

Type-Based Race Detection for Java

Cormac Flanagan
Compaq Systems Research Center
130 Lytton Ave.
Palo Alto, CA 94301
cormac.flanagan@compaq.com

Stephen N. Freund*
Department of Computer Science
Stanford University
Stanford, CA 94305-9045
freunds@cs.stanford.edu

Abstract

This paper presents a static race detection analysis for multithreaded Java programs. Our analysis is based on a formal type system that is capable of capturing many common synchronization patterns. These patterns include classes with internal synchronization, classes that require client-side synchronization, and thread-local classes. Experience checking over 40,000 lines of Java code with the type system demonstrates that it is an effective approach for eliminating races conditions. On large examples, fewer than 20 additional type annotations per 1000 lines of code were required by the type checker, and we found a number of races in the standard Java libraries and other test programs.

1 Introduction

Race conditions are common, insidious errors in multithreaded programs. A race condition occurs when two threads manipulate a shared data structure simultaneously, without synchronization. Race conditions often result in unexpected program behavior, such as program crashes or incorrect results. They can be avoided by careful programming discipline: protecting each data structure with a lock and acquiring that lock before manipulating the data structure [Bir89]. Since a lock can be held by at most one thread at any time, careful adherence to this lock-based synchronization discipline ensures a race-free program.

Current programming tools provide little support for this synchronization discipline. It is easy to write a

program that, by mistake, neglects to perform certain crucial synchronization operations. These synchronization errors are not detected by traditional compile-time checks. Furthermore, because the resulting race conditions are scheduler dependent, they are difficult to catch using testing techniques. A single synchronization error in an otherwise correct program may yield a race condition whose cause may take weeks to identify [SBN⁺97].

This paper investigates a static analysis system for detecting race conditions in Java programs. The analysis supports the lock-based synchronization discipline by tracking the protecting lock for each shared field in the program and verifying that the appropriate lock is held whenever a shared field is accessed. We express the reasoning and checks performed by this analysis as an extension of Java's type system.

This work builds on an earlier paper that describes a race-free type system for a concurrent object calculus [FA99a]. We start by adapting that type system to a core subset of Java. This initial type system is sufficient to verify some example programs as race free. In order to accommodate larger, more realistic, multithreaded programs, we extend the initial type system with a number of additional features. These features include:

1. classes parameterized by locks, which allow the fields of a class to be protected by some lock external to the class;
2. the notion of objects that are local to a particular thread and therefore safely accessible without synchronization; and
3. mechanisms for escaping the type system in places where it proves too restrictive, or where a particular race condition is considered benign.

To evaluate the utility of the resulting type system, we implemented a type checker and tested it on a variety of Java programs totaling over 40,000 lines of code. These programs include the standard Java input/output package `java.io`; an interpreter for the web

*This work was completed while the author was employed at the Compaq Systems Research Center.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
PLDI 2000, Vancouver, British Columbia, Canada.
Copyright 2000 ACM 1-58113-199-2/00/0006...\$5.00.

scripting language WebL; and Ambit, a mobile object calculus implementation.

Checking these programs using our type system requires adding some additional type annotations. This annotation burden is not excessive; typically fewer than 20 annotations were required per 1000 lines of code. Most of the annotations were inserted based on feedback from the type checker. This annotation process proceeded at a rate of roughly 1000 lines of code per programmer-hour. During this process, we discovered a number of race conditions in the programs being checked, including one race condition in `java.util.Vector`, four in the `java.io` package, and five in the WebL implementation. Although it is far from complete, the type system proved sufficiently expressive to accommodate the majority of synchronization patterns present in these programs.

The presentation of our results proceeds as follows. Section 2 introduces a small concurrent subset of Java, which we use to provide a formal description of our type system. Section 3 describes an initial race-free type system. We extend the system to include classes parameterized by locks in Section 4 and thread-local classes in Section 5. Section 6 describes our prototype implementation, including the escape mechanisms. Section 7 discusses our experiences checking several Java programs. We relate this work to other projects in Section 8, and we conclude in Section 9. The Appendix contains a formal description of the type system.

2 A Multithreaded Subset of Java

This section introduces a small multithreaded subset of Java, `CONCURRENTJAVA`. This language is derived from `CLASSICJAVA` [FKF98], a sequential subset of Java, and we adopt much of the type structure and semantics of `CLASSICJAVA`.

2.1 Syntax and Informal Semantics

`CONCURRENTJAVA` supports multithreaded programs by including the operation `fork e` which spawns a new thread for the evaluation of e . This evaluation is performed only for its effect; the result of e is never used. Locks are provided for thread synchronization: each object has an associated lock that has two states, locked and unlocked, and is initially unlocked. The expression `synchronized e_1 in e_2` is evaluated in a manner similar to Java’s `synchronized` statement: the subexpression e_1 is evaluated first, and should yield an object, whose lock is then acquired; the subexpression e_2 is then evaluated; and finally the lock is released. The result of e_2 is returned as the result of the synchronized expression. While evaluating e_2 , the current thread is said to hold

P	$::=$ <code>defn* e</code>	(program)
<code>defn</code>	$::=$ <code>class cn body</code>	(class decl)
<code>body</code>	$::=$ <code>extends c</code> <code>{ $field^* meth^*$ }</code>	(class body)
<code>field</code>	$::=$ <code>[final]_{opt} t $fd = e$</code>	(field decl)
<code>meth</code>	$::=$ <code>t $mn(arg^*)$ { e }</code>	(method decl)
<code>arg</code>	$::=$ <code>t x</code>	(variable decl)
s, t	$::=$ <code>c int</code>	(type)
c	$::=$ <code>cn Object</code>	(class type)
e	$::=$ <code>new c</code>	(allocate)
	<code>x</code>	(variable)
	<code>$e.f$</code>	(field access)
	<code>$e.f = e$</code>	(field update)
	<code>$e.mn(e^*)$</code>	(method call)
	<code>let $arg = e$ in e</code>	(variable binding)
	<code>synchronized e in e</code>	(synchronization)
	<code>fork e</code>	(fork)
	$cn \in$ class names	
	$fd \in$ field names	
	$mn \in$ method names	
	$x, y \in$ variable names	

Figure 1: The grammar for `CONCURRENTJAVA`.

the lock. Any other thread that attempts to acquire the lock blocks until the lock is released. A newly forked thread does not inherit locks held by its parent thread.

The syntax of the `synchronized` and `fork` expressions and the rest of `CONCURRENTJAVA` is shown in Figure 1. A program is a sequence of class definitions together with an initial expression, which is the starting point for the program’s execution. Each class definition associates a class name with a class body consisting of a super class, a sequence of field declarations, and a sequence of method declarations. A field declaration includes an initialization expression and an optional `final` modifier; if this modifier is present, then the field cannot be updated after initialization. A method declaration consists of the method name, its return type, number and types of its arguments, and an expression for the method body. Types include class types and integers. Class types include class names introduced by the program, as well as the predefined class `Object`, which serves as the root of the class hierarchy. Expressions include the typical operations for object allocation, field access and update, method invocation, and variable binding and reference, as well as the concurrency primitives.

2.2 Locks Against Races

We present example programs in an extended language with integer and boolean constants and operations, and the constant `null`. We use $e_1; e_2$ to abbreviate `let $x = e_1$ in e_2` , where x does not occur free in e_2 , and we sometimes enclose expressions in braces for clarity.

Multithreaded `CONCURRENTJAVA` programs are prone to race conditions, as illustrated by the following program, which allocates a new bank account, and makes two deposits into the account in parallel:

```
class Account {
  int balance = 0
  int deposit (int x) {
    this.balance = this.balance + x
  }
}

let Account a = new Account in {
  fork { a.deposit(10) }
  fork { a.deposit(10) }
}
```

The program may exhibit unexpected behavior. In particular, if the two calls to `deposit` are interleaved, the final value of `balance` may reflect only one of the two deposits made to the account, which is clearly not the intended behavior of the program. Thus, the program contains a race condition: two threads attempt to manipulate the field `balance` simultaneously, with incorrect results.

We can fix this error by protecting the field `balance` by the lock of the account object and only accessing or updating `balance` when that lock is held:

```
class Account {
  int balance = 0
  int deposit(int x) {
    synchronized this in {
      this.balance = this.balance + x
    }
  }
}
```

The modified account implementation is race free and will behave correctly even when multiple deposits are made to the account concurrently.

3 Types Against Races

In practice, race conditions are commonly avoided by the lock-based synchronization discipline used in the example above. This section presents a type system that supports this programming discipline. The type system

needs to verify that each field has an associated protecting lock that is held whenever the field is accessed or updated. In order to verify this property, the type system:

1. associates a protecting lock with each field declaration, and
2. tracks the set of locks held at each program point.

We rely on the programmer to aid the verification process by providing a small number of additional type annotations. The type annotation `guarded_by l` on a field declaration states that the field is protected by the lock expression l ; the type system then verifies that this lock is held whenever the field is accessed or updated. The type annotation `requires l_1, \dots, l_n` on a method declaration states that the locks l_1, \dots, l_n are held on method entry; the type system verifies that these locks are indeed held at each call-site of the method, and checks that the method body is race-free given this assumption. We extend the syntax of field and method declarations to include these type annotations.

$$\begin{aligned} \text{field} &::= [\text{final}]_{\text{opt}} t \text{ fd } \text{ guarded_by } l = e \\ \text{meth} &::= t \text{ mn}(\text{arg}^*) \text{ requires } ls \{ e \} \\ ls &::= l^* && \text{(lock set)} \\ l &::= e && \text{(lock expression)} \end{aligned}$$

We refer to the extended language as `RACEFREEJAVA`.

To ensure that each field is consistently protected by a particular lock, irrespective of any assignments performed by the program, the type system requires that the lock expression in a `guarded_by` clause be a *final expression*. A final expression is either a reference to an immutable variable¹, or a field access $e.f$, where e is a final expression and f is a final field. The type system also requires that the lock expressions in a `requires` clause be final for similar reasons.

The core of our type system is a set of rules for reasoning about the type judgment

$$P; E; ls \vdash e : t.$$

Here, P (the program being checked) is included in the judgment to provide information about class definitions in the program; E is an environment providing types for the free variables of e ; ls is a set of final expressions describing the locks that are held when the expression e is evaluated; and t is the type of e . Thus, the type rules track the set of locks held each program point. The rule [EXP FORK] for `fork e` checks the expression e using the empty lock set since new threads do not inherit locks held by their parent.

$$\frac{[\text{EXP FORK}] \quad P; E; \emptyset \vdash e : t}{P; E; ls \vdash \text{fork } e : \text{int}}$$

¹All variables are immutable in `RACEFREEJAVA`, but only final variables are in Java.

The rule [EXP SYNC] for **synchronized** e_1 in e_2 checks that e_1 is a final expression of some class type c , and then type checks e_2 in an extended lock set that includes e_1 , reflecting the fact that the lock e_1 is always held when evaluating e_2 .

$$\frac{[\text{EXP SYNC}] \quad \frac{P; E \vdash_{\text{final}} e_1 : c \quad P; E; ls \cup \{e_1\} \vdash e_2 : t}{P; E; ls \vdash \text{synchronized } e_1 \text{ in } e_2 : t}}$$

The antecedent $P; E \vdash_{\text{final}} e_1 : c$ checks that e_1 is a final expression of type c .

The rule [EXP REF] for $e.f.d$ checks that e is a well-typed expression of some class type c and that c declares or inherits a field fd of type t , guarded by lock l .

$$\frac{[\text{EXP REF}] \quad \frac{P; E; ls \vdash e : c \quad P; E \vdash ([\text{final}]_{\text{opt}} t \text{ fd guarded_by } l = e') \in c \quad P; E \vdash [e/\text{this}]l \in ls \quad P; E \vdash [e/\text{this}]t}{P; E; ls \vdash e.f.d : [e/\text{this}]t}}$$

It remains to check that the lock l guarding fd is held at this program point; i.e., that l denotes the same lock as some expression l' in the current lock set. This requires checking the semantic equivalence of the two expressions l and l' , which is in general undecidable.

One approach is to conservatively approximate semantic equivalence by syntactic equivalence, and simply to check that $l \equiv l'$. This approximation is overly conservative in many cases however. In particular, occurrences of **this** in the lock expression l refer to the object being dereferenced, which is the same object as that denoted by e . To account for the aliasing of **this** and e , the type system replaces all occurrences of **this** in l by e , and then checks that $[e/\text{this}]l \equiv l'$ for some l' in the current lock set. This check is performed by the antecedent $P; E \vdash [e/\text{this}]l \in ls$. It is a sound approximation of semantic equivalence and has been sufficient for all programs we have inspected.

A similar aliasing situation arises in the next section, where we introduce types containing lock expressions. To accommodate this future extension, we include the appropriate substitutions for types here, and yield $[e/\text{this}]t$ as the type of the field access after checking that it is a well-formed type.

The rule [EXP ASSIGN] for $e.f.d = e'$ ensures that the appropriate lock is held whenever a field is updated. The rule [EXP INVOKE] for a method invocation ensures that all locks in the **requires** clause of a method declaration are held at each call site of the method.

$$\frac{[\text{EXP ASSIGN}] \quad \frac{P; E; ls \vdash e : c \quad P; E \vdash (t \text{ fd guarded_by } l = e'') \in c \quad P; E \vdash [e/\text{this}]l \in ls \quad P; E; ls \vdash e' : [e/\text{this}]t}{P; E; ls \vdash e.f.d = e' : [e/\text{this}]t}}$$

$$\frac{[\text{EXP INVOKE}] \quad \frac{P; E; ls_1 \vdash e : c \quad P; E \vdash (t \text{ mn}(s_j \ y_j \ j \in \{1 \dots n\}) \text{ requires } ls_2 \ \{ e' \}) \in c \quad P; E; ls_1 \vdash e_j : [e/\text{this}]s_j \quad P; E \vdash [e/\text{this}]ls_2 \subseteq ls_1 \quad P; E \vdash [e/\text{this}]t}{P; E; ls_1 \vdash e.\text{mn}(e_1 \dots e_n) : [e/\text{this}]t}}$$

The remaining type rules are straightforward and similar to those of CLASSICJAVA. The complete set of type judgments and rules is contained in Appendix A.

3.1 Race-Free Bank Accounts

We can use this type system to verify that the synchronized bank account implementation is race free. We must first add a type annotation stating that the field **balance** is guarded by **this**. When no locks are required to invoke a method we omit its **requires** clause, as we have done for **deposit**.

```
class Account {
    int balance guarded_by this = 0
    int deposit(int x) {
        synchronized this in {
            this.balance = this.balance + x
        }
    }
}

let Account a = new Account in {
    fork { a.deposit(10) }
    fork { a.deposit(10) }
}
```

An alternative implementation of the bank account may rely on its clients to perform the necessary synchronization operations:

```
class Account {
    int balance guarded_by this = 0
    int deposit(int x) requires this {
        this.balance = this.balance + x
    }
}

let Account a = new Account in {
    fork { synchronized a in a.deposit(10) }
    fork { synchronized a in a.deposit(10) }
}
```

In this example, the method signature

```
int deposit(int x) requires this
```

explicitates the requirement that the object's lock must be acquired before calling **deposit**. Since the necessary lock is indeed held at each call site, this program is also well typed and race free.

4 External Locks

The type system of the previous section can verify the absence of races in a number of interesting examples. However, larger, more realistic programs frequently use a variety of synchronization patterns, some of which cannot be captured by the system presented so far. To accommodate such programs, we extend the RACEFREEJAVA type system with additional features. This section presents classes parameterized by locks, and Section 5 introduces thread-local classes.

The type system requires that every field be guarded by a final expression of the form $x.f d_1. \dots . f d_n$. Since the only variable in scope at a field declaration is **this**, the fields of an object must be protected by a lock that is accessible from the object. In some cases, however, we would like to protect the fields of an object by some lock external to the object. For example, all of the fields in a linked list might naturally be protected by some object external to the list.

To accommodate this programming pattern, we extend RACEFREEJAVA to allow classes to be parameterized by external locks:

$$\begin{aligned} \text{defn} &::= \text{class } cn \langle \text{garg}^* \rangle \text{ body} \\ \text{garg} &::= \text{ghost } t \ x \quad (\text{ghost decl}) \\ c &::= cn \langle l^* \rangle \mid \text{Object} \end{aligned}$$

A class definition now contains a (possibly empty) sequence of formal parameters or *ghost variables*. These ghost variables are used by the type system to verify that the program is race free; they do not affect the run-time behavior of the program. In particular, they can appear only in type annotations and not in regular code. A class type c consists of a class name cn parameterized by a sequence of final expressions. The number and type of these expressions must match the formal parameters of the class.

Type checking of parameterized classes is handled via substitution. If

$$\text{class } cn \langle \text{ghost } t_1 \ x_1, \dots, \text{ghost } t_n \ x_n \rangle \text{ body}$$

is a well-formed class definition, then for any final expressions l_1, \dots, l_n of the appropriate types, we consider $cn \langle l_1, \dots, l_n \rangle$ to be a valid instantiated class type, with associated instantiated class definition

$$\text{class } cn \langle l_1, \dots, l_n \rangle \ [l_1/x_1, \dots, l_n/x_n] \text{body}$$

A few modifications to the type rules are necessary to accommodate parameterized classes. These modifications are described in Appendix B.

4.1 Using External Locks

To illustrate the use of external locks, consider the dictionary implementation of Figure 2. A dictionary maps

keys to values. In our implementation, a dictionary is represented as an object containing a linked list of **Nodes**, where each **Node** contains a key, a value, and a next pointer.

For efficiency reasons, we would like to protect the entire dictionary, including its linked list, with the lock of the dictionary. To accomplish this, the class **Node** is parameterized by the enclosing dictionary d ; the fields of **Node** are guarded by d ; and each method of **Node** requires that d is held on entry. Each method of **Dictionary** first acquires the dictionary lock and then proceeds with the appropriate manipulation of the linked list. Since all fields of the linked list are protected by the dictionary lock, the type system verifies that this program is well typed and race free.

5 Thread-Local Classes

Large multithreaded programs typically have sections of code that operate on data that is not shared across multiple threads. For example, only a single thread in a concurrent web server may need to access the information about a particular request. Objects used in this fashion require no synchronization and should not need to have locks guarding their fields. To accommodate this situation, we introduce the concept of thread-local classes. We extend the grammar to allow an optional **thread_local** modifier on class definitions and to make the **guarded_by** clause on field declarations optional in a thread-local class:

$$\begin{aligned} \text{defn} &::= [\text{thread_local}]_{\text{opt}} \text{class } cn \langle \text{garg}^* \rangle \text{ body} \\ \text{field} &::= [\text{final}]_{\text{opt}} t \ fd \ [\text{guarded_by } l]_{\text{opt}} = e \end{aligned}$$

An example of a thread-local class appears in Figure 3. The class **Crawler** is a concurrent web crawler that processes a page by iterating over its links and forking new threads to process the linked pages. The **LinkEnumerator** class, which parses the text of the page to find links, is not shared among threads. Therefore, it is declared as a **thread_local** class and contains unguarded fields.

A simple form of escape analysis is used to enforce single-threaded use of thread-local objects. A type is *thread-shared* provided it is not a thread-local class type. The type system must ensure that thread-local objects are not accessible from thread-shared objects. Therefore, a thread-shared class declaration must (1) have a thread-shared superclass and (2) contain only shareable fields. A field is *shareable* only if it has a thread-shared type and is either final or protected by a lock. Also, the free variables of a forked expression must be of a thread-shared type. The rules for thread-shared types and **fork** appear in Appendix C.

```

class Node<ghost Dictionary d> {
  String key guarded_by d = null
  Object value guarded_by d = null
  Node<d> next guarded_by d = null

  void init(String k, Object v, Node<d> n)
    requires d {
      node.key = k;
      node.value = v;
      node.next = n
    }
  void update(String k, Object v) requires d {
    if (this.key.equals(k)) {
      this.value = v
    } else if (this.next != null) {
      this.next.update(k,v)
    }
  }
  ...
}

class Dictionary {
  Node<this> head guarded_by this = null

  void put(String k, Object v) {
    synchronized this in {
      if (this.contains(k)) {
        this.head.update(k,v)
      } else {
        let Node<this> node =
          new Node<this> in {
            node.init(k,v,this.head);
            this.head = node
          }
      }
    }
  }
  ...
}

```

Figure 2: A synchronized dictionary.

```

thread_local class LinkEnumerator {
  String text = null
  int index = 0

  void init(String t) {
    this.text = t
  }
  boolean hasMoreLinks() { ... }
  String nextLink() { ... }
}

class Crawler {
  final Set visited = new Set
  void process(String url) {
    if (!visited.add(url)) {
      let String text = loadPageText(url) in
      let LinkEnumerator enum =
        new LinkEnumerator in {
          enum.init(text);
          while (enum.hasMoreLinks()) {
            let String link = enum.nextLink() in
            fork { this.process(link) }
          }
        }
    }
  }
  ...
}

let Crawler c = new Crawler in
  c.process("http://www.research.compaq.com")

```

Figure 3: A concurrent web crawler using a thread-local enumeration class.

Interestingly, our type system permits a thread-local class to have a thread-shared superclass. This design permits us to maintain `Object` (which is thread-shared) as the root of the class hierarchy, as it is in Java. However, it also permits a thread-local object to be viewed as an instance of a thread-shared class and hence to be shared between threads. This sharing does not cause a problem unless the object is downcast back to the thread-local type in a thread other than the one in which it was created. This downcast would make unguarded fields in the subclass visible to more than one thread.

To eliminate this possibility, our type system forbids downcasts from a thread-shared type to a thread-local type. This restriction applies to explicit cast operations² and, also, to the implicit downcasts that occur during dynamic dispatch. To avoid such implicit downcasts, our type system requires a thread-local class not to override any methods declared in a thread-shared superclass.

Alternatively, if these static requirements are too restrictive, a compiler could insert code to track the allocating thread of each object and dynamically check that thread-shared to thread-local downcasts are only performed by the appropriate thread.

6 Implementation

We have implemented the RACEFREEJAVA type system for the full Java language [GJS96]. This race condition checker, `rccjava`, extends the type system outlined so far with the missing Java features, including arrays, interfaces, constructors, static fields and methods, inner classes, and so on. Only thread-local arrays posed any technical challenges, but space considerations prohibit a full discussion of those challenges here.

The additional type information required by `rccjava` is embedded in Java comments to preserve compatibility with existing Java tools, such as compilers. Specifically, comments that start with the character “#” are treated as type annotations by `rccjava`. See Figure 5 for an example.

The `rccjava` tool was built on top of an existing Java front-end that includes a scanner, parser, and type checker. The extensions for race detection were relatively straightforward to add to the existing code base and required approximately 5,000 lines of new code. The major additions were maintaining the lock set during type checking, implementing syntactic equality and substitution on abstract syntax trees, and incorporating classes parameterized by locks.

²RACEFREEJAVA does not contain explicit casts, but the Java language does.

An important goal in the design of `rccjava` was to provide a cost-effective way to detect race conditions statically. Thus, it was important to minimize both the number of annotations required and the number of false alarms produced. In order to attain this goal, `rccjava` was designed to be able to relax the formal type system in several ways and, also, to infer default annotations for unannotated code. These features are described below.

6.1 Escape mechanisms

`Rccjava` provides mechanisms for escaping from the type system when it proves too restrictive. The simplest escape mechanism is the `no_warn` annotation, which turns off certain kinds of warnings on a particular line of code, e.g.

```
f.a = 3; ///no_warn race
```

This annotation is commonly used if a particular race condition is considered benign.

Also, `rccjava` may be configured with a command line flag to ignore all errors of a particular kind. For example, the “`-no_warn thread_local_override`” flag turns off the restrictions whereby a thread-local class cannot override a method of its thread-shared superclass.

The `holds` annotation asserts that a particular lock is held at a given program point:

```
///holds f  
f.a = 3;
```

This annotation puts `f` into the lock set for the remainder of block of statements in which it appears. As with the `no_warn` annotations, `rccjava` may be configured to make global assumptions about when locks are held. For instance, when run with the command line flag “`-constructor_holds_lock`”, `rccjava` assumes that the lock `this` is held in constructors. This is sound as long as references to `this` are not passed to other threads before the constructor call returns. Violations of this assumption are unlikely, and using it eliminates a large number of spurious warnings.

6.2 Default Annotations

Although type inference for the `rccjava` type system remains for future work, `rccjava` does construct default annotations for unannotated classes and fields. The heuristics used to compute default annotations are:

- A class with no annotations and no synchronized methods is thread local by default, unless the class is `java.lang.Object` or a subclass of `java.lang.Thread`.

- Unguarded non-final instance fields in thread shared classes are guarded by `this`.
- Unguarded non-final static fields are guarded by the class object for the class to which they belong.
- A `guarded_by` annotation is permitted on a class declaration, and it applies to all fields of the class.

These heuristics are not guaranteed to produce the correct annotations, but experience has shown that they save a significant amount of time while annotating large programs. Roughly 90% of the classes in the test programs described below are treated correctly by these heuristics.

7 Evaluation

To test the effectiveness of `rccjava` as a static race detection tool, we used it to check several multithreaded Java programs. Our test cases include two representative single classes, `java.util.Hashtable` and `java.util.Vector`, and several large programs, including `java.io`, the Java input/output package (version 1.1) [Jav98]; `Ambit`, an implementation of a mobile object calculus [Car97]; and an interpreter and run-time environment for `WebL`, a language for automating web-based tasks [KM98].

These five programs use a variety of synchronization patterns, most of which were captured easily with `rccjava` annotations. `Rccjava` was run with the command line flags “`-no_warn thread_local_override`” and “`-constructor_holds_lock`” for these tests (see Section 6.1). Although these flags may cause `rccjava` to miss some potential races, they significantly reduce the number of false alarms reported and provide the most effective way to deal with existing programs that were not written with this type system in mind. Table 1 summarizes our experience in checking these programs. It shows the number of annotations and time required to annotate each program, as well as the number of race conditions found in each program. The time includes both the time spent by the programmer inserting annotations and the time to run the tool.

Figure 4 breaks down the annotation count into the different categories of annotations, normalized to the frequency with which they appear in 1000 lines of code. For the large programs, fewer than 20 annotations were required per 1000 lines. Most of these annotations were clustered in the small number of classes manipulated from different threads. The majority of classes typically required very few or no annotations. Evidence of this pattern is reflected in the statistics for the single class examples, which have higher annotation frequencies than the larger programs. `Hashtable` has a

high occurrence of annotations concerning class parameters and arguments because it contains a linked list similar to that of Figure 2. Interestingly, restructuring `Hashtable` to declare the linked list as an inner class within the scope of the protecting lock reduces the number of annotations to 25.

We discovered race conditions in three of the five case studies, despite most of the code in these examples being well tested and relatively mature. Of the four races found in `java.io`, one was fixed in JDK version 1.2. We also found benign race conditions in all test cases.

Figure 5 contains an excerpt from `java.util.Vector` that illustrates a typical race condition caught during our experiments. Suppose that there are two threads manipulating a shared `Vector` object. If one thread calls `lastIndexOf(elem)` for some `elem`, that method may access `elementCount` without acquiring the lock of the `Vector` object. However, the other thread may call `removeAllElements` (which sets `elementCount` to 0) and then call `trimToSize` (which resets `elementData` to an array of length 0). Thus, an array out of bounds exception will be triggered when the first thread enters the binary version of `lastIndexOf` and accesses the `elementData` array based on the old value of `elementCount`. Declaring both versions of `lastIndexOf` to be synchronized removes this race condition.

8 Related Work

A number of tools have been developed for detecting race conditions, both statically and dynamically. `Warlock` [Ste93] is a static race detection system for ANSI C programs. It supports the lock-based synchronization discipline through annotations similar to ours. However, `Warlock` uses a different analysis mechanism; it works by tracing execution paths through the program, but it fails to trace paths through loops or recursive function calls, and thus may not detect certain races. In addition, `Warlock` assumes, but does not verify, the thread-local annotations introduced by the programmer. However, these soundness issues have not prevented `Warlock` from being a practical tool. It has been used to catch races in several programs, including an X-windows library.

The extended static checker for Java (`Esc/Java`) is a tool for static detection of software defects [LSS99, DLNS98]. It uses an underlying automatic theorem prover to reason about the program’s behavior and to verify the absence of certain kinds of errors, such as null dereferences and index out of bounds errors. `ESC/Java` supports multithreaded programming via annotations similar to our `guarded_by` and `requires` clauses, and

Program	Lines Of Code	Programmer Time (hrs)	Annotations	Races Found
java.util.Hashtable	440	0.5	60	0
java.util.Vector	430	0.5	10	1
java.io.*	16,000	16.0	139	4
Ambit	4,500	4.0	38	0
WebL	20,000	12.0	358	5

Table 1: Programs analyzed using `rccjava`.

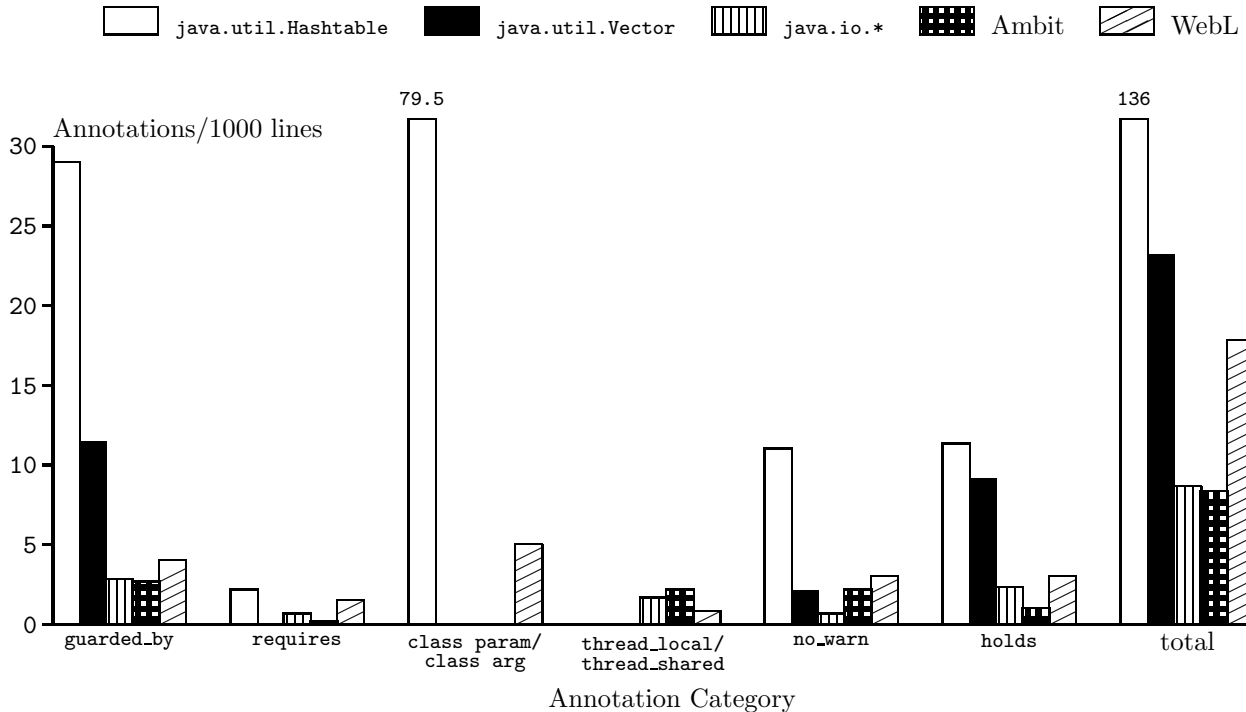


Figure 4: Number of `rccjava` annotations added to each program.

verifies that the appropriate lock is held whenever a guarded field is accessed. However, it may still permit race conditions on unguarded fields, since it does not verify that such unguarded fields only occur in thread-local classes. Overall, `Esc/Java` is a complex but powerful tool capable of detecting many kinds of errors, whereas `rccjava` is a lightweight tool tuned specifically for detecting race conditions.

Aiken and Gay [AG98] also investigate static race detection, in the context of SPMD programs. Since synchronization in these programs is performed using barriers, as opposed to locks, their system does not need to track the locks held at each program point or the association between locks the fields they protect. Their system has been used successfully on a number of SPMD programs.

Eraser is a tool for detecting race conditions and

deadlocks dynamically [SBN⁺97], rather than statically. This approach has the advantage of being able to check unannotated programs, but it may fail to detect certain errors because of insufficient test coverage.

A variety of other approaches have been developed for race and deadlock prevention; they are discussed in more detail in an earlier paper [FA99b].

A number of formal calculi for Java have been presented in recent literature. These include attempts to model the entire Java language [DE97, Sym97, NvO98] and, also, smaller systems designed to study specific features and extensions [IPW99]. We chose to use the CLASSICJAVA calculus of Flatt, Krishnamurthi, and Felleisen [FKF98] as the starting point for our study.

There have been many suggested language extensions for supporting Java classes parameterized by types [OW97, BOSW98, AFM97, BLM96, CJ98].

```

class Vector {
    Object elementData[] /*# guarded_by this */;
    int elementCount /*# guarded_by this */;

    synchronized void trimToSize() { ... }
    synchronized boolean removeAllElements() { ... }

    synchronized int lastIndexOf(Object elem, int n) {
        for (int i = n ; --i >= 0 ; )
            if (elem.equals(elementData[i])) { ... }
    }

    int lastIndexOf(Object elem) {
        return lastIndexOf(elem, elementCount); // race!!!
    }
    ...
}

```

Figure 5: Excerpt from `java.util.Vector`.

Our work uses a different notion of parameterization, namely, classes parameterized by *values* (more specifically, lock expressions). Apart from this distinction, our class parameterization approach most closely follows that of GJ [BOSW98], in that information about class parameters is not preserved at run time.

The lock sets used in our type systems are similar to *effects* [JG91, LG88, Nie96] since the locks held on entry to an expression constrain the effects that it may produce. It may be possible to adapt existing techniques for effect reconstruction [TT94, TT97, ANN97, TJ92] to our setting to reduce the number of type annotations required.

Our type system verifies that objects of a `thread_local` type are never shared between threads. Much work has been done on the related problem of inferring which objects are not shared between threads [CGS⁺99, Bla99, BH99, WR99, ACSE99]. This work has primarily focused on optimizing synchronization operations, but it may be possible to adapt this work to reduce or eliminate the need for `thread_local` annotations.

9 Conclusions and Future Work

Race conditions are difficult to catch using traditional testing techniques. They persist even in common, relatively mature Java programs. In this paper, we have presented a type system for catching race conditions statically and described `rccjava`, an implementation of this system for Java. Our experience with `rccjava` indicates that this technique is a promising approach for building more reliable multithreaded software.

Because the type system is modular, it enables race conditions to be detected early in the development cycle, before the entire program has been written. The type system does not require the programmer to write additional type annotations, but these annotations also function as documentation of the locking strategies used by the program.

To reduce the annotation overhead further, we are currently studying the issue of type inference. We are also considering `rccjava` extensions to support additional synchronization patterns. These extensions may include methods parameterized by locks and support for reader-writer locks.

Availability: We intend to make the `rccjava` prototype implementation available for download from <http://www.research.compaq.com>.

Acknowledgments: Thanks to John Mitchell for comments on an earlier draft of this paper, and to Hannes Marais and Martín Abadi for several useful discussions.

References

- [ACSE99] Jonathan Aldrich, Craig Chambers, Emin Gun Sirer, and Susan Eggers. Static analyses for eliminating unnecessary synchronization from Java programs. In *Proceedings of the Sixth International Static Analysis Symposium*, September 1999.
- [AFM97] Ole Agesen, Stephen N. Freund, and John C. Mitchell. Adding type parameterization to the Java language. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, October 1997.
- [AG98] Alexander Aiken and David Gay. Barrier inference. In *Proceedings of the 25th Symposium on Principles of Programming Languages*, pages 243–354, 1998.

- [ANN97] Torben Amtoft, Flemming Nielson, and Hanne Riis Nielson. Type and behaviour reconstruction for higher-order concurrent programs. *Journal of Functional Programming*, 7(3):321–347, 1997.
- [BH99] Jeff Bogda and Urs Hölzle. Removing unnecessary synchronization in Java. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, November 1999.
- [Bir89] Andrew D. Birrell. An introduction to programming with threads. Research Report 35, Digital Equipment Corporation Systems Research Center, 1989.
- [Bla99] Bruno Blanchet. Escape analysis for object-oriented languages. Application to Java. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, November 1999.
- [BLM96] J. Bank, B. Liskov, and A. Myers. Parameterized Types and Java. Technical Report MIT/LCS/TM-553, Massachusetts Institute of Technology, 1996.
- [BOSW98] Gilad Bracha, Martin Odersky, David Stoutamire, and Philip Wadler. Making the future safe for the past: Adding genericity to the Java programming language. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, October 1998.
- [Car97] Luca Cardelli. Mobile ambient synchronization. Technical Report 1997-013, Digital Systems Research Center, Palo Alto, CA, July 1997.
- [CGS⁺99] J.D. Choi, M. Gupta, M. Serrano, V.C. Sreedhar, and S. Midkiff. Escape analysis for Java. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, November 1999.
- [CJ98] Robert Cartwright and Guy L. Steele Jr. Compatible genericity with run-time types for the Java programming language. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, October 1998.
- [DE97] S. Drossopoulou and S. Eisenbach. Java is type safe — probably. In *European Conference On Object Oriented Programming*, pages 389–418, 1997.
- [DLNS98] David L. Detlefs, K. Rustan M. Leino, Greg Nelson, and James B. Saxe. Extended static checking. Research Report 159, Compaq Systems Research Center, 130 Lytton Ave., Palo Alto, CA 94301, USA, December 1998.
- [FA99a] Cormac Flanagan and Martín Abadi. Object types against races. In *Proceedings of CONCUR*, August 1999.
- [FA99b] Cormac Flanagan and Martín Abadi. Types for safe locking. In *Proceedings of European Symposium on Programming*, March 1999.
- [FKF98] Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. Classes and mixins. In *Proc. 25th ACM Symposium on Principles of Programming Languages*, pages 171–183, January 1998.
- [GJS96] James Gosling, Bill Joy, and Guy Steele. *The Java Language Specification*. Addison-Wesley, 1996.
- [IPW99] Atsushi Igarishi, Benjamin Pierce, and Philip Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, November 1999.
- [Jav98] JavaSoft. Java Developers Kit, version 1.1. <http://java.sun.com>, 1998.
- [JG91] Pierre Jouvelot and David Gifford. Algebraic reconstruction of types and effects. In *Proceedings of the 18th Symposium on Principles of Programming Languages*, pages 303–310, 1991.
- [KM98] Thomas Kistler and Johannes Marais. WebL — a programming language for the web. *Computer Networks and ISDN Systems*, 30:259–270, April 1998.
- [LG88] John M. Lucassen and David K. Gifford. Polymorphic effect systems. In *Proceedings of the ACM Conference on Lisp and Functional Programming*, pages 47–57, 1988.
- [LSS99] K. Rustan M. Leino, James B. Saxe, and Raymie Stata. Checking Java programs via guarded commands. Technical Report 1999-002, Compaq Systems Research Center, Palo Alto, CA, May 1999. Also appeared in Formal Techniques for Java Programs, workshop proceedings. Bart Jacobs, Gary T. Leavens, Peter Muller, and Arnd Poetzsch-Heffter, editors. Technical Report 251, Fernuniversitat Hagen, 1999.
- [Nie96] Flemming Nielson. Annotated type and effect systems. *ACM Computing Surveys*, 28(2):344–345, 1996. Invited position statement for the Symposium on Models of Programming Languages and Computation.
- [NvO98] Tobias Nipkow and David von Oheimb. Java_{light} is type-safe - definitely. In *Proc. 25th ACM Symposium on Principles of Programming Languages*, January 1998.
- [OW97] Martin Odersky and Philip Wadler. Pizza into Java: Translating theory into practice. In *Proc. 24th ACM Symposium on Principles of Programming Languages*, January 1997.
- [SBN⁺97] Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro, and Thomas E. Anderson. Eraser: A dynamic data race detector for multi-threaded programs. *ACM Transactions on Computer Systems*, 15(4):391–411, 1997.
- [Ste93] Nicholas Sterling. Warlock: A static data race analysis tool. In *USENIX Winter Technical Conference*, pages 97–106, 1993.
- [Sym97] Don Syme. Proving Java type soundness. Technical Report 427, University of Cambridge Computer Laboratory Technical Report, 1997.
- [TJ92] Jean-Pierre Talpin and Pierre Jouvelot. Polymorphic type, region and effect inference. *Journal of Functional Programming*, 2(3):245–271, 1992.
- [TT94] Mads Tofte and Jean-Pierre Talpin. Implementation of the typed call-by-value lambda-calculus using a stack of regions. In *Proceedings of the 21st Symposium on Principles of Programming Languages*, pages 188–201, 1994.
- [TT97] Mads Tofte and Jean-Pierre Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [WR99] John Whaley and Martin Rinard. Compositional pointer and escape analysis for Java programs. In *Proceedings of ACM Conference on Object Oriented Languages and Systems*, November 1999.

A The Initial Type System

This appendix presents the type system described in Section 3. We introduce class parameters and thread-local classes in Appendices B and C, respectively.

We first define a number of predicates used in the type system informally. These predicates are based on similar predicates from [FKF98], and we refer the reader to that paper for their precise formulation.

Predicate	Meaning
$ClassOnce(P)$	no class is declared twice in P
$WFClasses(P)$	there are no cycles in the class hierarchy
$FieldsOnce(P)$	no class contains two fields with the same name, either declared or inherited
$MethodsOncePerClass(P)$	no method name appears more than once per class
$OverridesOK(P)$	overriding methods have the same return type, parameter types, and requires set as the method being overridden

A typing environment is defined as

$$E ::= \emptyset \mid E, arg$$

We define the type system using the following judgments.

Judgment	Meaning
$\vdash P : t$	program P yields type t
$P \vdash defn$	$defn$ is a well-formed class definition
$P; E \vdash wf$	E is a well-formed typing environment
$P; E \vdash meth$	$meth$ is a well-formed method
$P; E \vdash field$	$field$ is a well-formed field
$P; E \vdash t$	t is a well-formed type
$P; E \vdash s <: t$	s is a subtype of t
$P; E \vdash defn$	$defn$ is a class defined in P
$P; E \vdash field \in c$	class c declares/inherits $field$
$P; E \vdash meth \in c$	class c declares/inherits $meth$
$P; E \vdash_{final} l : t$	l is a final expression with type t
$P; E \vdash ls$	ls is a well-formed lock set
$P; E \vdash l \in ls$	l appears in ls
$P; E \vdash ls_1 \subseteq ls_2$	lock set ls_1 is contained in ls_2
$P; E; ls \vdash e : t$	expression e has type t

The typing rules for these judgments are presented below.

<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$\vdash P : t$</div> <div style="text-align: center;">[PROG]</div> $\frac{\begin{array}{c} ClassOnce(P) \quad WFClasses(P) \\ FieldsOnce(P) \quad MethodsOncePerClass(P) \\ OverridesOK(P) \\ P = defn_{1\dots n} \ e \quad P \vdash defn_i \quad P; \emptyset; \emptyset \vdash e : t \end{array}}{\vdash P : t}$	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$P \vdash defn$</div> <div style="text-align: center;">[CLASS]</div> $\frac{\begin{array}{c} E = cn \ this \\ P; E \vdash c \\ P; E \vdash field_i \quad P; E \vdash meth_i \end{array}}{P \vdash class \ cn \ extends \ c \ \{ \ field_{1\dots j} \ meth_{1\dots k} \}}$
<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$P; E \vdash wf$</div> <div style="text-align: center;">[ENV EMPTY]</div> $\frac{}{P; \emptyset \vdash wf}$ <div style="text-align: center;">[ENV X]</div> $\frac{P; E \vdash t \quad x \notin Dom(E)}{P; E, t \ x \vdash wf}$	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$P; E \vdash defn$</div> <div style="text-align: center;">[CLASS DEFINITION]</div> $\frac{P; E \vdash wf \quad class \ c \ \dots \in P}{P; E \vdash class \ c \ \dots}$
<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$P; E \vdash t$</div> <div style="text-align: center;">[TYPE C]</div> $\frac{P; E \vdash class \ c \ \dots}{P; E \vdash c}$ <div style="text-align: center;">[TYPE OBJECT]</div> $\frac{P; E \vdash wf}{P; E \vdash Object}$ <div style="text-align: center;">[TYPE INT]</div> $\frac{P; E \vdash wf}{P; E \vdash int}$	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$P; E \vdash t_1 <: t_2$</div> <div style="text-align: center;">[SUBTYPE REFL]</div> $\frac{P; E \vdash t}{P; E \vdash t <: t}$ <div style="text-align: center;">[SUBTYPE CLASS]</div> $\frac{\begin{array}{c} P; E \vdash c_1 <: c_2 \\ P; E \vdash class \ c_2 \ extends \ c_3 \ \dots \end{array}}{P; E \vdash c_1 <: c_3}$
<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$P; E \vdash field \in c$</div> <div style="text-align: center;">[FIELD]</div> $\frac{\begin{array}{c} P; E \vdash_{final} l : c \\ P; E; \emptyset \vdash e : t \end{array}}{P; E \vdash [final]_{opt} \ t \ fd \ guarded.by \ l = e}$	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">$P; E \vdash field \in c$</div> <div style="text-align: center;">[FIELD DECLARED]</div> $\frac{P; E \vdash class \ c \ \dots \ \{ \ \dots \ field \ \dots \ \}}{P; E \vdash field \in c}$ <div style="text-align: center;">[FIELD INHERITED]</div> $\frac{\begin{array}{c} P; E \vdash class \ c \ extends \ c' \ \dots \\ P; E \vdash field \in c' \end{array}}{P; E \vdash field \in c}$

<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">$P; E \vdash \text{method}$</div> <div style="margin-bottom: 5px;">[METHOD]</div> $\frac{P; E \vdash t \quad P; E \vdash ls \quad P; E, \text{arg}_{1..n}; ls \vdash e : t}{P; E \vdash t \text{ mn}(\text{arg}_{1..n}) \text{ requires } ls \{ e \}}$	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">$P; E \vdash \text{meth} \in c$</div> <div style="margin-bottom: 5px;">[METHOD DECLARED]</div> $\frac{P; E \vdash \text{class } c \dots \{ \dots \text{meth} \dots \}}{P; E \vdash \text{meth} \in c}$	<div style="margin-bottom: 5px;">[METHOD INHERITED]</div> $\frac{P; E \vdash \text{class } c \text{ extends } c' \dots \quad P; E \vdash \text{meth} \in c'}{P; E \vdash \text{meth} \in c}$
<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">$P; E \vdash_{\text{final}} e : t$</div> <div style="margin-bottom: 5px;">[FINAL VAR]</div> $\frac{P; E \vdash wf \quad E = E_1, t x, E_2}{P; E \vdash_{\text{final}} x : t}$	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">$P; E \vdash ls$</div> <div style="margin-bottom: 5px;">[LOCK SET]</div> $\frac{P; E \vdash wf \quad \forall l \in ls. \exists c. P; E \vdash_{\text{final}} l : c}{P; E \vdash ls}$	<div style="margin-bottom: 5px;">[FINAL REF]</div> $\frac{P; E \vdash_{\text{final}} e : c \quad P; E \vdash (\text{final } t \text{ fd } \text{guarded_by } l = e') \in c}{P; E \vdash_{\text{final}} e.\text{fd} : t}$
<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">$P; E \vdash l \in ls$</div> <div style="margin-bottom: 5px;">[LOCK SET ELEM]</div> $\frac{l \in ls \quad P; E \vdash ls}{P; E \vdash l \in ls}$	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">$P; E \vdash ls_1 \subseteq ls_2$</div> <div style="margin-bottom: 5px;">[LOCK SET SUBSET]</div> $\frac{P; E \vdash ls_1 \quad P; E \vdash ls_2 \quad ls_1 \subseteq ls_2}{P; E \vdash ls_1 \subseteq ls_2}$	
<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">$P; E \vdash e : t$</div> <div style="margin-bottom: 5px;">[EXP SUB]</div> $\frac{P; E; ls \vdash e : s \quad P; E \vdash s <: t}{P; E; ls \vdash e : t}$	<div style="margin-bottom: 5px;">[EXP NEW]</div> $\frac{P; E \vdash ls \quad P; E \vdash c}{P; E; ls \vdash \text{new } c : c}$	<div style="margin-bottom: 5px;">[EXP VAR]</div> $\frac{P; E \vdash ls \quad E = E_1, t x, E_2}{P; E; ls \vdash x : t}$
<div style="margin-bottom: 5px;">[EXP REF]</div> $\frac{P; E; ls \vdash e : c \quad P; E \vdash ((\text{final})_{\text{opt}} t \text{ fd } \text{guarded_by } l = e') \in c \quad P; E \vdash [e/\text{this}]l \in ls \quad P; E \vdash [e/\text{this}]t}{P; E; ls \vdash e.\text{fd} : [e/\text{this}]t}$	<div style="margin-bottom: 5px;">[EXP ASSIGN]</div> $\frac{P; E; ls \vdash e : c \quad P; E \vdash (t \text{ fd } \text{guarded_by } l = e'') \in c \quad P; E \vdash [e/\text{this}]l \in ls \quad P; E; ls \vdash e' : [e/\text{this}]t}{P; E; ls \vdash e.\text{fd} = e' : [e/\text{this}]t}$	
<div style="margin-bottom: 5px;">[EXP INVOKE]</div> $\frac{P; E; ls_1 \vdash e : c \quad P; E \vdash (t \text{ mn}(s_j \ y_j^{j \in 1..n}) \text{ requires } ls_2 \{ e' \}) \in c \quad P; E; ls_1 \vdash e_j : [e/\text{this}]s_j \quad P; E \vdash [e/\text{this}]ls_2 \subseteq ls_1 \quad P; E \vdash [e/\text{this}]t}{P; E; ls_1 \vdash e.\text{mn}(e_{1..n}) : [e/\text{this}]t}$	<div style="margin-bottom: 5px;">[EXP LET]</div> $\frac{P; E; ls \vdash e_1 : t \quad P; E, t x; ls \vdash e_2 : s \quad P; E \vdash [e_1/x]s}{P; E; ls \vdash \text{let } t x = e_1 \text{ in } e_2 : [e_1/x]s}$	
<div style="margin-bottom: 5px;">[EXP SYNC]</div> $\frac{P; E \vdash_{\text{final}} e_1 : c \quad P; E; ls \cup \{e_1\} \vdash e_2 : t}{P; E; ls \vdash \text{synchronized } e_1 \text{ in } e_2 : t}$	<div style="margin-bottom: 5px;">[EXP FORK]</div> $\frac{P; E \vdash ls \quad P; E; \emptyset \vdash e : t}{P; E; ls \vdash \text{fork } e : \text{int}}$	

B Parameterized Classes

This section extends the type system with classes parameterized by lock expressions. We extend typing environments to include ghost variables:

$$E ::= \emptyset \mid E, \text{arg} \mid E, \text{garg}$$

An instantiated class definition has the form:

$$ci ::= \text{class } c \text{ extends } c \{ \text{field}^* \text{meth}^* \}$$

There is one new judgment form.

Judgment	Meaning
$P; E \vdash ci$	ci is a valid instantiated class definition for P

We redefine well-typed classes to include ghost parameters, and we also introduce new rules for constructing environments and instantiating parameterized classes. Rules that have subscripts in their names, such as [CLASS₂], replace earlier rules of the same name. Rules that do not have subscripts are used in addition to the previous rules.

$P \vdash \text{defn}$ [CLASS₂]

$$\frac{\begin{array}{l} P; \emptyset \vdash t_i \\ \text{garg}_i = \text{ghost } t_i \ x_i \\ E = \text{garg}_{1\dots n}, \text{cn} \langle x_{1\dots n} \rangle \text{ this} \\ P; E \vdash c \\ P; E \vdash \text{field}_i \quad P; E \vdash \text{meth}_i \end{array}}{P \vdash \text{class } \text{cn} \langle \text{garg}_{1\dots n} \rangle \text{ extends } c \{ \text{field}_{1\dots j} \ \text{meth}_{1\dots k} \}}$$

 $P; E \vdash \text{wf}$

[ENV GHOST]

$$\frac{P; E \vdash t \quad x \notin \text{Dom}(E)}{P; E, \text{ghost } t \ x \vdash \text{wf}}$$

 $P; E \vdash_{\text{final}} e : t$ [FINAL VAR₂]

$$\frac{P; E \vdash \text{wf} \quad E = E_1, [\text{ghost}]_{\text{opt}} t \ x, E_2}{P; E \vdash_{\text{final}} x : t}$$

 $P; E \vdash \text{ci}$

[CLASS INSTANTIATION]

$$\frac{\begin{array}{l} \text{class } \text{cn} \langle \text{ghost } t_i \ x_i^{i \in 1\dots n} \rangle \text{ body} \in P \\ P; E \vdash_{\text{final}} l_i : s_i \quad P; E \vdash s_i <: t_i \end{array}}{P; E \vdash \text{class } \text{cn} \langle l_{1\dots n} \rangle [l_i/x_i^{i \in 1\dots n}] \text{body}}$$

C Thread-Local Classes

This section extends the type system with thread-local classes. The following judgments are added to the system:

Judgment	Meaning
$P \vdash t \ \text{shared}$	values of type t can be shared between threads
$P \vdash \text{field} \ \text{shareable}$	field has a thread-shared type and is guarded by a lock or is final
$P; E; c \vdash \text{meth}$	meth does not override a method from any thread-shared super type of c

The typing rules for these judgments are presented below. Rules [EXP REF] and [EXP ASSIGN] must also be updated to support unguarded fields.

 $P \vdash t \ \text{shared}$

[CLASS SHARED]

$$\frac{\text{class } \text{cn} \langle \text{garg}_{1\dots n} \rangle \text{ extends } c \{ \text{field}_{1\dots j} \ \text{meth}_{1\dots k} \} \in P}{P \vdash \text{cn} \langle x_{1\dots n} \rangle \ \text{shared}}$$

[OBJECT SHARED]

$$P \vdash \text{Object} \ \text{shared}$$

[INT SHARED]

$$P \vdash \text{int} \ \text{shared}$$

 $P \vdash \text{field} \ \text{shareable}$

[SHAREABLE GUARDED FIELD]

$$\frac{P \vdash t \ \text{shared}}{P \vdash [\text{final}]_{\text{opt}} t \ \text{fd} \ \text{guarded_by } l = e \ \text{shareable}}$$

[SHAREABLE FINAL FIELD]

$$\frac{P \vdash t \ \text{shared}}{P \vdash \text{final } t \ \text{fd} = e \ \text{shareable}}$$

 $P \vdash \text{defn}$ [CLASS₃]

$$\frac{\begin{array}{l} P; \emptyset \vdash t_i \\ \text{garg}_i = \text{ghost } t_i \ x_i \\ E = \text{garg}_{1\dots n}, \text{cn} \langle x_{1\dots n} \rangle \text{ this} \\ P; E \vdash c \\ P; E \vdash \text{field}_i \quad P; E \vdash \text{meth}_i \\ P \vdash c \ \text{shared} \quad P \vdash \text{field}_i \ \text{shareable} \end{array}}{P \vdash \text{class } \text{cn} \langle \text{garg}_{1\dots n} \rangle \text{ extends } c \{ \text{field}_{1\dots j} \ \text{meth}_{1\dots k} \}}$$

[LOCAL CLASS]

$$\frac{\begin{array}{l} P; \emptyset \vdash t_i \\ \text{garg}_i = \text{ghost } t_i \ x_i \\ E = \text{garg}_{1\dots n}, \text{cn} \langle x_{1\dots n} \rangle \text{ this} \\ P; E \vdash c \\ P; E \vdash \text{field}_i \quad P; E \vdash \text{meth}_i \\ P; E; c \vdash \text{meth}_i \end{array}}{P \vdash \text{thread_local class } \text{cn} \langle \text{garg}_{1\dots n} \rangle \text{ extends } c \{ \text{field}_{1\dots j} \ \text{meth}_{1\dots k} \}}$$

 $P; E; c \vdash \text{meth}$

[OVERRIDE OK]

$$\frac{\forall c'. \left[\begin{array}{l} P; E \vdash c <: c' \\ \wedge P; E \vdash (\dots \text{mn}(\dots) \ \text{requires } \dots) \in c' \end{array} \right] \Rightarrow P \not\vdash c' \ \text{shared}}{P; E; c \vdash \text{mn}(\text{arg}_{1\dots n}) \ \text{requires } \text{ls} \{ e \}}$$

 $P; E \vdash \text{ci}$

[LOCAL CLASS INSTANTIATION]

$$\frac{\begin{array}{l} \text{thread_local class } \text{cn} \langle \text{ghost } t_i \ x_i^{i \in 1\dots n} \rangle \text{ body} \in P \\ P; E \vdash_{\text{final}} l_i : s_i \quad P; E \vdash s_i <: t_i \end{array}}{P; E \vdash \text{class } \text{cn} \langle l_{1\dots n} \rangle [l_i/x_i^{i \in 1\dots n}] \text{body}}$$

 $P; E; \text{ls} \vdash e : t$ [EXP FORK₂]

$$\frac{P; E; \emptyset \vdash e : t \quad \forall x \in \text{FV}(e). P \vdash E(x) \ \text{shared}}{P; E; \text{ls} \vdash \text{fork } e : \text{int}}$$