

Atomizer: A dynamic atomicity checker for multithreaded programs

Cormac Flanagan^{a,*}, Stephen N. Freund^b

^a *Department of Computer Science, University of California at Santa Cruz, Santa Cruz, CA 95064, United States*

^b *Department of Computer Science, Williams College, Williamstown, MA 01267, United States*

Received 12 October 2005; received in revised form 5 January 2007; accepted 9 December 2007

Available online 4 March 2008

Abstract

Ensuring the correctness of multithreaded programs is difficult, due to the potential for unexpected interactions between concurrent threads. Much previous work has focused on detecting race conditions, but the absence of race conditions does not by itself prevent undesired interactions between threads.

A more fundamental noninterference property is *atomicity*. A method is atomic if its execution is not affected by and does not interfere with concurrently-executing threads. Atomic methods can be understood according to their sequential semantics, which significantly simplifies both formal and informal correctness arguments.

This paper presents a dynamic analysis for detecting atomicity violations. This analysis combines ideas from both Lipton's theory of reduction and earlier dynamic race detectors. Experience with a prototype checker for multithreaded Java code demonstrates that this approach is effective for detecting errors due to unintended interactions between threads. In particular, our atomicity checker detects errors that would be missed by standard race detectors. Our experimental results also indicate that the majority of methods in our benchmarks are atomic, indicating that atomicity is a standard methodology in multithreaded programming.

© 2008 Elsevier B.V. All rights reserved.

Keywords: Atomicity; Dynamic analysis; Reduction; Concurrency

1. Reliable multithreaded programming

Multiple threads of control are widely used in software development because they help reduce latency, increase throughput, and provide better utilization of multiprocessor machines. However, reasoning about the behaviour and correctness of multithreaded code is difficult, due to the need to consider all possible interleavings of the executions of the various threads. Thus, methods for specifying and controlling the interference between threads are crucial to the cost-effective development of reliable multithreaded software.

Much previous work on controlling thread interference has focused on *race conditions*. A race condition occurs when two threads simultaneously access the same data variable, and at least one of the accesses is a write [1]. In practice, race conditions are commonly avoided by protecting each data structure with a lock [2]. This lock-based synchronization discipline is supported by a variety of type systems [3–9] and other static [10–13] and dynamic [1,14–17] analyses.

* Corresponding author.

E-mail address: cormac@cs.ucsc.edu (C. Flanagan).

```

public final class StringBuffer {
    ...
    private int count;    // guarded by ‘this’
    private char[] value; // guarded by ‘this’

    public synchronized StringBuffer append(StringBuffer sb) {
        if (sb == null) { sb = NULL; }
        int len = sb.length();
        int newcount = count + len;
        if (newcount > value.length) expandCapacity(newcount);

        // other threads may have changed sb.length(),
        // so len does not reflect the length of sb
        sb.getChars(0, len, value, count);
        count = newcount;
        return this;
    }

    public synchronized int length() { return count; }
    public synchronized void getChars(...) { ... }
    ...
}

```

Fig. 1. Excerpt from `java.lang.StringBuffer`.

Unfortunately, the absence of race conditions is not sufficient to ensure the absence of errors due to unexpected interference between threads. As a concrete illustration of this limitation, consider the excerpt shown in Fig. 1 from the class `java.lang.StringBuffer`. All fields of a `StringBuffer` object are protected by the implicit lock associated with the object, and all `StringBuffer` methods should be safe for concurrent use by multiple threads.

The `append` method shown above first calls `sb.length()`, which acquires the lock `sb`, retrieves the length of `sb`, and releases the lock. The length of `sb` is stored in the variable `len`. At this point, a second thread could remove characters from `sb`. In this situation, `len` is now *stale* [18] and no longer reflects the current length of `sb`, and so the `getChars` method is called with an invalid `len` argument and may throw an exception. Thus, `StringBuffer` objects cannot be safely used by multiple threads, even though the implementation is free of race conditions.

Recent results have shown that subtle defects of a similar nature are common, even in well-tested libraries [19]. Havelund reports finding similar errors in NASA’s Remote Agent spacecraft controller [20], and Burrows and Leino [18] and von Praun and Gross [15] have detected comparable defects in Java applications. Clearly, the construction of reliable multithreaded software requires the development and application of more systematic methods for controlling the interference between concurrent threads.

This paper focuses on a strong yet widely-applicable noninterference property called *atomicity*. A method (or in general a code block) is atomic if for every (arbitrarily interleaved) program execution, there is an equivalent execution with the same overall behaviour where the atomic method is executed *serially*, that is, the method’s execution is not interleaved with actions of other threads.

Atomicity corresponds to a natural programming methodology, essentially dating back to Hoare’s monitors¹ [21]. Many existing classes and library interfaces already follow this methodology, and our experimental results indicate that the vast majority of methods in our benchmarks are atomic.

Atomicity provides a strong, indeed maximal, guarantee of noninterference between threads. This guarantee reduces the challenging problem of reasoning about an atomic method’s behaviour in a *multithreaded* context to the simpler problem of reasoning about the method’s *sequential* behaviour. The latter problem is significantly more amenable to standard techniques such as manual code inspection, dynamic testing and static analysis.

¹ Monitors are less general in that they rely on syntactic scope restrictions and do not support dynamically-allocated shared data.

```

StringBuffer.append is not atomic:
  Atomic block entered
    at StringBuffer.append(StringBuffer.java:445)
    at UnitTest.main(UnitTest.java:21)

Atomic block commits at lock release:
  at StringBuffer.length(StringBuffer.java:144)
  at StringBuffer.append(StringBuffer.java:451)
  at UnitTest.main(UnitTest.java:21)

Atomicity violation at lock acquire:
  at StringBuffer.getChars(StringBuffer.java:326)
  at StringBuffer.append(StringBuffer.java:455)
  at UnitTest.main(UnitTest.java:21)

```

Fig. 2. Atomizer error report.

In summary, atomicity is a widely-applicable and fundamental correctness property of multithreaded code. However, traditional testing techniques are inadequate to verify atomicity. While testing may discover a particular interleaving on which an atomicity violation results in erroneous behaviour, the exponentially-large number of possible interleavings makes obtaining adequate test coverage essentially impossible.

This paper presents a dynamic analysis for detecting atomicity violations. For each code block annotated as being atomic, our analysis dynamically verifies that every execution of that code block is not affected by and does not interfere with other threads. Intuitively, this approach increases the coverage of traditional dynamic testing. Instead of waiting for a particular interleaving on which an atomicity violation causes erroneous behavior, such as a program crash, the checker actively looks for evidence of atomicity violations that may cause errors under other interleavings. Our approach synthesizes ideas from dynamic race detectors (such as Eraser’s *Lockset* algorithm) and Lipton’s theory of reduction (described in Section 3.1). For the `StringBuffer` class described above, our technique detects that `append` contains a window of vulnerability between where the lock `sb` is released inside `length` and then reacquired inside `getChars`, and produces the warning in Fig. 2, even on executions where this window of vulnerability is not exploited by concurrent threads.

We have implemented this dynamic analysis in an automatic checking tool called the *Atomizer*. The application of this tool to over 200,000 lines of Java code demonstrates that it provides an effective approach for detecting defects in multithreaded programs, including some defects that would be missed by existing race-detection tools. In addition, the *Atomizer* avoids false alarms on benign races that do not cause atomicity violations.² Our results also suggest that a large majority of the methods in our benchmarks are atomic, which validates our hypothesis that atomicity is a widely-used programming methodology.

We believe that the application of this dynamic analysis during the development and validation of multithreaded programs may provide multiple benefits, including:

- detecting atomicity violations that are resistant to both traditional testing and existing race detection tools;
- facilitating safe code reuse in multithreaded settings by validating atomicity properties of interfaces;
- simplifying code inspection and debugging, since atomic methods can be understood according to their sequential semantics;
- improving concurrent programming methodology by encouraging programmers to document the atomicity guarantees provided by their code.

In concurrent work, Wang and Stoller [23] have developed several algorithms for checking atomicity dynamically, including the basic algorithm we describe in Section 3.4 as well as more precise but more expensive block-based algorithms. Their original block-based algorithms used an offline, trace-based analysis, but they have more recently explored an online approach [24] similar to ours.

² For simplicity, we assume a sequentially consistent memory model in our prototype. Some such race conditions may not be benign under Java’s relaxed memory model [22].

$$\begin{array}{ll}
u, t \in Tid & \sigma \in GlobalStore = (Var \rightarrow Value) \cup (Lock \rightarrow (Tid \cup \{\perp\})) \\
x \in Var & \pi \in LocalStore \\
v \in Value & \Pi \in LocalStores = Tid \rightarrow LocalStore \\
m \in Lock & \Sigma \in State = GlobalStore \times LocalStores
\end{array}$$

Fig. 3. Domains.

Several static analysis techniques for atomicity have been recently developed, including a type system for atomicity [19,25,26] and the Calvin-R tool [27]. Dynamic atomicity checking complements these static techniques, since most software is validated using a combination of static type checking and dynamic testing. For large, legacy programs, a benefit of the dynamic approach is that it avoids both the overhead of type annotations and the cost of type inference [28].

The presentation of our results proceeds as follows: Section 2 introduces a model of concurrent programs that we use as the basis for our development. Section 3 describes our dynamic analysis for atomicity. Section 4 describes how the Atomizer implements this analysis, and Section 5 presents our experimental results. Section 6 discusses related work, and we conclude with Section 7.

2. Multithreaded programs

We provide a formal basis for reasoning about interference between threads by first formalizing an execution semantics for multithreaded programs. In this semantics, a multithreaded program consists of a number of concurrently executing threads, each of which has an associated thread identifier $t \in Tid$, as defined in Fig. 3. The threads communicate through a global store σ , which is shared by all threads. The global store maps program variables x to values v . The global store also records the state of each lock variable $m \in Lock$. If $\sigma(m) = t$, then the lock m is held by thread t ; if $\sigma(m) = \perp$, then that lock is not held by any thread.

In addition to operating on the shared global store, each thread also has its own local store π containing data not manipulated by other threads, such as the program counter and stack of that thread. A state $\Sigma = (\sigma, \Pi)$ of the multithreaded system consists of a global store σ and a mapping Π from thread identifiers t to the local store $\pi = \Pi(t)$ of each thread. Program execution starts in an initial state $\Sigma_0 = (\sigma_0, \Pi_0)$.

2.1. Standard semantics

We model the behaviour of each thread in a multithreaded program as the transition relation T :

$$T \subseteq Tid \times LocalStore \times Operation \times LocalStore$$

The relation $T(t, \pi, a, \pi')$ holds if the thread t can take a step from a state with local store π , performing the operation $a \in Operation$ on the global store, yielding a new local store π' . The set of possible operations on the global store includes:

- $rd(x, v)$, which reads a value v from a variable x ;
- $wr(x, v)$, which writes a value v to a variable x ;
- $acq(m)$ and $rel(m)$, which acquire and release a lock m , respectively;
- $begin$ and end , which mark the beginning and end of an atomic block; and
- ϵ , the empty operation.

$$a \in Operation ::= rd(x, v) \mid wr(x, v) \mid acq(m) \mid rel(m) \mid begin \mid end \mid \epsilon$$

The relation $\sigma \xrightarrow{a}_t \sigma'$ defined in Fig. 4 models the effect of an operation a by thread t on the global store σ . We use the notation $\sigma[x := v]$ to denote the global store that is identical to σ except that it maps the variable x to the value v .

The transition relation $\Sigma \rightarrow \Sigma'$ performs a single step of an arbitrarily chosen thread (see Fig. 5). We use \rightarrow^* to denote the reflexive-transitive closure of \rightarrow . A transition sequence $\Sigma_0 \rightarrow^* \Sigma$ models the arbitrary interleaving of the various threads of a multithreaded program, starting from the initial state Σ_0 . Although dynamic thread creation is not explicitly supported by the semantics, it can be modelled within the semantics in a straightforward way.

2.2. Serialized semantics

We assume the function $A : LocalStore \rightarrow Nat$ indicates the number of atomic blocks that are currently active, perhaps by examining the program counter and thread stack recorded in the local store. This count should be zero in

$$\begin{array}{c}
\text{[ACT READ]} \\
\frac{\sigma(x) = v}{\sigma \xrightarrow{rd(x,v)}_t \sigma} \\
\\
\text{[ACT ACQUIRE]} \\
\frac{\sigma(m) = \perp}{\sigma \xrightarrow{acq(m)}_t \sigma[m := t]}
\end{array}
\qquad
\begin{array}{c}
\text{[ACT WRITE]} \\
\frac{}{\sigma \xrightarrow{wr(x,v)}_t \sigma[x := v]} \\
\\
\text{[ACT RELEASE]} \\
\frac{\sigma(m) = t}{\sigma \xrightarrow{rel(m)}_t \sigma[m := \perp]}
\end{array}
\qquad
\begin{array}{c}
\text{[ACT OTHER]} \\
\frac{a \in \{begin, end, \epsilon\}}{\sigma \xrightarrow{a}_t \sigma}
\end{array}$$

Fig. 4. Effect of operations: $\sigma \xrightarrow{a}_t \sigma'$.

$$\begin{array}{c}
\text{[STD STEP]} \\
\frac{T(t, \Pi(t), a, \pi') \quad \sigma \xrightarrow{a}_t \sigma'}{(\sigma, \Pi) \rightarrow (\sigma', \Pi[t := \pi'])}
\end{array}$$

Fig. 5. Standard semantics: $\Sigma \rightarrow \Sigma'$.

$$\begin{array}{c}
\text{[SERIAL STEP]} \\
\frac{T(t, \Pi(t), a, \pi') \quad \sigma \xrightarrow{a}_t \sigma' \quad \forall u \neq t. A(\Pi(u)) = 0}{(\sigma, \Pi) \mapsto (\sigma', \Pi[t := \pi'])}
\end{array}$$

Fig. 6. Serialized semantics: $\Sigma \mapsto \Sigma'$.

the initial state, and should only change when entering or leaving an atomic block. We formalize these requirements as follows:

- $A(\Pi_0(t)) = 0$ for all $t \in Tid$;
- if $T(t, \pi, begin, \pi')$ then $A(\pi') = A(\pi) + 1$;
- if $T(t, \pi, end, \pi')$, then $A(\pi) > 0$ and $A(\pi') = A(\pi) - 1$; and
- if $T(t, \pi, a, \pi')$ for $a \notin \{begin, end\}$, then $A(\pi) = A(\pi')$.

The relation $\mathcal{A}(\Pi)$ holds if any thread is inside an atomic block:

$$\mathcal{A}(\Pi) \stackrel{\text{def}}{=} \exists t \in Tid. A(\Pi(t)) \neq 0$$

The *serialized* transition relation \mapsto defined in Fig. 6 is similar to the standard relation \rightarrow , except that a thread cannot perform a step if another thread is inside an atomic block. Thus, the serialized relation \mapsto does not interleave the execution of an atomic block with instructions of concurrent threads.

Reasoning about program behaviour and correctness is much easier under the serialized semantics (\mapsto) than under the standard semantics (\rightarrow), since each atomic block can be understood sequentially, without the need to consider all possible interleaved actions of concurrent threads. However, standard language implementations only provide the standard semantics (\rightarrow), which admits additional transition sequences and behaviors. In particular, a program that behaves correctly according to the serialized semantics may still behave erroneously under the standard semantics. Thus, in addition to being correct with respect to the serialized semantics, the program should also use sufficient synchronization to ensure the atomicity of each block of code that is intended to be atomic. Thus, for any program execution $(\sigma_0, \Pi_0) \rightarrow^* (\sigma, \Pi)$ where $\neg \mathcal{A}(\Pi)$, there should exist an equivalent serialized execution $(\sigma_0, \Pi_0) \mapsto^* (\sigma, \Pi)$. We call this the *atomicity requirement* on program executions, and any execution of a correctly synchronized program should satisfy this requirement. (The restriction $\neg \mathcal{A}(\Pi)$ avoids consideration of partially-executed atomic blocks.)

3. Dynamic atomicity checking

In this section, we present an instrumented semantics that dynamically detects violations of the atomicity requirement. We start by reviewing Lipton's theory of reduction [29], which forms the basis of our approach.




Fig. 7. Reduced execution sequence.

3.1. Reduction

The theory of reduction is based on the notion of right-mover and left-mover actions. An action b is a *right-mover* if, for any execution where the action b performed by one thread is immediately followed by an action c of a concurrent thread, the actions b and c can be swapped without changing the resulting state.

For example, if the operation b is a lock acquire, then the action c of the second thread neither acquires nor releases the lock, and so cannot affect the state of that lock. Hence, the acquire operation can be moved to the right of c without changing the resulting state, and we classify each lock acquire operation as a right-mover.

Conversely, an action c is a *left-mover* if whenever c immediately follows an action b of a different thread, the actions b and c can be swapped, again without changing the resulting state. Suppose the operation c by the second thread is a lock release. During b , the second thread holds the lock, and b can neither acquire nor release the lock. Hence the lock release operation can be moved to the left of b without changing the resulting state, and we classify each lock release operation as a left-mover.

Next, consider an access (read or write) to a variable that is shared by multiple threads. If the variable is protected by some lock that is held whenever the variable is accessed, then two threads can never access the variable at the same time, and we classify each access to that variable as a *both-mover*, which means that it is both a right-mover and a left-mover. If the variable is not consistently protected by some lock, we classify the variable access as a *non-mover*. In summary, we classify operations performed by a thread as follows:

Operation	Mover status
lock acquire	right-mover
lock release	left-mover
access to protected data	both-mover
access to unprotected data	non-mover

To illustrate how the classification of actions as various kinds of movers enables us to verify atomicity, consider the first execution trace shown in Fig. 7. In this trace, a thread:

- (1) acquires a lock m ,
- (2) reads a variable x protected by that lock,
- (3) updates x , and
- (4) then releases m .

The execution path of this thread is interleaved with arbitrary actions b_1, b_2, b_3 of other threads. Because the acquire operation is a right-mover and the write and release operations are left movers, there exists an equivalent serial execution in which the operations of this path are not interleaved with operations of other threads, as illustrated by the diagram in Fig. 7. Thus the execution path is atomic.

More generally, suppose a path through a code block contains a sequence of right-mover actions, followed by at most one non-mover action, followed by a sequence of left-mover actions. Then this path can be *reduced* to an equivalent serial execution, with the same resulting state, where the path is executed without any interleaved actions by other threads.

The non-mover action on a reducible path is called the *commit point* of that path. (If the path does not contain a non-mover action, then the first left-mover action is the commit point.) This commit action thus divides the states of the path into *pre-commit* states (where all preceding actions are right movers) and *post-commit* states (where all succeeding actions are left movers).

$\frac{\begin{array}{l} \text{[INS ACCESS PROTECTED]} \\ a \in \{rd(x, v), wr(x, v)\} \\ P(x) \text{ defined} \\ \sigma(P(x)) = t \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi}$	$\frac{\begin{array}{l} \text{[INS RACE COMMIT]} \\ a \in \{rd(x, v), wr(x, v)\} \\ P(x) \text{ undefined} \\ \varphi(t) = \text{PreCommit} \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi[t := \text{PostCommit}]}$
$\frac{\begin{array}{l} \text{[INS RACE OUTSIDE]} \\ a \in \{rd(x, v), wr(x, v)\} \\ P(x) \text{ undefined} \\ \varphi(t) = \text{Outside} \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi}$	$\frac{\begin{array}{l} \text{[INS ACQUIRE]} \\ \varphi(t) \in \{\text{PreCommit}, \text{Outside}\} \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{acq(m)} \varphi}$
$\frac{\begin{array}{l} \text{[INS RELEASE COMMIT]} \\ \varphi(t) \in \{\text{PreCommit}, \text{PostCommit}\} \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rel(m)} \varphi[t := \text{PostCommit}]}$	$\frac{\begin{array}{l} \text{[INS RELEASE OUTSIDE]} \\ \varphi(t) = \text{Outside} \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rel(m)} \varphi}$
$\frac{\begin{array}{l} \text{[INS ENTER]} \\ A(\Pi(t)) = 0 \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{begin} \varphi[t := \text{PreCommit}]}$	$\frac{\begin{array}{l} \text{[INS NESTED ENTER]} \\ A(\Pi(t)) > 0 \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{begin} \varphi}$
$\frac{\begin{array}{l} \text{[INS EXIT]} \\ A(\Pi(t)) = 1 \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{end} \varphi[t := \text{Outside}]}$	$\frac{\begin{array}{l} \text{[INS NESTED EXIT]} \\ A(\Pi(t)) > 1 \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{end} \varphi}$
$\frac{}{(\sigma, \varphi, \Pi) \Rightarrow_t^c \varphi}$	$\frac{\begin{array}{l} \text{[WRONG ACCESS PROTECTED]} \\ a \in \{rd(x, v), wr(x, v)\} \\ P(x) \text{ defined} \\ \sigma(P(x)) \neq t \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^a \text{wrong}}$
$\frac{\begin{array}{l} \text{[WRONG RACE VIOLATE]} \\ a \in \{rd(x, v), wr(x, v)\} \\ P(x) \text{ undefined} \\ \varphi(t) = \text{PostCommit} \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^a \text{wrong}}$	$\frac{\begin{array}{l} \text{[WRONG ACQUIRE]} \\ \varphi(t) = \text{PostCommit} \end{array}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{acq(m)} \text{wrong}}$

Fig. 8. Instrumented operations: $\Sigma \Rightarrow_t^a \varphi'$ and $\Sigma \Rightarrow_t^a \text{wrong}$.

3.2. Checking atomicity via reduction

We next leverage the theory of reduction to verify atomicity dynamically. In an initial presentation of our approach, we assume the programmer provides a partial function

$$P : \text{Var} \dashrightarrow \text{Lock}$$

that maps protected shared variables to associated locks; if $P(x)$ is undefined, then x is not protected by any lock.

We develop an instrumented semantics that only admits code paths that are reducible, and which goes wrong on irreducible paths. To record whether each thread is in the pre-commit or post-commit part of an atomic block, we extend the state space with an *instrumentation store*:

$$\varphi : \text{Tid} \rightarrow \{\text{PreCommit}, \text{PostCommit}, \text{Outside}\}$$

We use *Outside* to denote when a thread is outside any atomic block. Each state is now a triple (σ, φ, Π) . If $A(\Pi(t)) \neq 0$, then thread t is inside an atomic block, and $\varphi(t)$ indicates whether the thread is in the pre-commit or post-commit part of that atomic block. The initial instrumentation store φ_0 is given by $\varphi_0(t) = \text{Outside}$ for all $t \in \text{Tid}$.

$$\begin{array}{c}
\text{[INS STEP T]} \\
\frac{T(t, \Pi(t), a, \pi') \quad \sigma \xrightarrow{a}_t \sigma' \quad (\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi'}{(\sigma, \varphi, \Pi) \Rightarrow_t (\sigma', \varphi', \Pi[t := \pi'])} \\
\\
\text{[INS WRONG]} \\
\frac{T(t, \Pi(t), a, \pi') \quad \sigma \xrightarrow{a}_t \sigma' \quad (\sigma, \varphi, \Pi) \Rightarrow_t^a \text{wrong}}{(\sigma, \varphi, \Pi) \Rightarrow \text{wrong}} \\
\\
\text{[INS SERIAL STEP]} \\
\frac{(\sigma, \varphi, \Pi) \Rightarrow_t (\sigma', \varphi', \Pi') \quad \forall u \neq t. A(\Pi(u)) = 0}{(\sigma, \varphi, \Pi) \models (\sigma', \varphi', \Pi')}
\end{array}$$


Fig. 9. Instrumented semantics: $\Sigma \Rightarrow_t \Sigma'$, $\Sigma \Rightarrow \Sigma'$, $\Sigma \Rightarrow \text{wrong}$ and $\Sigma \models \Sigma'$.

Fig. 10. Instrumented execution sequence.

The relation $\Sigma \Rightarrow_t^a \varphi'$ in Fig. 8 updates the instrumentation store whenever thread t performs operation a . The rule [INS ACCESS PROTECTED] deals with an access to a protected variable while holding the appropriate lock. This action is a both-mover and so the instrumentation store φ does not change. Accesses to unprotected variables are non-movers, and they cause an atomic block to commit: see [INS RACE COMMIT]. Unprotected accesses are also allowed outside atomic blocks: see [INS RACE OUTSIDE]. Acquire operations are right movers, and they can occur outside or in the pre-commit part of an atomic block, and conversely for release operations.

The relation $\Sigma \Rightarrow_t^a \text{wrong}$ holds if the operation a by thread t would *go wrong* by accessing a protected variable without holding the correct lock [WRONG RACE VIOLATE], or by performing a nonleft-mover action in the left-mover part of an atomic block. Nonleft-mover actions include accessing an unprotected variable [WRONG RACE VIOLATE] or acquiring a lock [WRONG ACQUIRE].

Fig. 9 defines four additional instrumented transition relations:

- the transition relation $\Sigma \Rightarrow_t \Sigma'$ performs an instrumented step of thread t ;
- the transition relation $\Sigma \Rightarrow \Sigma'$ performs an instrumented step of an arbitrary thread;
- the transition relation $\Sigma \Rightarrow \text{wrong}$ holds if a step from Σ could violate the synchronization discipline or the atomicity requirement; and
- the transition relation $\Sigma \models \Sigma'$ is a serialized variant of the instrumented semantics \Rightarrow .

As an illustration of the instrumented semantics, Fig. 10 shows the information tracked by the instrumented semantics for thread 1 during the execution of an atomic fragment of code. In that figure, $\Sigma_i = (\sigma_i, \varphi_i, \Pi_i)$ and in the initial state Σ_0 we have $A(\Pi_0(1)) = 0$. Thus, thread 1 is not initially inside an atomic block. Note that the transition from PreCommit to PostCommit for thread 1 is delayed as long as possible, until the release of the lock, which is the first left-mover operation of the code sequence.

3.3. Correctness

The following theorem states that the instrumented semantics is identical to the standard semantics, except that the instrumented semantics records the additional information φ and may go wrong.

Theorem 1 (Equivalence of Semantics).

- (1) If $(\sigma, \varphi, \Pi) \Rightarrow^* (\sigma', \varphi', \Pi')$, then $(\sigma, \Pi) \rightarrow^* (\sigma', \Pi')$.
 (2) If $(\sigma, \Pi) \rightarrow^* (\sigma', \Pi')$ then $\forall \varphi$ such that either
 (a) $(\sigma, \varphi, \Pi) \Rightarrow^*$ wrong, or
 (b) $\exists \varphi'$ such that $(\sigma, \varphi, \Pi) \Rightarrow^* (\sigma', \varphi', \Pi')$.

Proof. The proof of both parts proceeds by induction over transition sequences and by case analysis on the first transition rule in the sequence.

- (1) Suppose $(\sigma, \varphi, \Pi) \Rightarrow (\sigma', \varphi', \Pi')$. Then, by [INS STEP], we have that $T(t, \Pi(t), a, \pi')$ and $\sigma \rightarrow_t^a \sigma'$ and $\Pi' = \Pi[t := \pi']$. Hence, by [STD STEP], $(\sigma, \Pi) \rightarrow (\sigma', \Pi')$.
 (2) Suppose $(\sigma, \Pi) \rightarrow (\sigma', \Pi')$. Then, by the rule [STD STEP], we have that $T(t, \Pi(t), a, \pi')$ and $\sigma \rightarrow_t^a \sigma'$ and $\Pi' = \Pi[t := \pi']$.

Let $\Sigma = (\sigma, \varphi, \Pi)$. An inspection of the rules in Fig. 8 shows that either $\Sigma \Rightarrow_t^a$ wrong or there exists φ' such that $\Sigma \Rightarrow_t^a \varphi'$. In the former case, $(\sigma, \varphi, \Pi) \Rightarrow$ wrong via [INS WRONG]; in the latter case, $(\sigma, \varphi, \Pi) \Rightarrow (\sigma', \varphi', \Pi')$ as required. \square

In addition, any instrumented execution that does not go wrong satisfies the atomicity requirement.

Theorem 2 (Instrumented Reduction). If $(\sigma_0, \varphi_0, \Pi_0) \Rightarrow^* (\sigma, \varphi, \Pi)$ and $\neg \mathcal{A}(\Pi)$, then $(\sigma_0, \Pi_0) \mapsto^* (\sigma, \Pi)$.

Proof. We first note that any reachable state has the following well-formedness property stating that φ is consistent with the function A :

$$\text{WellFormedState} = \{(\sigma, \varphi, \Pi) \mid A(\Pi(t)) = 0 \Leftrightarrow \varphi(t) = \text{Outside}\}.$$

Clearly, $(\sigma_0, \varphi_0, \Pi_0) \in \text{WellFormedState}$, and it easy to show that well-formedness is preserved by the relation \Rightarrow .

We now introduce three state predicates $Out(t)$, $Pre(t)$, and $Post(t)$, where $Out(t)$ means that thread t is not in an atomic block, and $Pre(t)$ and $Post(t)$ mean that thread t is in the pre-commit and post-commit parts of an atomic block, respectively:

$$Out(t) \stackrel{\text{def}}{=} \{(\sigma, \varphi, \Pi) \in \text{WellFormedState} \mid \varphi(t) = \text{Outside}\}$$

$$Pre(t) \stackrel{\text{def}}{=} \{(\sigma, \varphi, \Pi) \in \text{WellFormedState} \mid \varphi(t) = \text{PreCommit}\}$$

$$Post(t) \stackrel{\text{def}}{=} \{(\sigma, \varphi, \Pi) \in \text{WellFormedState} \mid \varphi(t) = \text{PostCommit}\}$$

We introduce some additional notation to specify properties of these state predicates. For two actions $b, c \subseteq \text{WellFormedState} \times \text{WellFormedState}$, we say that b right-commutes with c if for all $\Sigma_1, \Sigma_2, \Sigma_3$, whenever $(\Sigma_1, \Sigma_2) \in b$ and $(\Sigma_2, \Sigma_3) \in c$, then there exists Σ_2' such that $(\Sigma_1, \Sigma_2') \in c$ and $(\Sigma_2', \Sigma_3) \in b$. The action b left-commutes with the action c if c right-commutes with b . We also define the left restriction $\rho \cdot b$ and the right restriction $b \cdot \rho$ of an action b with respect to a set of states $\rho \subseteq \text{WellFormedState}$.

$$\rho \cdot b \stackrel{\text{def}}{=} \{(\Sigma, \Sigma') \in b \mid \Sigma \in \rho\}$$

$$b \cdot \rho \stackrel{\text{def}}{=} \{(\Sigma, \Sigma') \in b \mid \Sigma' \in \rho\}$$

Using this notation, the following Reduction Theorem formalizes five conditions (A1–A5) that are sufficient to conclude that all atomic blocks are reducible. We next prove that these predicates satisfy the five requirements of the Reduction Theorem, for $t, u \in \text{Tid}$ with $t \neq u$:

- A1. They clearly partition WellFormedState .
 A2. $(Post(t) \cdot \Rightarrow_t \cdot Pre(t))$ is empty, since $\varphi(t)$ is never set to PreCommit while within an atomic block.
 A3. $(\Rightarrow_t \cdot Pre(t))$ right-commutes with \Rightarrow_u , since if $Pre(t)$ holds after an action of thread t , then that action must be by one of the rules [INS ACCESS PROTECTED], [INS ACQUIRE], [INS ENTER], [INS NESTED ENTER], or [INS NESTED EXIT], [INS NO-OP], all of which right-commute with \Rightarrow_u .

- A4. $(Post(t) \cdot \Rightarrow_t)$ left-commutes with \Rightarrow_u , since if $Post(t)$ holds before an action of thread t , then that action must be by one of the rules [INS ACCESS PROTECTED], [INS RELEASE COMMIT], [INS NESTED ENTER], [INS NESTED EXIT], [INS EXIT], or [INS NO-OP], all of which left-commute with \Rightarrow_u .
- A5. if $\Sigma \Rightarrow_t \Sigma'$, then $\Sigma \in Pre(u) \Leftrightarrow \Sigma' \in Pre(u)$ and $\Sigma \in Post(u) \Leftrightarrow \Sigma' \in Post(u)$, since a step by thread t does not change $\varphi(u)$ or $\Pi(u)$.

Hence by [Theorem 3](#) (Reduction), $(\sigma_0, \varphi_0, \Pi_0) \mapsto^* (\sigma, \varphi, \Pi)$, and therefore $(\sigma_0, \Pi_0) \mapsto^* (\sigma, \Pi)$. \square

Theorem 3 (Reduction). *Suppose that for all $t, u \in Tid$ with $t \neq u$:*

- A1. $Pre(t), Post(t)$, and $Out(t)$ form a partition of $WellFormedState$.
- A2. $(Post(t) \cdot \Rightarrow_t \cdot Pre(t))$ is empty.
- A3. $(\Rightarrow_t \cdot Pre(t))$ right-commutes with \Rightarrow_u .
- A4. $(Post(t) \cdot \Rightarrow_t)$ left-commutes with \Rightarrow_u .
- A5. if $\Sigma \Rightarrow_t \Sigma'$, then $\Sigma \in Pre(u) \Leftrightarrow \Sigma' \in Pre(u)$, and $\Sigma \in Post(u) \Leftrightarrow \Sigma' \in Post(u)$.

Suppose further that $\Sigma_0 \Rightarrow^* \Sigma$ and Σ_0 and Σ are in $Out(t)$ for all $t \in Tid$. Then $\Sigma_0 \mapsto^* \Sigma$.

Proof. See [30].

If the instrumented semantics admits a particular execution, then not only is that execution reducible, but many similar executions are also reducible. In particular, when an atomic block is being executed, the only scheduling decision that affects program behaviour is when the commit operation (the transition from `PreCommit` to `PostCommit`) is scheduled. Scheduling decisions regarding when other operations in the atomic block are scheduled are irrelevant, in that they do not affect program behaviour or reducibility. Hence, one test run under our instrumented semantics can simultaneously verify the reducibility of many executions of the standard semantics.

3.4. Inferring protecting locks

The instrumented semantics of the previous section relies on the programmer to specify protecting locks for shared variables. To remove this limitation, we next extend the instrumented semantics to infer protecting locks, using a variant of Eraser’s *Lockset* algorithm [1]. We extend the instrumentation store φ to map each variable x to a set of *candidate locks* for x , such that these candidate locks have all been held on every access to x seen so far:

$$\varphi : (Tid \rightarrow \{PreCommit, PostCommit, Outside\}) \cup (Var \rightarrow 2^{Lock})$$

The initial candidate lock set for each variable is the set of all locks, that is, $\varphi_0(x) = Lock$ for all $x \in Var$.

The relation $\Sigma \Rightarrow_t^a \varphi'$ updates the extended instrumentation store whenever thread t performs operation a on the global store: see [Fig. 11](#). The rule [INS2 ACCESS] for a variable access removes from the variable’s candidate lock set all locks not held by the current thread. We use $H(t, \sigma)$ to denote the set of locks held by thread t in state σ :

$$H(t, \sigma) = \{m \in Lock \mid \sigma(m) = t\}$$

If the candidate lock set for a variable becomes empty, then all accesses to that variable should be treated as non-movers, but previous accesses may already have been incorrectly classified as both-movers. For example, if $\varphi(x) = \{m\}$ when thread t enters the following function `doubleIt`, then the first access to x by thread t will be classified as a both-mover. If, at that point, an action of a concurrent thread causes $\varphi(x)$ to become empty, the analysis will classify the second access to x by t as a non-mover, but will not reclassify the first access, and thus the analysis will fail to recognize that this execution of `doubleIt` may not be reducible.

```

/** atomic */ void doubleIt() {
  synchronized (m) {
    int t = x;
    x = 2 * t;
  }
}

```

$$\begin{array}{c}
\text{[INS2 ACCESS]} \\
\frac{a \in \{rd(x, v), wr(x, v)\} \quad \varphi(x) \cap H(t, \sigma) \neq \emptyset}{(\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi[x := \varphi(x) \cap H(t, \sigma)]} \\
\\
\text{[INS2 RELEASE COMMIT]} \\
\frac{\varphi(t) \in \{\text{PreCommit}, \text{PostCommit}\}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rel(m)} \varphi[t := \text{PostCommit}]} \\
\\
\text{[INS2 ENTER]} \\
\frac{A(\Pi(t)) = 0}{(\sigma, \varphi, \Pi) \Rightarrow_t^{begin} \varphi[t := \text{PreCommit}]} \\
\\
\text{[INS2 EXIT]} \\
\frac{A(\Pi(t)) = 1}{(\sigma, \varphi, \Pi) \Rightarrow_t^{end} \varphi[t := \text{Outside}]} \\
\\
\text{[INS2 NO-OP]} \\
\frac{}{(\sigma, \varphi, \Pi) \Rightarrow_t^c \varphi} \\
\\
\text{[WRONG2 RACE]} \\
\frac{a \in \{rd(x, v), wr(x, v)\} \quad \varphi(x) \cap H(t, \sigma) = \emptyset}{(\sigma, \varphi, \Pi) \Rightarrow_t^a \text{wrong}} \\
\\
\text{[INS2 ACQUIRE]} \\
\frac{\varphi(t) \in \{\text{PreCommit}, \text{Outside}\}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{acq(m)} \varphi} \\
\\
\text{[INS2 RELEASE OUTSIDE]} \\
\frac{\varphi(t) = \text{Outside}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rel(m)} \varphi} \\
\\
\text{[INS2 NESTED ENTER]} \\
\frac{A(\Pi(t)) > 0}{(\sigma, \varphi, \Pi) \Rightarrow_t^{begin} \varphi} \\
\\
\text{[INS2 NESTED EXIT]} \\
\frac{A(\Pi(t)) > 1}{(\sigma, \varphi, \Pi) \Rightarrow_t^{end} \varphi} \\
\\
\text{[WRONG2 ACQUIRE]} \\
\frac{\varphi(t) = \text{PostCommit}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{acq(m)} \text{wrong}}
\end{array}$$

Fig. 11. Instrumented operations 2: $\Sigma \Rightarrow_t^a \varphi'$ and $\Sigma \Rightarrow_t^a \text{wrong}$.

$$\begin{array}{c}
\text{[INS2 STEP]} \\
\frac{T(t, \Pi(t), a, \pi') \quad \sigma \xrightarrow_t^a \sigma' \quad (\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi'}{(\sigma, \varphi, \Pi) \Rightarrow (\sigma', \varphi', \Pi[t := \pi'])} \\
\\
\text{[INS2 WRONG]} \\
\frac{T(t, \Pi(t), a, \pi') \quad \sigma \xrightarrow_t^a \sigma' \quad (\sigma, \varphi, \Pi) \Rightarrow_t^a \text{wrong}}{(\sigma, \varphi, \Pi) \Rightarrow \text{wrong}}
\end{array}$$

Fig. 12. Instrumented semantics 2: $\Sigma \Rightarrow \Sigma'$ and $\Sigma \Rightarrow \text{wrong}$.

Thus, to ensure soundness, the lock inference semantics does not support unprotected variables and instead requires every variable to have a protecting lock. If the candidate lock set becomes empty, then that state goes wrong, via [WRONG2 RACE].

The relation $\Sigma \Rightarrow \Sigma'$ in Fig. 12 performs an instrumented step (with lock inference) of an arbitrarily chosen thread; the relation $\Sigma \Rightarrow \text{wrong}$ describes states that go wrong.

Like the previous instrumented semantics (\Rightarrow), the lock-inference semantics (\Rightarrow) is equivalent to the standard semantics (\rightarrow) except that it only admits execution sequences that satisfy the atomicity requirement. The following two theorems formalize these correctness properties. Their proofs are analogous to those of Theorems 1 and 2.

Theorem 4 (Equivalence of Semantics 2).

- (1) If $(\sigma, \varphi, \Pi) \Rightarrow^* (\sigma', \varphi', \Pi')$, then $(\sigma, \Pi) \rightarrow^* (\sigma', \Pi')$.
- (2) If $(\sigma, \Pi) \rightarrow^* (\sigma', \Pi')$ then $\forall \varphi$ either
 - (a) $(\sigma, \varphi, \Pi) \Rightarrow^* \text{wrong}$, or
 - (b) $\exists \varphi'$ such that $(\sigma, \varphi, \Pi) \Rightarrow^* (\sigma', \varphi', \Pi')$.

Theorem 5 (Instrumented Reduction 2). If $(\sigma_0, \varphi_0, \Pi_0) \Rightarrow^* (\sigma, \varphi, \Pi)$ and $\neg \mathcal{A}(\Pi)$, then $(\sigma_0, \Pi_0) \mapsto^* (\sigma, \Pi)$.

Again, if the instrumented semantics admits a particular execution, then all executions that are equivalent to that execution modulo irrelevant scheduling decisions are reducible.

4. Implementation

We have developed an implementation, called the *Atomizer*, of the dynamic analysis outlined in the previous section. The Atomizer takes as input a multithreaded Java [31] program and rewrites the program to include additional instrumentation code. This instrumentation code calls methods of the Atomizer runtime that implement the Lockset and reduction algorithms and issue warning messages when atomicity violations are detected.

The Atomizer performs the instrumentation on Java source code. This approach has a number of advantages: it supports programmer-supplied annotations, it works at the high level of abstraction of the Java language, and it is portable across Java virtual machines. This approach does require source code, but the instrumentation could also be performed at the bytecode level.

The target program can include annotations in comments to indicate that a method is atomic, as in

```
/*# atomic */ void getChars() {...}
```

The `/*# assume_atomic */` and `/*# assume_mover */` annotations can be used to indicate that the Atomizer should assume that a specific method is atomic or a both-mover, respectively, without checking this requirement. The annotation `/*# assume_guarded */` on a field indicates that the Atomizer should assume there are no simultaneous access race conditions on that field; any potential race conditions detected by the Lockset algorithm are assumed to be false alarms and are ignored. This annotation is useful when the Lockset algorithm does not properly handle the observed access pattern, such as when data local to a thread is transferred to another thread through a global queue [1]. To facilitate suppressing false alarms, the `/*# no_warn */` annotation turns off all Lockset and reduction checking for a single line of code.

When checking code linked with uninstrumented libraries, the Atomizer assumes that all methods defined in the uninstrumented code are movers by default. Alternatively, some or all of the library source code may be instrumented by the Atomizer.

In order to reduce the burden of annotating large programs, the Atomizer also provides two built-in heuristics for deciding which blocks should be checked for atomicity:

- (1) *Export+Synch*: the first heuristic is that (1) all methods *exported* by classes should be atomic, and (2) all synchronized blocks and synchronized methods should be atomic. Exported methods are those that are public or package protected. However, this heuristic is not used for `main` and the `run` methods of `Runnable` objects, because these methods typically are not atomic.
- (2) *Synch*: the second heuristic is to ignore exported methods and only assume that all synchronized blocks and synchronized methods should be atomic.

Although these heuristic are quite simple, they provide a reasonable starting point for identifying atomicity errors in unannotated code. These two heuristics are experimentally compared in the following section.

In the rest of this section, we describe our Lockset and reduction implementation, demonstrate how the tool identifies and reports errors, and present several improvements to the basic algorithm.

4.1. Lockset algorithm

For each field of each allocated object, the Atomizer tracks a *state* that reflects the degree to which the field has been shared among multiple threads.

The possible states of our algorithm, as shown in Fig. 13, are similar to the states in earlier race detectors [1,32]:

- *Thread-Local*: The field has only been accessed by the object's creating thread.
- *Thread-Local (2)*: Ownership has transferred to a second thread, and the field is no longer accessed by the creating thread. This state supports common initialization patterns in Java [32].
- *Read-Shared*: The field has been read, but not written, by multiple threads.
- *Shared-Modified*: The field has been read and written by multiple threads, and a candidate lock set records which locks have been consistently held when accessing this field. When entering this state, the candidate set is initialized with all locks held by the current thread.

Some of these extensions introduce some degree of unsoundness [1,32] into the algorithm. However, we do not believe these extensions miss a large number of errors, and they substantially reduce the false alarm rate.




Fig. 13. Lockset algorithm states for each allocated field.

$$\begin{array}{c}
 \text{[INS2 RACE OUTSIDE]} \\
 a \in \{rd(x, v), wr(x, v)\} \\
 \varphi(x) \cap H(t, \sigma) = \emptyset \\
 \varphi(t) = \text{Outside} \\
 \hline
 (\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi[x := \emptyset]
 \end{array}
 \qquad
 \begin{array}{c}
 \text{[INS2 RACE COMMIT]} \\
 a \in \{rd(x, v), wr(x, v)\} \\
 \varphi(x) \cap H(t, \sigma) = \emptyset \\
 \varphi(t) = \text{PreCommit} \\
 \hline
 (\sigma, \varphi, \Pi) \Rightarrow_t^a \varphi[t := \text{PostCommit}, x := \emptyset]
 \end{array}$$

$$\begin{array}{c}
 \text{[WRONG2 RACE]} \\
 a \in \{rd(x, v), wr(x, v)\} \\
 \varphi(x) \cap H(t, \sigma) = \emptyset \\
 \varphi(t) = \text{PostCommit} \\
 \hline
 (\sigma, \varphi, \Pi) \Rightarrow_t^a \text{wrong}
 \end{array}$$

Fig. 14. Relaxed instrumentation: $\Sigma \Rightarrow_t^a \varphi'$ and $\Sigma \Rightarrow_t^a \text{wrong}$.

4.2. Reduction algorithm

The instrumented semantics for lock inference in Section 3.4 goes wrong on any race condition. Since programs frequently have benign races, the Atomizer implements a relaxed version of this semantics that accommodates such benign race conditions. If the candidate lock set for a variable becomes empty, then subsequent accesses to that variable are considered non-movers. Note that previous accesses to that variable, which were earlier classified as movers, will not be reclassified as non-movers, since storing a history of all variable accesses would be expensive. Thus, as mentioned in Section 3.4, these relaxed rules introduce a degree of unsoundness. We believe this unsoundness rarely causes the Atomizer to miss atomicity violations in practice because it requires an unlucky scheduling of operations and because the Atomizer will report the problem on the next execution of the nonatomic code fragment. The rules in Fig. 14 adapt the relations $\Sigma \Rightarrow_t^a \varphi'$ and $\Sigma \Rightarrow_t^a \text{wrong}$ to express this relaxed semantics.

To produce clear error messages like that in Section 1, the Atomizer can optionally capture stack traces (in the form of Exception objects) at the entry and commit points of each atomic block, and include these stack traces in error messages. Since the Atomizer supports nested atomic blocks, a single operation could result in multiple atomicity violations.

4.3. Redundant locking

The Atomizer may produce false alarms due to imprecisions in the Lockset and reduction algorithms. We next present several improvements that eliminate many of these false alarms. We start by revisiting the treatment of synchronization operations during reduction. The classification of lock acquires and releases as right-movers and left-movers, respectively, is correct but overly conservative in some cases. In particular, modular programs typically include redundant synchronization operations that we can more precisely characterize as both-movers.

- *Re-entrant locks.* Lock acquires are in general only right-movers and not left-movers. However, Java provides re-entrant locks, and a re-entrant lock acquire is a both-mover, because this operation cannot interact with other threads. Similarly, a re-entrant release is also a both-mover.
- *Thread-local locks.* If a lock is used by only a single thread, acquires and releases of that lock are both-movers.
- *Thread-local (2) locks.* Adding another *Thread-local* state, as in our Lockset algorithm, eliminates false alarms caused by initialization patterns in which one thread creates and initializes a protected object, and then transfers ownership of both the object and its protecting lock to another thread.
- *Protected locks.* Suppose each thread always holds some lock m_1 before acquiring lock m_2 . In this case, two threads cannot attempt to acquire m_2 simultaneously, and so operations on the lock m_2 are also both-movers.

In essence, these improvements lead us to treat lock operations in a similar fashion to field accesses. The only major difference is that locks do not have a notion of *Read-Shared*. The *Thread-local* and *Thread-local (2)* extensions are unsound for reasons similar to the analogous states in the race detector described in Section 4.1.

4.4. Write-protected data

Consider the following two methods, in which the variable x is protected by a lock for all writes, but not protected for reads.

```

    /**# atomic */ int read() {           /**# atomic */ void inc() {
        return x;                         synchronized (lock) {
    }                                       x = x+1;
                                           }
                                           }

```

If x is a 32-bit variable, then the `read()` method is atomic on a sequentially-consistent machine, even though no protecting lock is held. Despite the presence of such unprotected reads, the `inc()` method is also atomic. In particular, when the lock is held, a read of x is a both-mover, since no other thread can write to x without holding the lock.

To handle examples like this one, we use a variant of the Lockset algorithm that implies a pair of lock sets for each program variable. That is, we extend the instrumentation store φ to map each variable x to two lock sets:

- (1) an *access-protecting* lock set $\varphi(x, A)$, which contains locks held on every access (read or write) to that field, and
- (2) a *write-protecting* lock set $\varphi(x, W)$, which contains locks held on every write to that field.

The instrumentation store now has the type

$$\varphi : (Tid \rightarrow \{\text{PreCommit, PostCommit, Outside}\}) \cup (Var \times \{A, W\} \rightarrow 2^{Lock})$$

The access-protecting lock set $\varphi(x, A)$ is always a subset of the write-protecting lock set $\varphi(x, W)$. Both lock sets initially contain all locks, that is, $\varphi_0(x, A) = \varphi_0(x, W) = Lock$.

The new relation $\Sigma \Rightarrow_t^a \varphi'$ shown in Fig. 15 implies these lock-set pairs and uses them to more precisely determine which operations are movers and which are non-movers. A field read is a both-mover if the current thread holds at least one of the write-protecting locks (see [INS3 READ]); otherwise the read is a non-mover and may be a benign race condition (see [INS3 READ RACE]) or may cause an atomicity violation (see [WRONG3 READ]). In contrast, a field write is a both-mover only if the access-protecting lock set is nonempty (see [INS3 WRITE]); otherwise the write is a non-mover (see [INS3 WRITE RACE] and [WRONG3 WRITE]).

Using the new instrumented semantics, the Atomizer can imply that `inc()` is atomic, since it consists of a right-mover (the lock acquire); a both-mover (the read of x); an atomic action (since the write of x does not commute with concurrent reads of other threads); and a left-mover (the lock release). In comparison, existing race-detection tools would produce a warning about the race condition in `read()`, even though this race condition is benign and does not affect the atomicity of either method.

$\frac{[\text{INS3 READ}]}{\frac{\varphi(x, W) \cap H(t, \sigma) \neq \emptyset}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rd(x,v)} \varphi[(x, A) := \varphi(x, A) \cap H(t, \sigma)]}}$	$\frac{[\text{INS3 READ RACE OUTSIDE}]}{\frac{\varphi(x, W) \cap H(t, \sigma) = \emptyset}{\varphi(t) = \text{Outside}}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rd(x,v)} \varphi[(x, A) := \varphi(x, A) \cap H(t, \sigma)]}}$
$\frac{[\text{INS3 READ RACE COMMIT}]}{\frac{\varphi(x, W) \cap H(t, \sigma) = \emptyset}{\varphi(t) = \text{PreCommit}}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rd(x,v)} \varphi[(x, A) := \varphi(x, A) \cap H(t, \sigma), t := \text{PostCommit}]}}$	$\frac{[\text{INS3 WRITE}]}{\frac{\varphi(x, A) \cap H(t, \sigma) \neq \emptyset}{(\sigma, \varphi, \Pi) \Rightarrow_t^{wr(x,v)} \varphi[(x, W) := \varphi(x, W) \cap H(t, \sigma), (x, A) := \varphi(x, A) \cap H(t, \sigma)]}}$
$\frac{[\text{INS3 WRITE RACE OUTSIDE}]}{\frac{\varphi(x, A) \cap H(t, \sigma) = \emptyset}{\varphi(t) = \text{Outside}}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{wr(x,v)} \varphi[(x, W) := \varphi(x, W) \cap H(t, \sigma), (x, A) := \emptyset]}}$	$\frac{[\text{INS3 WRITE RACE COMMIT}]}{\frac{\varphi(x, A) \cap H(t, \sigma) = \emptyset}{\varphi(t) = \text{PreCommit}}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{wr(x,v)} \varphi[(x, W) := \varphi(x, W) \cap H(t, \sigma), (x, A) := \emptyset, t := \text{PostCommit}]}}$
$\frac{[\text{INS3 ACQUIRE}]}{\frac{\varphi(t) \in \{\text{PreCommit}, \text{Outside}\}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{acq(m)} \varphi}}$	$\frac{[\text{INS3 NO-OP}]}{(\sigma, \varphi, \Pi) \Rightarrow_t^c \varphi}$
$\frac{[\text{INS3 RELEASE OUTSIDE}]}{\frac{\varphi(t) = \text{Outside}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rel(m)} \varphi}}$	$\frac{[\text{INS3 RELEASE COMMIT}]}{\frac{\varphi(t) \in \{\text{PreCommit}, \text{PostCommit}\}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rel(m)} \varphi[t := \text{PostCommit}]}}$
$\frac{[\text{INS3 ENTER}]}{\frac{A(\varphi(t)) = 0}{(\sigma, \varphi, \Pi) \Rightarrow_t^{begin} \varphi[t := \text{PreCommit}]}}$	$\frac{[\text{INS3 NESTED ENTER}]}{\frac{A(\varphi(t)) > 0}{(\sigma, \varphi, \Pi) \Rightarrow_t^{begin} \varphi}}$
$\frac{[\text{INS3 EXIT}]}{\frac{A(\varphi(t)) = 1}{(\sigma, \varphi, \Pi) \Rightarrow_t^{end} \varphi[t := \text{Outside}]}}$	$\frac{[\text{INS3 NESTED EXIT}]}{\frac{A(\varphi(t)) > 1}{(\sigma, \varphi, \Pi) \Rightarrow_t^{end} \varphi}}$
$\frac{[\text{WRONG3 READ RACE VIOLATE}]}{\frac{\varphi(x, W) \cap H(t, \sigma) = \emptyset}{\varphi(t) = \text{PostCommit}}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{rd(x,v)} \text{wrong}}$	$\frac{[\text{WRONG3 WRITE RACE VIOLATE}]}{\frac{\varphi(x, A) \cap H(t, \sigma) = \emptyset}{\varphi(t) = \text{PostCommit}}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{wr(x,v)} \text{wrong}}$
$\frac{[\text{WRONG3 ACQUIRE}]}{\frac{\varphi(t) = \text{PostCommit}}{(\sigma, \varphi, \Pi) \Rightarrow_t^{acq(m)} \text{wrong}}}$	

Fig. 15. Instrumented operations 3: $\Sigma \Rightarrow_t^a \varphi'$ and $\Sigma \Rightarrow_t^a \text{wrong}$.

5. Evaluation

This section summarizes our experience applying the Atomizer to the following twelve benchmark programs:

- `elevator`, a discrete event simulator for elevators [15] configured with the small default data file;
- `hedc`, a tool to access astrophysics data from Web sources [15] configured with the default configuration;
- `tsp`, a Travelling Salesman Problem solver [15] configured to solve the problem for the `map16c` input file with eight threads;
- `sor`, a scientific computing program [15] configured to perform 50 iterations of computation with two threads;
- `mtrt`, a multithreaded ray-tracing program from the SPEC JVM98 benchmark suite [33];
- `jigsaw`, an open source web server [34] configured to serve a fixed number of pages to a crawler;
- `specJBB`, the SPEC JBB2000 business object simulator [33] modified to process a fixed number of transactions;
- `moldyn`, `montecarlo`, and `raytracer` from the Java Grande benchmark suite [35] (all configured to use four threads and the small problem size);

Benchmark	Lines	Num. Threads	Num. Locks	Max. Locks Held	Num. Lock Set Pairs	Base Time (s)	Slowdown
elevator	529	5	8	1	17	11.14	—
hedc	29,948	26	385	3	728	8.36	—
tsp	706	10	2	1	5	0.94	48.2
sor	17,690	4	1	1	2	0.70	7.3
modyn	1,291	5	1	1	2	3.62	11.8
montecarlo	3,557	5	1	1	2	7.94	2.2
raytracer	1,859	5	5	1	7	5.96	36.6
mtrt	11,315	6	7	2	7	2.33	46.4
jigsaw	90,100	53	706	31	4,531	13.49	4.7
specJBB	30,490	10	262,000	6	340,088	18.01	11.2
webl	22,284	5	402,445	3	452,685	60.35	—
lib-java	75,305	39	816,617	6	986,855	96.50	—

Fig. 16. Summary of test programs and performance, when configured with the *Export+Synch* heuristic.

- `webl`, a scripting language interpreter for processing web pages, configured to execute a simple web crawler [36];
- `lib-java`, an uninstrumented test harness (comprised of `webl`, `jbb`, and `hedc`) that tests an instrumented version of the standard Java libraries `java.lang`, `java.io`, `java.net`, and `java.util`. (All programs other than `lib-java` use uninstrumented libraries.)

The Atomizer instrumented these programs using both the *Synch* and *Export+Synch* heuristics described in Section 4. To ensure that our measurements would accurately reflect the cost of the underlying analysis, for these tests the Atomizer did not record stack histories for atomic block entry and commit points. We performed the experiments on a Red Hat Linux 8.0 computer with dual 3.06 GHz Pentium 4 Xeon processors and 2 GB of memory. We used the Sun JDK 1.4.2 compiler and virtual machine for all benchmarks except `lib-java`, for which we used the Sun JDK 1.3.1 virtual machine due to compatibility problems.

Fig. 16 presents several statistics for the test programs using all the extensions from Sections 4.3 and 4.4 and the *Export+Synch* heuristic. The “Num. Threads” column shows the number of threads created during execution; the “Num. Locks” column shows the number of distinct locks acquired during execution; and the “Max. Locks Held” column shows the maximum number of locks held at the same time by any thread during execution. The number of locks and distinct lock set pairs were relatively small for most programs, although the larger programs used many objects as locks, in some cases several orders of magnitude more than in comparably-sized C programs [1].

The slowdown incurred by the instrumentation varied from 2.2 x to roughly 50 x . We only report slowdowns for compute-bound programs. Those programs with very little slowdown, such as `sor` and `montecarlo`, spent most of the time in uninstrumented library code. We believe that slowdowns of 20 x –40 x are representative for most programs. We did not focus on efficiency in this prototype, however, and there is much room for improvement. In particular, static analyses have reduced the overhead of dynamic race detection to under 50% [15], which suggests that similar performance could be achieved when checking atomicity.


Fig. 18 illustrates the benefit of the various analysis extensions described in Sections 4.3 and 4.4, under the *Export+Synch* heuristic. The “Basic” column indicates the number of warnings reported for each program using the basic Lockset and reduction algorithms. The succeeding columns show the number of warnings as each refinement to the algorithm is added. Cumulatively, these five refinements are quite effective: they reduce the number of warnings by roughly 70% (from 341 to 97).

We next consider the behaviour of the Atomizer with all of these extensions enabled. Fig. 17 presents the number of warnings and errors identified by the Atomizer under both the *Export+Synch* and *Synch* heuristics. For each heuristic, the “Warnings” column reports the number of atomic blocks that failed the Atomizer’s atomicity requirements, and the “Errors” column reports the number of these warnings that we consider errors, either because they lead to undesirable program behaviour or because they violate documented atomicity properties. Despite checking only mature software, the Atomizer identified a number of potentially damaging errors, which we discuss in more detail below.

On the `lib-java` benchmark, the Atomizer detected an atomicity violation in the synchronized method `PrintStream.println(String s)`, which uses two method calls to write the string `s` and the following new-line

Benchmark	<i>Export+Synch</i> Heuristic		<i>Synch</i> Heuristic	
	Warnings	Errors	Warnings	Errors
elevator	2	0	0	0
hedc	4	1	1	0
tsp	7	0	7	0
sor	0	0	0	0
moldyn	0	0	0	0
montecarlo	1	0	0	0
raytracer	1	1	1	1
mtrt	6	0	0	0
jigsaw	34	1	16	0
specJBB	4	0	0	0
webl	19	0	0	0
lib-java	19	4	4	1

Fig. 17. Summary of atomizer warnings.

Fig. 18. Warnings reported by the Atomizer under different configurations, using the *Export+Synch* heuristic.

character to a stream stored in the instance variable `out`. Another thread could concurrently write to `out`, thus potentially corrupting the output stream. Since the `println` method is synchronized, this error was detected under both heuristics. Notably, this error was not caused by a race condition, and so would not be detected by a race detector. A comparable error in `PrintWriter` had been previously identified by a static type system for atomicity [19], but that system requires a significant number of programmer-inserted annotations. Other errors in `lib-java` include known problems with iterators for collections (such as `Hashtable`) in the presence of concurrent modifications [23].

On the `jigsaw` benchmark, the Atomizer detected an error in the method `ResourceStoreManager.loadResourceStore`, where a specific interleaving could allow an entry to be added to a resource store after the store had been closed as part of the shutdown process. This error was only detected under the *Export+Synch* heuristic, since this method is not synchronized. As before, this error was not caused by race conditions, and would not be detected by a race detector. This particular error was previously detected using a static *view consistency* analysis [15].

On the `raytracer` benchmark, the Atomizer detected an atomicity violation in the `JGFRayTracerBench` class, which was caused by race conditions on a checksum field. Similarly, on the `hedc` benchmark, the Atomizer detected an atomicity error caused by race conditions on a `java.util.Calendar` object that was improperly shared among multiple threads.

In most programs, the warnings that did not indicate defects could be suppressed by inserting a handful of annotations. A significant number of false alarms were due to the overly-optimistic heuristics employed to identify

atomic blocks. Identifying atomic blocks via explicit `/** atomic */` declarations instead of using these heuristics would produce fewer false alarms. For example, atomicity violations were often reported on methods called near the top-level entry points of the program (the `main` and `run` methods), but many such methods are not intended to be atomic and would not be labeled as atomic by a programmer.

Other common sources of false alarms include double-checked locking patterns, lazy initialization patterns, and various caching idioms. These programming idioms are notoriously problematic for analysis tools based on race detection and are discussed in more detail in [16]. Although some of these practices, such as double-checked locking, are incompatible with the Java's relaxed memory model specification [22], we classify them as false alarms since they do not cause problems in most current Java environments [16].

During these tests, the Atomizer also recognized five fields with benign race conditions that did not lead to atomicity violations (under our simplifying assumption that the memory model is sequentially consistent). The Atomizer does not report spurious warnings for these benign races.

Overall, the Atomizer found no potential atomicity violations in over 90% of the methods annotated as atomic that were exercised during our test runs. These statistics suggest that atomicity is a fundamental design principle in many multithreaded systems, especially library classes and reusable application components.

6. Related work

Lipton [29] first proposed reduction as a way to reason about concurrent programs without considering all possible interleavings. He focused primarily on checking deadlock freedom. Doepfner [37], Back [38], and Lamport and Schneider [39] extended this work to allow proofs of general safety properties. Cohen and Lamport [40] extended reduction to allow proofs of liveness properties. Misra [41] has proposed a reduction theorem for programs built with monitors [21] communicating via procedure calls.

Eraser [1] introduced the Lockset algorithm for dynamic race detection. This approach has been extended to object-oriented languages [32] and has been improved for precision and performance [14,17]. O'Callahan and Choi [16] recently combined the Lockset algorithm with a happens-before analysis to reduce false alarms in a dynamic race detector for Java programs.

A number of static race detectors have also been developed. Warlock [10] is a static race detection system for C programs. ESC/Java [11] statically catches a variety of software defects, including race conditions. Other approaches for static race and deadlock prevention are discussed in earlier papers [4,5,3]; these include model checking [12,13,42], dataflow analysis [43], abstract interpretation [44], and type systems for process calculi [45,46].

In previous work, we produced a type system [3] that prevents violations of the lock-based synchronization discipline. Since then, similar type systems have been developed that include a notion of object ownership [7], and that target other languages such as Cyclone [9], a type-safe variant of C. Compared to dynamic techniques, these static type systems provide stronger soundness guarantees and detect errors earlier in the development cycle, but require more effort from programmer. Type inference techniques have also been developed for such type systems [28].

Agarwal and Stoller [47] have developed a dynamic analysis tool that computes many of the annotations necessary for applying a type-based static analysis [7], thereby reducing the overhead of applying static tools to large, unannotated legacy systems. We believe that the Atomizer can be extended to output annotations regarding the locking discipline and the atomicity requirements for methods in much the same way.

While some of these race detection tools have been quite effective, they may fail to detect atomicity violations and may yield false alarms on benign race conditions that do not violate atomicity.

Bacon et al. developed Guava [48], an extension to the Java language with a form of monitor capable of sharing object state in a way that prevents race conditions. The Atomizer would work very well for languages like Guava, since language-enforced race freedom would eliminate several common sources of false alarms observed while checking programs written in languages that permit races.

In recent work, Flanagan and Qadeer developed a static type system to verify atomicity in Java programs [30,19]. In comparison to the Atomizer, the type system provides better coverage and soundness guarantees, but is less expressive (for example, it does not fully support redundant locking). The type system also requires programmer-inserted annotations that specify properties such as the locking discipline followed by the program. The expressiveness of this type system has been extended by exploiting the notion of pure expressions [25]. We have recently explore

type inference for this system [26] as a way to eliminate some of the overhead of using type-based techniques. In similar work, Sasturkar et al. [49] have explored static atomicity inference combined with a runtime race-condition analysis.

This type system for atomicity was inspired by the Calvin-R [27] static checking tool for multithreaded programs. Calvin-R supports modular verification of multithreaded programs by annotating each procedure with a specification; this specification is related to the procedure implementation via abstraction relation that combines the notions of simulation and reduction. In ongoing work, the notions of reduction and atomicity are used by Qadeer et al. [50] to infer concise procedure summaries in an analysis for multithreaded programs.

An alternative approach for verifying atomicity using model checking is being explored by Hatcliff et al. [51]. In addition to using Lipton’s theory of reduction, they also investigate an approach based on partial order reductions. Their experimental results suggest that the model-checking approach for verifying atomicity is feasible for unit testing, where the reachable state space is smaller than in integration testing. A more general (but more expensive) technique for verifying atomicity during model checking is *commit-atomicity* [52].

Wang and Stoller have developed an alternative *block-based* approach to verifying atomicity. In comparison to our approach based on reduction, their block-based approach is more precise but is significantly slower for some programs. A detailed experimental comparison of the two approaches is presented in [24].

Atomicity is a semantic correctness condition for multithreaded software. In this respect, it is similar to strict serializability [53], a correctness condition for database transactions, and linearizability [54], a correctness condition for concurrent objects. Verifying that an object is linearizable requires full program verification. We hope that our analysis for atomicity can be leveraged to develop lightweight checking tools for related correctness conditions.

Artho et al. [20] have developed a dynamic analysis tool to identify one class of “higher-level races”. The analysis is based on the notion of *view consistency*. Intuitively, a view is the set of variables accessed within a synchronized block. Thread A is view consistent with B if all views from the execution of A, intersected with the maximal view of B, are ordered by subset inclusion. Violations of view consistency can indicate that a program may be using shared variables in a problematic way. View consistency violations can also be detected statically [55]. ESC/Java has been extended to catch a different notion of higher-level races, where a stale value from one synchronized block is used in a subsequent synchronized block [18].

While our tool checks atomicity, other researchers have proposed using atomicity as a language primitive, essentially implementing the serialized semantics \mapsto . Lomet [56] first proposed the use of atomic blocks for synchronization. The Argus [57] and Avalon [58] projects developed language support for implementing atomic objects. Persistent languages [59,60] are attempting to augment atomicity with data persistence in order to introduce transactions into programming languages. A more recent approach to supporting atomicity uses lightweight transactions implemented in the runtime system [61]. An alternative is to generate synchronization code automatically from high-level specifications [62].

7. Conclusions

Developing reliable multithreaded software is notoriously difficult, because concurrent threads often interact in unexpected and erroneous ways. Programmers try to avoid unintended interactions by designing methods and interfaces that are atomic, but traditional testing techniques are inadequate for verifying atomicity.

This paper presents a dynamic analysis designed to catch atomicity violations would be missed by traditional testing or (static or dynamic) race-detection techniques. This analysis has been implemented and applied to a range of benchmark programs, and has successfully detected atomicity violations in these programs. In addition, our experimental results suggest that over 90% of the methods in our benchmarks are atomic, which validates our hypothesis that atomicity is a fundamental design principle in multithreaded programs.

For future work, we hope to study *hybrid* atomicity checkers based on a synthesis of the dynamic and static approaches. In one combination, a static type-based analysis may verify many expected race-freedom and atomicity properties, and the dynamic atomicity checker could then focus on the unverified residue. For race detection, this hybrid approach has reduced the instrumentation overhead by an order of magnitude [15,16]; we expect comparable improvements when checking atomicity.

Acknowledgments

We thank Bill Thies, Martín Abadi, Shaz Qadeer, Rob O’Callahan, Mayur Naik and Scott Stoller for valuable comments on this work. We also thank Christof von Praun for his assistance in collecting test programs. This work was partly supported by the National Science Foundation under Grants CCF-0341179, CCF-0341387, and CCF-0644130, and by faculty research funds granted by the University of California at Santa Cruz and by Williams College. The first author was also supported by a Sloan Foundation Fellowship.

References

- [1] S. Savage, M. Burrows, G. Nelson, P. Sobalvarro, T.E. Anderson, Eraser: A dynamic data race detector for multi-threaded programs, *ACM Transactions on Computer Systems* 15 (4) (1997) 391–411.
- [2] A.D. Birrell, An introduction to programming with threads, Research Report 35, Digital Equipment Corporation Systems Research Center, 1989.
- [3] C. Flanagan, S.N. Freund, Type-based race detection for Java, in: *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2000, pp. 219–232.
- [4] C. Flanagan, M. Abadi, Types for safe locking, in: S.D. Swierstra (Ed.), *Proceedings of European Symposium on Programming*, in: *Lecture Notes in Computer Science*, vol. 1576, Springer-Verlag, 1999, pp. 91–108.
- [5] C. Flanagan, M. Abadi, Object types against races, in: J.C.M. Baeten, S. Mauw (Eds.), *Proceedings of the International Conference on Concurrency Theory*, in: *Lecture Notes in Computer Science*, vol. 1664, Springer-Verlag, 1999, pp. 288–303.
- [6] C. Flanagan, S.N. Freund, Detecting race conditions in large programs, in: *Workshop on Program Analysis for Software Tools and Engineering*, 2001, pp. 90–96.
- [7] C. Boyapati, M. Rinard, A parameterized type system for race-free Java programs, in: *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages and Applications*, 2001, pp. 56–69.
- [8] C. Boyapati, R. Lee, M. Rinard, A type system for preventing data races and deadlocks in Java programs, in: *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages and Applications*, 2002, pp. 211–230.
- [9] D. Grossman, Type-safe multithreading in Cyclone, in: *Proceedings of the ACM Workshop on Types in Language Design and Implementation*, 2003, pp. 13–25.
- [10] N. Sterling, Warlock: A static data race analysis tool, in: *Proceedings of the USENIX Winter Technical Conference*, 1993, pp. 97–106.
- [11] C. Flanagan, K.R.M. Leino, M. Lillibridge, G. Nelson, J.B. Saxe, R. Stata, Extended static checking for Java, in: *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2002, pp. 234–245.
- [12] A.T. Chamillard, L.A. Clarke, G.S. Avrunin, An empirical comparison of static concurrency analysis techniques, Technical Report 96-084, Department of Computer Science, University of Massachusetts at Amherst, 1996.
- [13] J.C. Corbett, Evaluating deadlock detection methods for concurrent software, *IEEE Transactions on Software Engineering* 22 (3) (1996) 161–180.
- [14] J.-D. Choi, K. Lee, A. Loginov, R. O’Callahan, V. Sarkar, M. Sridhara, Efficient and precise datarace detection for multithreaded object-oriented programs, in: *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2002, pp. 258–269.
- [15] C. von Praun, T. Gross, Static conflict analysis for multi-threaded object-oriented programs, in: *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2003, pp. 115–128.
- [16] R. O’Callahan, J.-D. Choi, Hybrid dynamic data race detection, in: *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 2003, pp. 167–178.
- [17] E. Pozniansky, A. Schuster, Efficient on-the-fly data race detection in multithreaded C++ programs, in: *Proceedings of the ACM Symposium on Principles and Practice of Parallel Programming*, 2003, pp. 179–190.
- [18] M. Burrows, K.R.M. Leino, Finding stale-value errors in concurrent programs, Technical Note 2002-004, Compaq Systems Research Center, 2002.
- [19] C. Flanagan, S. Qadeer, A type and effect system for atomicity, in: *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2003, pp. 338–349.
- [20] C. Artho, K. Havelund, A. Biere, High-level data races, in: *The First International Workshop on Verification and Validation of Enterprise Information Systems*, 2003.
- [21] C. Hoare, Monitors: An operating systems structuring concept, *Communications of the ACM* 17 (10) (1974) 549–557.
- [22] J. Manson, W. Pugh, S.V. Adve, The java memory model, in: *Proceedings of the ACM Symposium on the Principles of Programming Languages*, 2005, pp. 378–391.
- [23] L. Wang, S.D. Stoller, Run-time analysis for atomicity, in: *Proceedings of the Workshop on Runtime Verification*, in: *Electronic Notes in Computer Science*, vol. 89 (2), Elsevier, 2003.
- [24] L. Wang, S.D. Stoller, Runtime analysis of atomicity for multi-threaded programs, *IEEE Transactions on Software Engineering* 32 (2006) 93–110.
- [25] C. Flanagan, S.N. Freund, S. Qadeer, Exploiting purity for atomicity, *IEEE Transactions on Software Engineering* 31 (4) (2005) 275–291.
- [26] C. Flanagan, S.N. Freund, M. Lifshin, Type inference for atomicity, in: *Proceedings of the ACM Workshop on Types in Language Design and Implementation*, 2005, pp. 47–58.
- [27] S.N. Freund, S. Qadeer, Checking concise specifications for multithreaded software, *Journal of Object Technology* 3 (6) (2004) 81–101.
- [28] C. Flanagan, S.N. Freund, Type inference against races., in: *Proceedings of the Static Analysis Symposium*, 2004, pp. 116–132.

- [29] R.J. Lipton, Reduction: A method of proving properties of parallel programs, *Communications of the ACM* 18 (12) (1975) 717–721.
- [30] C. Flanagan, S. Qadeer, Types for atomicity, in: *Proceedings of the ACM Workshop on Types in Language Design and Implementation*, 2003, pp. 1–12.
- [31] J. Gosling, B. Joy, G. Steele, *The Java Language Specification*, Addison-Wesley, 1996.
- [32] C. von Praun, T. Gross, Object race detection, in: *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages and Applications*, 2001, pp. 70–82.
- [33] Standard Performance Evaluation Corporation, SPEC benchmarks. Available from: <http://www.spec.org/>, 2003.
- [34] World Wide Web Consortium, Jigsaw. Available from: <http://www.w3c.org>, 2001.
- [35] Java Grande Forum, Java Grande benchmark suite. Available from: <http://www.javagrande.org/>, 2003.
- [36] T. Kistler, J. Marais, WebL — A programming language for the web, in: *Proceedings of the International World Wide Web Conference*, in: *Computer Networks and ISDN Systems*, vol. 30, Elsevier, 1998, pp. 259–270.
- [37] T.W. Doepfner Jr., Parallel program correctness through refinement, in: *Proceedings of the ACM Symposium on the Principles of Programming Languages*, 1977, pp. 155–169.
- [38] R.-J. Back, A method for refining atomicity in parallel algorithms, in: *PARLE 89: Parallel Architectures and Languages Europe*, in: *Lecture Notes in Computer Science*, vol. 366, Springer-Verlag, 1989, pp. 199–216.
- [39] L. Lamport, F.B. Schneider, Pretending atomicity, *Research Report 44*, DEC Systems Research Center, 1989.
- [40] E. Cohen, L. Lamport, Reduction in TLA, in: *Proceedings of the International Conference on Concurrency Theory*, in: *Lecture Notes in Computer Science*, vol. 1466, Springer-Verlag, 1998, pp. 317–331.
- [41] J. Misra, *A Discipline of Multiprogramming: Programming Theory for Distributed Applications*, Springer-Verlag, 2001.
- [42] L. Fajstrup, E. Goubault, M. Raussen, Detecting deadlocks in concurrent systems, in: D. Sangiorgi, R. de Simone (Eds.), *Proceedings of the International Conference on Concurrency Theory*, in: *Lecture Notes in Computer Science*, vol. 1466, Springer-Verlag, 1998, pp. 332–347.
- [43] M.B. Dwyer, L.A. Clarke, Data flow analysis for verifying properties of concurrent programs, *Technical Report 94-045*, Department of Computer Science, University of Massachusetts at Amherst, 1994.
- [44] Polyspace technologies. Available at: <http://www.polyspace.com>, 2003.
- [45] N. Kobayashi, A partially deadlock-free typed process calculus, *ACM Transactions on Programming Languages and Systems* 20 (2) (1998) 436–482.
- [46] N. Kobayashi, S. Saito, E. Sumii, An implicitly-typed deadlock-free process calculus, in: C. Palamidessi (Ed.), *Proceedings of the International Conference on Concurrency Theory*, in: *Lecture Notes in Computer Science*, vol. 1877, Springer-Verlag, 2000, pp. 489–503.
- [47] R. Agarwal, S.D. Stoller, Type inference for parameterized race-free Java, in: *Proceedings of the Conference on Verification, Model Checking, and Abstract Interpretation*, in: *Lecture Notes in Computer Science*, vol. 2937, Springer, 2004, pp. 149–160.
- [48] D.F. Bacon, R.E. Strom, A. Tarafdar, Guava: A dialect of Java without data races, in: *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages and Applications*, 2001, pp. 382–400.
- [49] A. Sasturkar, R. Agarwal, L. Wang, S.D. Stoller, Automated type-based analysis of data races and atomicity., in: *Proceedings of the ACM Symposium on Principles and Practice of Parallel Programming*, 2005, pp. 83–94.
- [50] S. Qadeer, S.K. Rajamani, J. Rehof, Summarizing procedures in concurrent programs, in: *Proceedings of the ACM Symposium on the Principles of Programming Languages*, 2004, pp. 245–255.
- [51] J. Hatcliff Robby, M.B. Dwyer, Verifying atomicity specifications for concurrent object-oriented software using model-checking, in: *Proceedings of the International Conference on Verification, Model Checking and Abstract Interpretation*, 2004, pp. 175–190.
- [52] C. Flanagan, Verifying commit-atomicity using model-checking, in: *SPIN 2004: 11th International SPIN Workshop on Model Checking of Software*, 2004, pp. 252–266.
- [53] C. Papadimitriou, *The theory of database concurrency control*, Computer Science Press, 1986.
- [54] M.P. Herlihy, J.M. Wing, Linearizability: A correctness condition for concurrent objects, *ACM Transactions on Programming Languages and Systems* 12 (3) (1990) 463–492.
- [55] C. von Praun, T. Gross, Static detection of atomicity violations in object-oriented programs, in: *Workshop on Formal Techniques for Java-like Programs*, 2003.
- [56] D.B. Lomet, Process structuring, synchronization, and recovery using atomic actions, *Language Design for Reliable Software (1977)* 128–137.
- [57] B. Liskov, D. Curtis, P. Johnson, R. Scheifler, Implementation of Argus, in: *Proceedings of the Symposium on Operating Systems Principles*, 1987, pp. 111–122.
- [58] J.L. Eppinger, L.B. Mummert, A.Z. Spector, Camelot and Avalon: A Distributed Transaction Facility, Morgan Kaufmann, 1991.
- [59] M.P. Atkinson, K.J. Chisholm, W.P. Cockshott, PS-Algol: An Algol with a persistent heap, *ACM SIGPLAN Notices* 17 (7) (1981) 24–31.
- [60] M.P. Atkinson, D. Morrison, Procedures as persistent data objects, *ACM Transactions on Programming Languages and Systems* 7 (4) (1985) 539–559.
- [61] T.L. Harris, K. Fraser, Language support for lightweight transactions, in: *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages and Applications*, 2003, pp. 388–402.
- [62] X. Deng, M. Dwyer, J. Hatcliff, M. Mizuno, Invariant-based specification, synthesis, and verification of synchronization in concurrent programs, in: *International Conference on Software Engineering*, 2002, pp. 442–452.