Journal of Object Technology*, 3(6), 2004.
2. J. Berdine, C. Calcagno, P. W. O'Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *Formal Methods for Components and Objects*, 2005.
3. K. Bierhoff, J. Aldrich. Modular typestate verification of aliased objects. In *ACM Conference on Object-Oriented Programming Systems, Languages, and Applications*, 2007.
4. R. Bornat, P. O'Hearn, C. Calcagno, M. Parkinson. Permission accounting in separation logic. In *Principles of Programming Languages*, New York, NY, USA, 2005. ACM Press.
5. J. Boyland. Checking interference with fractional permissions. In R. Cousot, ed., *Static Analysis Symposium*, vol. 2694 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
6. J. Boyland. Semantics of fractional permissions with nesting. Technical report, University of Wisconsin at Milwaukee, 2007.
7. J. Boyland, W. Retert. Connecting effects and uniqueness with adoption. In *Principles of Programming Languages*, 2005.
8. J. Boyland, W. Retert, Y. Zhao. Iterators can be independent "from" their collections. International Workshop on Aliasing, Confinement and Ownership in object-oriented programming, 2007.
9. W. Chin, C. David, H. Nguyen, S. Qin. Enhancing modular OO verification with separation logic. In *Principles of Programming Languages*, 2008.
10. K. Cray, D. Walker, G. Morrisett. Typed memory management in a calculus of capabilities. In *Principles of Programming Languages*, 1999.
11. R. DeLine, M. Fähndrich. Enforcing high-level protocols in low-level software. In *Programming Languages Design and Implementation*, 2001.
12. R. DeLine, M. Fähndrich. Typestates for objects. In *European Conference on Object-Oriented Programming*, 2004.
13. J.-Y. Girard. Linear logic: Its syntax and semantics. In J.-Y. Girard, Y. Lafont, L. Regnier, eds., *Advances in Linear Logic*. Cambridge University Press, 1995.
14. A. Gotsman, J. Berdine, B. Cook, N. Rinetzky, M. Sagiv. Local reasoning for storable locks and threads. In *Asian Programming Languages and Systems Symposium*, 2007.
15. C. Haack, C. Hurlin. Resource usage protocols for iterators. <http://www.cs.ru.nl/~chaack/papers/iterators.pdf>.
16. C. Haack, C. Hurlin. Separation logic contracts for a Java-like language with fork/join. Technical Report 6430, INRIA, 2008.
17. A. Igarashi, B. Pierce, P. Wadler. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Trans. Program. Lang. Syst.*, 23(3), 2001.
18. S. Ishtiaq, P. O'Hearn. BI as an assertion language for mutable data structures. In *Principles of Programming Languages*, 2001.
19. G. Krishnaswami. Reasoning about iterators with separation logic. In *Specification and Verification of Component-Based Systems*, 2006.
20. G. T. Leavens, A. L. Baker, C. Ruby. Preliminary design of JML: a behavioral interface specification language for Java. *SIGSOFT Software Engineering Notes*, 31(3), 2006.
21. K. R. M. Leino. Data groups: Specifying the modification of extended state. In *ACM Conference on Object-Oriented Programming Systems, Languages, and Applications*, 1998.
22. P. O'Hearn. Resources, concurrency and local reasoning. *Theor. Comp. Science*, 375(1–3), 2007.
23. P. W. O'Hearn, D. J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2), 1999.
24. P. W. O'Hearn, H. Yang, J. C. Reynolds. Separation and information hiding. In *Principles of Programming Languages*, Venice, Italy, 2004. ACM Press.
25. M. Parkinson. Local reasoning for Java. Technical Report UCAM-CL-TR-654, University of Cambridge, 2005.
26. M. Parkinson, G. Bierman. Separation logic and abstraction. In *Principles of Programming Languages*, 2005.
27. M. Parkinson, G. Bierman. Separation logic, abstraction and inheritance. In *Principles of Programming Languages*, 2008.
28. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science*, Copenhagen, Denmark, 2002. IEEE Press.
29. F. Smith, D. Walker, G. Morrisett. Alias types. In G. Smolka, ed., *European Symposium on Programming*, vol. 1782 of *Lecture Notes in Computer Science*. Springer-Verlag, 2000.
30. P. Wadler. A taste of linear logic. In *Mathematical Foundations of Computer Science*, 1993.