DoubleChecker: Efficient Sound and Precise Atomicity Checking *

Swarnendu Biswas Jipeng Huang Aritra Sengupta Michael D. Bond

Department of Computer Science and Engineering
Ohio State University
{biswass,huangjip,sengupta,mikebond}@cse.ohio-state.edu

Abstract

Atomicity is a key correctness property that allows programmers to reason about code regions in isolation. However, programs often fail to enforce atomicity correctly, leading to atomicity violations that are difficult to detect. Dynamic program analysis can detect atomicity violations based on an atomicity specification, but existing approaches slow programs substantially.

This paper presents DoubleChecker, a novel sound and precise atomicity checker whose key insight lies in its use of two new cooperating dynamic analyses. Its *imprecise* analysis tracks cross-thread dependences soundly but imprecisely with significantly better performance than a fully precise analysis. Its *precise* analysis is more expensive but only needs to process a subset of the execution identified as potentially involved in atomicity violations by the imprecise analysis. If DoubleChecker operates in *single-run* mode, the two analyses execute in the same program run, which guarantees soundness and precision but requires logging program accesses to pass from the imprecise to the precise analysis. In *multi-run* mode, the first program run executes only the imprecise analysis, and a second run executes both analyses. Multi-run mode trades accuracy for performance; each run of multi-run mode outperforms single-run mode, but can potentially miss violations.

We have implemented DoubleChecker and an existing state-of-the-art atomicity checker called Velodrome in a high-performance Java virtual machine. DoubleChecker's single-run mode significantly outperforms Velodrome, while still providing full soundness and precision. DoubleChecker's multi-run mode improves performance further, without significantly impacting soundness in practice. These results suggest that DoubleChecker's approach is a promising direction for improving the performance of dynamic atomicity checking over prior work.

Categories and Subject Descriptors D.2.4 [*Software Engineering*]: Software/Program Verification—reliability; D.2.5 [*Software Engineering*]: Testing and Debugging—monitors, testing tools; D.3.4 [*Programming Languages*]: Processors—compilers, runtime environments

Keywords atomicity checking; dynamic program analysis

* This material is based upon work supported by the National Science Foundation under Grants CAREER-1253703 and CSR-1218695.

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PLDI '14, June 9–11, 2014, Edinburgh, United Kingdom.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2784-8/14/06...$15.00.
http://dx.doi.org/10.1145/2594291.2594323

1. Introduction

Modern multicore hardware trends have made parallelism necessary for performance, but developing parallel programs that are correct and scalable is notoriously challenging. Concurrency bugs have caused several high-profile failures (e.g., [33]), a testament to the fact that concurrency errors are present even in well-tested code (e.g., [15]). According to a study of real-world concurrency bugs, 65% of concurrency errors are due to atomicity violations [23].

Atomicity is a fundamental non-interference property that eases reasoning about program behavior in multithreaded programs. For programming language semantics, atomicity is synonymous with *serializability*: program execution must be equivalent to some serial execution of atomic regions. That is, the code block's execution *appears* not to be interleaved with statements from other concurrently executing threads. Programmers can thus reason about atomic regions without considering effects of other threads. However, modern general-purpose languages provide crude support for enforcing atomicity—programmers are basically stuck using locks to control how threads' shared-memory accesses can interleave. Programmers try to maximize scalability by minimizing synchronization, often mistakenly writing code that does not correctly enforce needed atomicity properties.

An *atomicity specification* denotes particular code regions that are expected to execute atomically. Program analysis can check atomicity by checking whether a program conforms to the atomicity specification. A *violation* indicates that the program or specification is wrong (or both). Writing an atomicity specification may seem burdensome, but prior work shows that specifications can be derived mostly automatically [11, 14].

Existing analyses. Static analysis can check atomicity but is inherently imprecise, and type-based approaches rely on annotations [8, 10, 13, 15]. Existing dynamic analyses are precise but slow programs by up to an order of magnitude or more [9, 11, 14, 24, 35–37]. Dynamic approaches incur high costs to track cross-thread dependences, which is especially expensive because it requires inserting intrusive synchronization to ensure correctness. We compare most closely with the state-of-the-art *Velodrome* algorithm, which soundly and precisely checks *conflict serializability*, a sufficient condition for serializability [14]. Velodrome slows programs by about an order of magnitude on average [14], mainly because of the high cost of identifying cross-thread data dependences soundly and precisely (Section 2).

Motivation. Atomicity violations are common but serious errors that are sensitive to inputs, environments, and thread interleavings, so violations manifest unexpectedly and only in certain settings. *Low-overhead* checking is needed in order to use it liberally to find bugs during in-house, alpha, and beta testing, and perhaps even some production settings. Greathouse et al. note that high dynamic analysis overheads "reduce the degree to which programs can be tested within a reasonable amount of time. Beyond that,

high overheads slow debugging efforts, as repeated runs of the program to hunt for root causes and verify fixes also suffer these slowdowns" [17]. Our goal is to reduce the cost of sound and precise atomicity checking significantly in order to increase its practicality for various use cases.

Our Approach

This paper presents a sound and precise dynamic conflict serializability checker called DoubleChecker that significantly reduces overhead compared with existing state-of-the-art detectors. The key insight of DoubleChecker lies in its *dual-analysis* approach that avoids the high costs of precisely tracking cross-thread dependences and performing synchronized metadata updates, by over-approximating dependences between transactions (a *transaction* is a dynamically executing atomic region) and then recovering precision only for those transactions that might be involved in violations.

DoubleChecker achieves low overhead by staging work between two new analyses, one imprecise and the other precise. The *imprecise* analysis constructs a graph that soundly but imprecisely captures dependences among transactions. The imprecise analysis (1) detects cross-thread dependences by extending an existing concurrency control mechanism [3]; (2) computes dependence edges that soundly imply the true cross-thread dependences; (3) detects cycles in the graph, which indicate potential atomicity violations and are a superset of the true (precise) cycles; and (4) (when the precise analysis executes in the same run as the imprecise analysis) captures enough information about program accesses to allow reconstruction of precise dependences. The *precise* analysis computes precise cross-thread dependences and checks for cycles in a precise dependence graph. However, the precise analysis processes *only those transactions that the imprecise analysis identified* as being involved in a cycle. In practice, these transactions are a reasonable overapproximation of the precise cycles, eliminating most of the expensive work that would otherwise normally be performed by *any* sound and precise analysis.

DoubleChecker supports two execution modes. In *single-run* mode, the imprecise and precise analyses operate on the *same* program execution. Single-run mode requires the imprecise analysis to record all program accesses so that the precise analysis can determine the precise dependences among the transactions identified by the imprecise analysis. Single-run mode is thus fully sound and precise: it detects all atomicity violations in an execution.

In *multi-run* mode, the first and second runs operate on *different* program runs. The *first run* executes only the imprecise analysis, while the *second run* executes both the imprecise and precise analyses (similar to single-run mode). The first run is thus imprecise whereas the second run is precise. The first run passes *static* program information about imprecise cycles to the second run to help reduce the instrumentation introduced by the second run. The multi-run mode is unsound since the first and second runs operate on different program executions, which could differ due to different program inputs and thread interleavings. The multi-run mode can thus miss atomicity violations that occur in either program run.

We have implemented DoubleChecker and prior work's Velodrome in a high-performance Java virtual machine. We evaluate correctness, performance, and other characteristics of DoubleChecker on large, real-world multithreaded programs, and compare with Velodrome. In single-run mode, DoubleChecker is a fully sound and precise analysis that slows programs by 3.6X on average, a significant improvement over Velodrome's 6.1X slowdown. A limitation of single-run mode is that the imprecise analysis must log precise information about reads and writes, which not only slows execution, but also adds high memory overhead, sometimes exhausting the virtual memory of our 32-bit platform. DoubleChecker's multi-run mode does not guarantee soundness for either run, although we show it can find a high fraction of atomicity violations in practice. Its first and second runs slow programs by 1.9X and 2.4X, respectively. As such, the *overhead* added by DoubleChecker in its single- and multi- run modes is 1.9 and 3.7–5.6 times *less* than Velodrome's, respectively. These results suggest that DoubleChecker's novel approach is a promising direction for providing significantly better performance for dynamic atomicity checking.

Contributions. This paper makes the following contributions:

- a novel sound and precise dynamic atomicity checker based on using two new, cooperating analyses:
  1. an imprecise analysis that shows it can be cheaper to over-approximate dependence edges rather than compute them precisely, and thus detect cycles whose transactions are a superset of the true (precise) cycles, and
  2. a precise analysis that processes an execution history of only those transactions that are involved in potential cycles;
- two modes of execution that provide two choices for balancing soundness and performance;
- publicly available implementations of DoubleChecker and Velodrome; and
- an evaluation that shows DoubleChecker outperforms Velodrome significantly, with multi-run mode providing better performance without sacrificing much soundness in practice.

2. Background: Checking Conflict Serializability

Velodrome is a dynamic analysis that checks conflict serializability soundly and precisely [14]. (A related dynamic approach also checks conflict serializability soundly and precisely [9]; Section 6.) Each code region that is supposed to execute atomically (according to the atomicity specification) executes as a *transaction*. Other accesses each execute as a *unary transaction*. Velodrome's dynamic analysis builds a graph of transactions at run time. When a new (regular or unary) transaction starts, the analysis adds an intra-thread dependence edge from the thread's prior transaction to the new transaction. At each access, the analysis detects cross-thread data dependences: write–read, read–write, and write–write dependences between threads, as well as release–acquire synchronization dependences. (Although treating synchronization edges as cross-thread dependences can lead to false positives when checking conflict serializability, as Section 6 discusses, DoubleChecker follows Velodrome and includes synchronization edges.) Velodrome adds cross-thread dependence edges between transactions as the program executes. It detects cycles in the graph; a cycle is a sound and precise condition for a conflict serializability violation.

Velodrome slows programs by 12.7X in prior work [14] and 6.1X using our Velodrome implementation and experiments. These slowdowns are due largely to tracking cross-thread dependences soundly and precisely, which has two main costs. First, tracking dependences involves maintaining the last transaction to write, and each thread's last transaction to read, each variable. Second, to preserve correctness in the face of accesses potentially involved in data races, the analysis must use atomic operations and memory fences to ensure that an access and its corresponding analysis execute together atomically. Atomic operations and memory fences slow execution by limiting reordering and serializing in-flight instructions and by triggering remote cache misses.

3. Design of DoubleChecker

This section describes our dynamic conflict serializability checker that uses two cooperating dynamic analyses to check atomicity without incurring the full costs of tracking cross-thread depen-

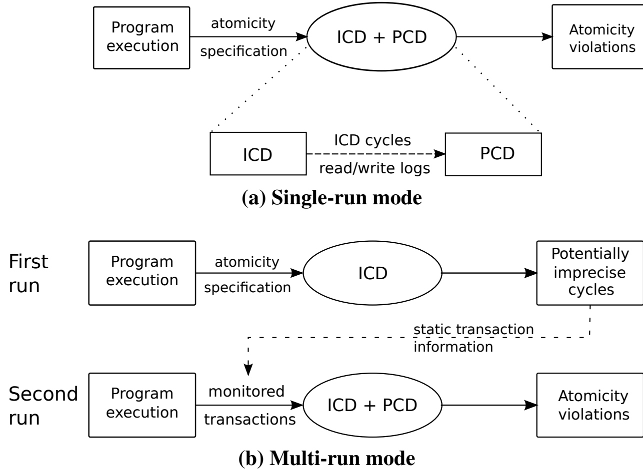

Figure 1. An overview of DoubleChecker's two execution modes.

dences soundly and precisely for all program accesses. Following an overview of DoubleChecker's analyses and execution modes, Section 3.2 describes the imprecise analysis, and Section 3.3 describes the precise analysis.

3.1 Overview

DoubleChecker's imprecise analysis, called *imprecise cycle detection* (ICD), monitors all program accesses to track cross-thread dependences soundly but imprecisely, i.e., the dependences imply the execution's actual dependences as well as other false dependences. ICD is inherently imprecise mainly because it identifies dependence edges by tracking shared-memory "ownership"; a transfer of ownership indicates a possible dependence, but does not guarantee a dependence nor identify the source of the dependence edge. ICD constructs a *dependence graph* whose nodes are transactions and whose edges correspond to the cross-thread dependences that ICD detects. ICD checks for cycles in this graph.

The second analysis, *precise cycle detection* (PCD), is a sound and precise analysis that limits its monitoring to a *subset* of all transactions: the transactions identified by ICD as being involved in potential cycles—which preserves soundness because every precise cycle's transactions will always be a subset of some (potentially imprecise) cycle identified by ICD. Note that PCD is *not* a standalone analysis: it performs its analysis on an execution's access log, provided by ICD.

DoubleChecker can operate in either of two modes. Figure 1 overviews the two modes of DoubleChecker.

Single-run mode. In *single-run* mode, ICD and PCD run on the same program execution. ICD logs all program reads and writes and ordering dependences between them, so PCD can identify precise cycles. A key cost of single-run mode is logging all program accesses.

Multi-run mode. In testing and deployment situations, programs are run multiple times with various inputs. DoubleChecker's *multi-run* mode takes advantage of this situation by splitting work across multiple program runs.[1] One run can identify transactions that might be involved in a dependence cycle, and another run can focus its monitoring on this set of transactions. In contrast to single-run mode, multi-run mode avoids logging all accesses during the first run by instead performing precise checking during a second run of the program. The *first* run of multi-run mode uses only ICD. This run identifies all *regular* (non-unary) transactions that are involved in imprecise cycles according to their static starting locations (e.g., method signatures). Rather than identifying precisely which *unary* transactions were involved in cycles, the first run identifies only whether *any* unary transactions were involved in any cycle. It would be expensive to identify unary transactions precisely, since it would essentially require recording the program location of every non-transactional access.

The *second* run takes this static transaction information—set of regular transactions plus a boolean about unary transactions—as input, and limits its analysis to the identified regular transactions and instruments *all* unary transactions if the unary transaction boolean is true. We find this approximation yields acceptable performance in practice since most accesses are *not* unary, i.e., they occur inside regular transactions. In our experiments, the second run uses both ICD and PCD—similar to the single-run mode—for the best performance, but the second run can also potentially use a different precise checker such as Velodrome.

In multi-run mode, DoubleChecker guarantees soundness if the two program runs execute identically. In practice, two executions in the wild will take different inputs and execute different thread interleavings. The set of (static) transactions identified by the first run may not be involved in a cycle in the second run; similarly, the second run may execute transactional cycles not present in the first run. To increase efficacy, the second run can take as input all transactions identified across multiple executions of the first run. The multi-run mode can still be effective in practice if the same regions tend to be involved in cycles across multiple runs.

3.2 Imprecise Cycle Detection

Imprecise cycle detection (ICD) is a dynamic analysis that analyzes all program execution in order to detect cycles among transactions. ICD constructs a sound but imprecise graph called the *imprecise dependence graph* (IDG) to model dependences among the transactions in a multithreaded program. The nodes in an IDG are regular transactions (which correspond to atomic regions) or unary transactions (which correspond to single accesses outside of atomic regions). A *cross-thread* edge between two nodes in different threads indicates a (potentially imprecise) cross-thread dependence between the transactions. Two consecutive nodes (i.e., transactions) in the same thread are connected by an intra-thread edge that effectively captures any intra-thread dependences.

We first describe an existing concurrency control mechanism that ICD extends to help detect cross-thread dependences but that makes detection inherently imprecise. We then describe how ICD builds the IDG and detects cycles.

3.2.1 Background: Concurrency Control

This section describes *Octet*, a recently developed software-based concurrency control mechanism [3] that ICD uses to help detect cross-thread dependences. Octet establishes and identifies *happens-before* relationships [21] that soundly but imprecisely imply all of an execution's cross-thread dependences.

At run time, Octet maintains a locality state for each object[2] that can be one of the following: $\mathsf{WrEx_T}$ (write-exclusive for thread T), $\mathsf{RdEx_T}$ (read-exclusive for thread T), or $\mathsf{RdSh_c}$ (read-shared; we explain the counter c later). Table 1 shows the possible state transitions based on an access and the object's current state. To maintain each object's state at run time, the compiler inserts instrumentation called a write barrier[3] before every store:

```
if (obj.state != WrEx_T) { // fast path
  /* slow path: change obj.state */
}
obj.f = ... ;  // program write
```

[1] Prior bug detection work has split work across runs using sampling (e.g., [22]), which is complementary to our work.

[2] We use the term "object" to refer to any unit of shared memory.

[3] A barrier is instrumentation added to every program load and store [38].

Trans. type	Old state	Access	New state	Cross-thread dependence?
Same state	$\mathsf{WrEx_T}$	R or W by T	Same	No
	$\mathsf{RdEx_T}$	R by T	Same	
	$\mathsf{RdSh_c}$	R by T *	Same	
Upgrading	$\mathsf{RdEx_T}$	W by T	$\mathsf{WrEx_T}$	No
	$\mathsf{RdEx_{T1}}$	R by T2	$\mathsf{RdSh_{gRdShCnt}}$	Possibly
Fence	$\mathsf{RdSh_c}$	R by T *	Same *	Possibly
Conflicting	$\mathsf{WrEx_{T1}}$	W by T2	$\mathsf{WrEx_{T2}}$	Possibly
	$\mathsf{WrEx_{T1}}$	R by T2	$\mathsf{RdEx_{T2}}$	
	$\mathsf{RdEx_{T1}}$	W by T2	$\mathsf{WrEx_{T2}}$	
	$\mathsf{RdSh_c}$	W by T	$\mathsf{WrEx_T}$	

Table 1. Octet state transitions. *A read to a $\mathsf{RdSh_c}$ object by T triggers a fence transition if and only if per-thread counter T.rdShCnt < c. The fence transition updates T.rdShCnt to c.

and a read barrier before every load:

```
if (obj.state != WrEx_T && obj.state != RdEx_T && // fast
    !(obj.state == RdSh_c && T.rdShCnt >= c)) {   // path
  /* slow path: change obj.state */
}
... = obj.f;   // program read
```

The state check, called the *fast path*, checks whether the state needs to change (the *Same state* rows in Table 1). The key to Octet's performance is that the fast path is simple and performs no writes or synchronization. If the state needs to change, the *slow path* executes in order to change the state.

Conflicting transitions. The last four rows of Table 1 show *conflicting* state transitions, which indicate a conflicting access and require a coordination protocol to perform the state change. For example, in Figure 2, before thread T2 can change an object o's state from $\mathsf{WrEx_{T1}}$ to $\mathsf{RdEx_{T2}}$, T2 must *coordinate* with T1 to ensure that T1 does not continue accessing o racily without synchronization. As part of this coordination protocol, T1 does not respond to T2's request until it reaches a *safe point*: a program point definitely *not* between an Octet barrier and its corresponding program access.

The coordination protocol for conflicting transitions first puts o into an *intermediate* state, which helps simplify the protocol by ensuring that only one thread at a time tries to change an object's state. For example, if T2 wants to read an object that is in the $\mathsf{WrEx_{T1}}$ state, T2 first puts the object into the $\mathsf{RdEx^{Int}_{T2}}$ state. The coordination protocol is then performed in one of two ways:

- The threads perform the *explicit* protocol if T1 is executing code normally. T2 sends a request to T1, and T1 responds to the request when it reaches a safe point. When T2 observes the response, a roundtrip happens-before relationship has been established, so T2 can change the state to $\mathsf{RdEx_{T2}}$ and proceed.
- Otherwise, thread T1 is "blocking," e.g., waiting for synchronization or I/O. Rather than waiting for T1, T2 *implicitly* coordinates with T1 by atomically setting a flag that T1 will observe when it unblocks. This protocol establishes a happens-before relationship, so T2 can change the state to $\mathsf{RdEx_{T2}}$ and proceed.

Upgrading and fence transitions. Upgrading and fence transitions (middle rows of Table 1) do not require coordination because other threads can safely continue accessing the object under the old state. In Figure 2, T3 atomically *upgrades* an object's state from $\mathsf{RdEx_{T2}}$ to $\mathsf{RdSh_c}$. The value c is the new value of a global counter gRdShCnt that gets incremented atomically every time an object transitions to RdSh, establishing a global order of all transitions to RdSh. This state change establishes a happens-before relationship from the read on T2 to the current program point on T3, ensuring a transitive happens-before relationship from T1's write to T3's read.

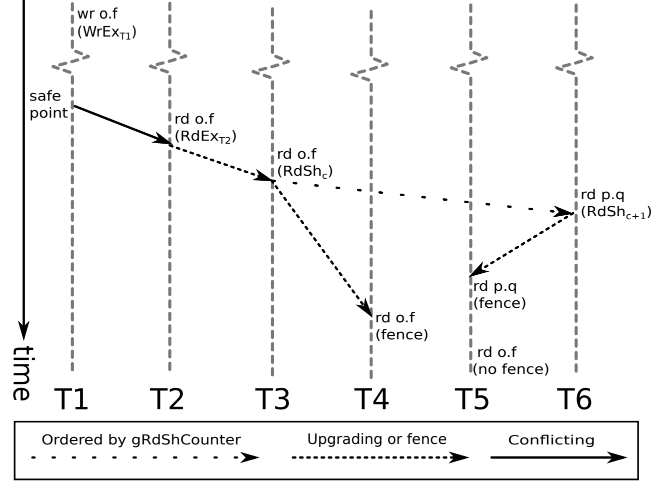

Figure 2. A possible interleaving of six concurrent threads accessing shared objects o and p, and the corresponding Octet state transitions they trigger (with new states shown in parentheses).

In Figure 2, T4 reads o in the $\mathsf{RdSh_c}$ state. To ensure a happens-before relationship from the last write to o (by T1) to this read in T4, a *fence* transition is triggered. The fence transition is triggered when a thread's local counter T.rdShCnt is not up-to-date with the counter c in $\mathsf{RdSh_c}$. T4 updates T4.rdShCnt to c and issues a memory fence to ensure a happens-before relationship from T3's transition to $\mathsf{RdSh_c}$ to T4's read.

In Figure 2, T5 reads o but does *not* trigger a fence transition because T5 has already read an object (p) in the $\mathsf{RdSh_{c+1}}$ state. However, a *transitive* happens-before relationship exists from T1's write to T5's read of o because there is a happens-before relationship from o's state transition to $\mathsf{RdSh_c}$ in T3 to p's transition to $\mathsf{RdSh_{c+1}}$ in T6 (since both transitions update gRdShCnt atomically).

Octet's state transitions thus establish happens-before edges that transitively imply all cross-thread dependences [3]. ICD can piggyback on Octet's state transitions to identify potential cross-thread dependences. Next, we address the challenge of actually identifying the dependence edges that ICD should add to the IDG.

3.2.2 Identifying Cross-Thread Dependences

ICD uses Octet to help detect cross-thread dependences. While Octet establishes happens-before relationships that soundly imply all cross-thread dependences, it does not precisely identify the exact points in the execution with which happens-before relationships are established. ICD addresses the challenge of how to identify these program points and thus add cross-thread edges to the IDG that soundly imply all cross-thread dependences, so that any true dependence cycle will lead to a cycle in the IDG. In this way, ICD detects atomicity violations soundly but imprecisely with substantially lower overhead than a fully precise approach.

The challenge of identifying each cross-thread edge is in identifying its *source*; the *sink* is obvious since it is the current execution point on the thread triggering the state change. ICD keeps track of a few "last transaction to do X" facts, to help identify sources of cross-thread edges later:

T.lastRdEx – Per-thread variable that stores the last transaction of thread T to transition an object into the $\mathsf{RdEx_T}$ state.

gLastRdSh – Global variable that stores the last transaction among all threads to transition an object into a RdSh state.

We also define the following helper function:

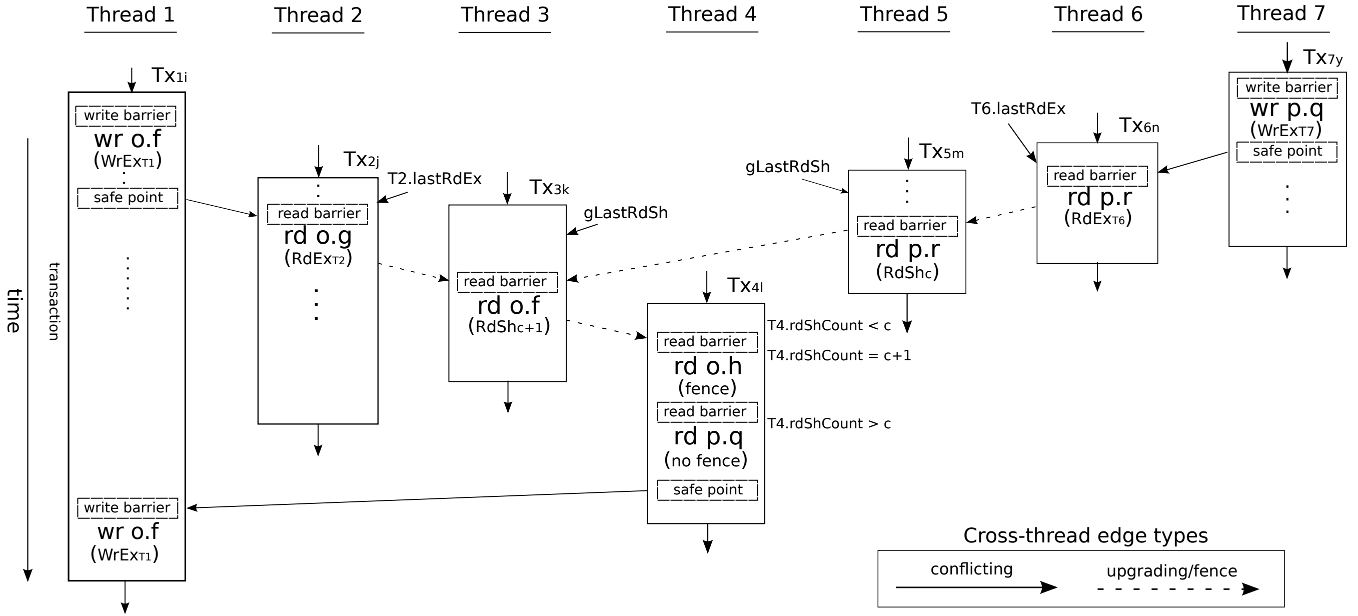

Figure 3. An example interleaving of threads executing atomic regions of code as transactions. The figure shows the Octet states after each access and the IDG edges added by ICD.

```
procedure handleConflictingTransition(respT, reqT, oldState, newState)
    IDG.addEdge(currTX(respT) → currTX(reqT))
    if newState = RdEx_reqT then
        reqT.lastRdEx := currTX(reqT)
    end if
end procedure
procedure handleUpgradingTransition(T, oldState, newState)
    Let rdExThread be the thread T such that oldState = RdEx_T
    IDG.addEdge(rdExThread.lastRdEx → currTX(T))
    IDG.addEdge(gLastRdSh → currTX(T))
    gLastRdSh := currTX(T)
end procedure
procedure handleFenceTransition(T)
    IDG.addEdge(gLastRdSh → currTX(T))
end procedure
```

Figure 4. ICD procedures called from Octet state transitions.

currTX(T) – Returns the transaction currently executing on T.

Creating cross-thread edges for conflicting transitions. A conflicting transition involves one requesting thread reqT, which coordinates with each responding thread respT. ICD piggybacks on each invocation of the coordination protocol, using the procedure handleConflictingTransition() in Figure 4, in order to add an edge to the IDG.

Either reqT or respT will invoke the procedure as part of the coordination protocol, depending on whether the explicit or implicit protocol is used. For the explicit protocol, respT invokes the procedure before it responds, which is safe since both threads are stopped at that point. For the implicit protocol, reqT invokes the procedure since respT is blocked; reqT first atomically places a "hold" on respT so respT will not unblock while reqT invokes the procedure.

Figure 3 shows a possible thread interleaving among seven concurrent threads executing transactions. The edges among transactions are IDG edges that ICD adds. The access rd o.g in Tx_{2j} conflicts with the first write to object o in transaction Tx_{1i}. The handleConflictingTransition() procedure creates a cross-thread edge in the IDG from Tx_{1i} (the transaction executing the responding safe point) to Tx_{2j} (the transaction triggering the conflicting transition).

To help upgrading transitions (explained next), handleConflictingTransition() updates the per-thread variable T.lastRdEx, the last transaction to put an object into $RdEx_T$. In Figure 3, this procedure updates T2.lastRdEx to Tx_{2j} (not shown).

Creating cross-thread edges for upgrading transitions. To see why ICD needs to add cross-thread edges for upgrading transitions (and not just for conflicting transitions), consider the upgrading transition from $RdEx_{T2}$ to $RdSh_{c+1}$ in Figure 3. Creating a cross-thread edge is necessary to capture the dependence from T1's write to o to T3's read of o *transitively*. To create this edge, T3 invokes the procedure handleUpgradingTransition() in Figure 4.

This procedure also creates a second edge: from the last transaction to transfer an object to the RdSh state, referenced by gLastRdSh, to the current transaction. This edge orders all transitions to RdSh state, and is needed in order to capture dependences related to fence transitions, discussed next. For rd o.f in Tx_{3k}, the procedure creates an edge from gLastRdSh, which is Tx_{5m}, to the current transaction. Finally, the procedure updates gLastRdSh to point to the current transaction, Tx_{3k}.

ICD safely ignores $RdEx_T \rightarrow WrEx_T$ upgrading transitions. Any new dependences created by this transition are already captured transitively by existing intra-thread and cross-thread edges.

Creating cross-thread edges for fence transitions. ICD adds edges to the IDG for fence transitions, in order to capture a possible write–read dependence for RdSh objects. Each fence transition calls handleFenceTransition() (Figure 4), which creates an edge from the last transaction to transition an object to RdSh (gLastRdSh) to the current transaction.

In Figure 3, T4's read of o.h triggers a fence transition and a call to handleFenceTransition(), which creates an edge from gLastRdSh (Tx_{3k}) to Tx_{4l}. This edge helps capture the possible dependence from T1's write to T4's read (in this case, no true dependence exists since the accesses are to different fields).

After T4 reads o.h, it reads p.q, which does not trigger a fence transition because T4 has already read an object (o) with a *more recent* RdSh counter (c+1) than p's RdSh counter (c). However, because RdEx → RdSh transitions create edges between all transactions that transition an object to RdSh (e.g., the edge from Tx_{5m} to Tx_{3k}), all write–read dependences are captured by IDG edges even if they do not trigger a fence transition. In the figure, the IDG edges added by the procedures transitively capture the dependence from T7's write to p.q to T4's read of p.q.

Handling synchronization operations. Like Velodrome [14], DoubleChecker captures dependences not only between reads and writes to program variables, but also between synchronization operations: lock release–acquire, notify–wait, and thread fork and join. ICD handles these operations by treating acquire-like operations as reads and release-like operations as writes, on the object being synchronized on.

Sources of imprecision. ICD is imprecise because the edges it adds to the IDG are imprecise in several ways. First, ICD does not maintain the last transaction to read and write each object, so it identifies last accesses conservatively. For a conflicting transition, ICD adds an edge from the responding thread's last safe point. For an upgrading transition from $RdEx_T$ to RdSh, it adds an edge from T's last transition to $RdEx_T$, which may involve a different object.

ICD not only does not maintain the last reader transactions, but it does not maintain even the last reader *threads* for a RdSh object. ICD adds edges between all upgrading transitions to RdSh (to help enable sound tracking of write–read dependences for RdSh objects). For conflicting transitions from RdSh to $WrEx_T$, ICD adds edges from all threads to T's current transaction.

Finally, ICD tracks dependences at object granularity instead of field granularity.

ICD's imprecision is inherent in its use of Octet, which gives up precise detection of dependences for better performance. Eliminating some but not all sources of ICD's imprecision would be of little use, since ICD would still be imprecise.

3.2.3 Cycle detection

Rather than triggering cycle detection each time it creates a cross-thread edge (as Velodrome does [14]), ICD waits until a transaction ends to detect cycles. Consider the following example.

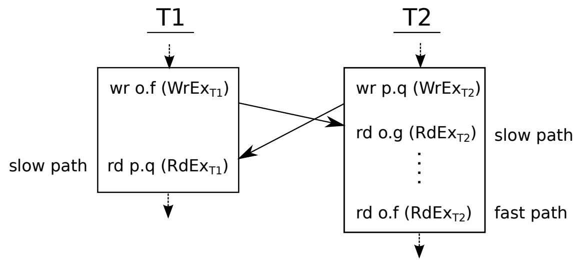

Even if T1 and T2 each trigger cycle detection when they add cross-thread edges, no precise cycle exists until rd o.f executes. In single-run mode, to ensure that PCD sees the precise cycle, ICD should report the cycle only after the transaction finishes. By invoking cycle detection when transactions end, ICD is guaranteed to detect each cycle at least once. In the first run of multi-run mode, deferring cycle detection until transactions finish is not strictly necessary but leads to fewer invocations of cycle detection.

Detecting strongly connected components. A side effect of delayed cycle detection is that a transaction might be involved in multiple cycles. ICD therefore computes the maximal strongly connected component (SCC) [7] starting from the transaction that just ended, which identifies the set of all transactions that are part of a cycle. The SCC computation explores a transaction tx only if tx has finished. This rule is sound because if tx is indeed involved in cycles, an SCC computation launched when tx finishes will detect those cycles. Avoiding processing unfinished transactions helps prevent identifying the same cycles multiple times, and it avoids races with threads updating their current transaction's state.

In Figure 3, ICD detects an SCC (in this case, a simple cycle) of size four when transaction Tx_{1i} ends. In single-run mode or the second run of multi-run mode, ICD passes these transactions to PCD for further processing. Note that PCD detects a precise cycle involving Tx_{1i} and Tx_{3k}. In contrast, if Tx_{3k} did not execute rd o.f, ICD would still detect an imprecise cycle, but PCD would not detect a precise cycle since none exists.

3.2.4 Maintaining Read/Write Logs

In single-run mode or the second run of multi-run mode, when ICD detects a cycle, it passes the set of transactions involved in the cycle to PCD. PCD also needs to know the exact accesses that have executed as well as the cross-thread ordering between them. To provide this information, ICD records read/write logs for every transaction: the exact memory accesses (e.g., object fields) read and written by the transaction. To accomplish this, ICD adds instrumentation before each program access but after Octet's instrumentation that records the access in the current transaction's read/write log. Synchronization operations are recorded as reads or writes to the objects being synchronized on. ICD provides cross-thread ordering of accesses by recording, for each IDG edge, not only the source and sink transactions of the edge, but also the exact read/write log entries that correspond to the edge's source and sink.

3.2.5 Soundness Argument

We now argue that ICD is a sound first-pass filter. In particular, we show that for any actual (precise) cycle of dependences, there exists an (imprecise) IDG cycle that is a superset of the precise cycle.

Let C be any set of executed nodes $tx_1, tx_2, \ldots, tx_n$ whose (sound and precise) dependence edges form a (sound and precise) cycle $tx_1 \rightarrow tx_2 \rightarrow \ldots \rightarrow tx_n \rightarrow tx_1$.

Let $tx_i \rightarrow tx_j$ be any dependence edge in C. Since ICD adds edges to the IDG that imply all dependences soundly, there must exist a path of edges from tx_i to tx_j in the IDG.

Thus there exists a path $tx_1 \rightarrow tx_2 \rightarrow \ldots \rightarrow tx_n \rightarrow tx_1$ in the IDG. ICD will detect this as a cycle $C' \supseteq C$ and pass C' to PCD. Since C' contains all nodes in C, and PCD computes all dependences between nodes in C', PCD will compute the dependences $tx_1 \rightarrow tx_2 \rightarrow \ldots \rightarrow tx_n \rightarrow tx_1$, and it will thus detect the cycle C.

3.3 Precise Cycle Detection

Precise cycle detection (PCD) is a sound and precise analysis that identifies cycles of dependences on a set of transactions provided as input. DoubleChecker invokes PCD with the following input from ICD: (1) a set of transactions identified by ICD as being involved in an SCC, (2) the read/write logs of the transactions, and (3) the cross-thread edges added by ICD recorded relative to read/write log entries (to order conflicting accesses). PCD processes each SCC identified by ICD separately. Using these inputs, PCD essentially "replays" the subset of execution corresponding to the transactions in the IDG cycle, and performs a sound and precise analysis similar to Velodrome [14]. PCD uses the read/write ordering information to replay accesses in an order that reflects the actual execution order. As PCD simulates replaying execution from logs, it tracks the last access(es) to each field:

- $\mathcal{W}$(f) is the last transaction to write field f.
- $\mathcal{R}$(T,f) is the last transaction of thread T to read field f.

PCD constructs a *precise dependence graph* (PDG) based on the last-access information. A helper function thread(tx) returns the thread that executes transaction tx. At each read or write of a

```
procedure READ(f, tx)
    if W(f) ≠ null and thread(tx) ≠ thread(W(f)) then
        Add PDG edge: W(f) → tx
    end if
    R(T,f) := tx                              ▷ Update last read for T
end procedure

procedure WRITE(f, tx)
    if W(f) ≠ null and thread(tx) ≠ thread(W(f)) then
        Add PDG edge: W(f) → tx
    end if
    for all T,  R(T,f) ≠ null do
        if thread(R(T,f)) ≠ thread(tx) then
            Add PDG edge: R(T,f) → tx
        end if
    end for
    W(f) := tx                                ▷ Update last write
    ∀ T,  R(T,f) := null                      ▷ Clear all reads
end procedure
```

Figure 5. Rules to update last-access information for a read and write of field f by a transaction tx.

field f, the analysis (1) adds a cross-thread edge to the PDG (if a dependence exists) and (2) updates the last-access information of f, as shown in Figure 5.

PCD detects cycles in the PDG after adding each cross-thread edge. A detected cycle indicates a precise atomicity violation. As part of the error log, PCD reports all the transactions and edges involved in the precise PDG cycle. For example, in Figure 3, PCD processes an IDG cycle of size four, computes the PDG, and identifies a precise cycle with just two transactions, Tx_{1i} and Tx_{3k}.

PCD aids debugging by performing *blame assignment* [14], which "blames" a transaction for an atomicity violation if its outgoing edge is created *earlier* than its incoming edge, implying that the transaction *completes* a cycle. In Figure 3, PCD blames Tx_{1i}.

4. Implementation

We have implemented a prototype of DoubleChecker in Jikes RVM 3.1.3 [1], a high-performance Java virtual machine (JVM) that provides performance competitive with commercial JVMs.[4] Our implementation builds on the publicly available Octet implementation [3]. For comparison purposes, we have also implemented Velodrome in Jikes RVM. Flanagan et al.'s implementation [14] is not available, and in any case it is implemented on top of the JVM-independent RoadRunner framework [12], so its performance characteristics could differ significantly. We have made our implementations of DoubleChecker and Velodrome publicly available on the Jikes RVM Research Archive.[5]

Specifying atomic regions. DoubleChecker takes an atomicity specification as input, specified as a list of methods to be *excluded* from the specification; all other methods are part of the specification, i.e., they are expected to execute atomically. DoubleChecker extends Jikes RVM's dynamic compilers so each compiled method is statically either *transactional* or *non-transactional*. Methods specified as atomic are always transactional, and non-atomic methods are compiled as transactional or non-transactional depending on the caller's context. The compilers compile two versions of non-atomic methods called from both contexts.

Constructing the IDG. The compilers insert instrumentation to start and end transactions in each *atomic* method called from a *non-transactional* context. At method start, instrumentation creates a new regular transaction. At method end, it creates a new unary transaction. While each non-transactional access conceptually executes in its own unary transaction, our implementation reuses prior work's optimization [14], which merges consecutive unary transactions not interrupted by an incoming or outgoing cross-thread edge.

Each transaction maintains (1) a list of its outgoing edges to other transactions and (2) (for single-run mode or the second run of multi-run mode) a read/write log that is an ordered list of precise memory access entries. Each read/write log entry records information about one access: the base object reference, field address, and read versus write. The read/write log has special entries that correspond to incoming and outgoing cross-thread edges, since PCD needs to know the access order with respect to cross-thread edges.

Transactions and their read/write logs are regular Java objects in our implementation, so garbage collection (GC) naturally collects them as they become transitively unreachable from each thread's current transaction reference. The implementation treats read/write log entries as weak references to avoid memory leaks. When a reference in a read/write log entry dies, our modified GC replaces the reference in the log with the old field address and the current GC invocation count, distinguishing the field precisely.

Instrumenting program accesses. The compilers add read and write barriers at (object and static) field and array accesses. Our experiments focus on evaluating only instrumenting field accesses and only in application (not Java library) methods, which imitates closely related prior work [11, 14], although we also evaluate the performance of instrumenting array accesses. The compilers instrument program synchronization by treating acquire operations like object reads, and release operations like object writes.

In single-run mode or the second run of multi-run mode, ICD adds instrumentation to record read/write logs. Although logs are ordered, duplicate entries with no incoming or outgoing edges between them can be elided to save space. To elide duplicate entries on the fly, ICD tracks, for each field, the value of a per-thread timestamp of the last access (and whether it was a read or write) to the object; RdSh objects have up to one timestamp per thread. Every time a new transaction starts, or a transaction has an incoming or outgoing edge, a thread increments its current timestamp. It stores this information in per-field metadata for WrEx and RdEx objects and per-thread hash tables for RdSh objects.

Velodrome implementation. Our DoubleChecker and Velodrome implementations share features as much as possible: they instrument the same accesses, demarcate transactions the same way, and represent dependence graphs the same way. The Velodrome implementation does not use Octet. It adds two words for each object and static field: one references the transaction to write the field, and the other references the last transaction(s) (up to one per thread) to read the field since the last write. To capture release–acquire dependences, each object has an extra header word to track the last transaction to release the object's lock. The implementation treats metadata references as weak references to avoid memory leaks in the transaction dependence graph.

At each access, instrumentation detects cross-thread dependences, adds them to the dependence graph, detects cycles (and reports a precise atomicity violation for each cycle), and updates the field's last-access information. To provide atomicity of the instrumentation together with the program access and thus track cross-thread dependences accurately, the instrumentation and access execute in a small critical section that "locks" a word of the field's metadata using an atomic operation.

[4] http://dacapo.anu.edu.au/regression/perf/9.12-bach.html

[5] http://www.jikesrvm.org/Research+Archive

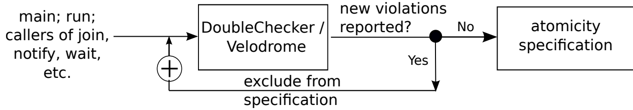

Figure 6. Iterative refinement methodology to generate a program's atomicity specification.

5. Evaluation

This section evaluates the correctness and performance of our prototype implementation of DoubleChecker in both single- and multi-run modes and compares with Velodrome.

5.1 Methodology

Benchmarks. Our experiments run the following programs: the multithreaded DaCapo benchmarks [2] that Jikes RVM 3.1.3 can execute successfully: eclipse6, hsqldb6, lusearch6, xalan6, avrora9, jython9, luindex9, lusearch9,[6] pmd9, sunflow9, and xalan9 (suffixes '6' and '9' distinguish benchmarks from versions 2006-10-MR2 and 9.12-bach, respectively). We also execute the following programs used in prior work [11, 14]: the microbenchmarks elevator, hedc, philo, sor, and tsp [34]; and moldyn, montecarlo, and raytracer from the Java Grande benchmark suite [30].

Experimental setup. We build a high-performance configuration of the JVM (FastAdaptive) that adaptively optimizes the application and uses a high-performance, generational, stop-the-world garbage collector. We let the JVM adjust the heap size automatically. Our experiments use the *small* workload size of the DaCapo benchmarks, since otherwise DoubleChecker's single-run mode and (to a lesser extent) Velodrome run out of memory with larger workload sizes for a few benchmarks. DoubleChecker's single-run mode also runs out of memory with the standard small size of moldyn and raytracer, so we modify the benchmarks to use an even smaller input, which all atomicity checkers execute. The JVM, which targets the IA-32 platform, is limited to a heap of roughly 1.5–2 GB; a 64-bit implementation could avoid these out-of-memory errors. For the other benchmarks, we use the same input configurations described in prior work [11, 14].

For DoubleChecker's multi-run mode, we execute 10 trials of the first run, take the union of the transactions reported as part of ICD cycles, and use it as input for the second run. This methodology represents a point in the accuracy–performance tradeoff that we anticipate would be used in practice: combining information from multiple first runs should help a second run find more atomicity violations but also increase its overhead.

Platform. The experiments execute on a workstation with a 3.30 GHz 4-core Intel i5 processor, with 4 GB memory running 64-bit RedHat Enterprise Linux 6.5, kernel 2.6.32.

Deriving atomicity specifications. Atomicity specifications for the benchmarks either have not been determined by prior work (DaCapo) or are not publicly available (non-DaCapo). We derive specifications for all the programs using an *iterative refinement* methodology used successfully by prior work [11, 13, 14, 35]. Figure 6 illustrates iterative refinement. First, all methods are assumed to be atomic with a few exceptions: top-level methods (e.g., main() and Thread.run()) and methods that contain interrupting calls (e.g., to wait() or notify()). We also exclude the DaCapo benchmarks' driver thread (which launches worker threads that actually run the benchmark program) from the atomicity specification, since we know it executes non-atomically. Iterative refinement repeatedly removes methods from the specification when they are "blamed" for detected atomicity violations. We terminate iterative refinement only when no new atomicity violations are reported after 10 trials, in order to approximate well-tested software, which has an accurate atomicity specification and *few, if any, known* atomicity violations.

[6] We use a version of lusearch9 that fixes a serious memory leak [39].

	Velodrome		DoubleChecker		
	Total	(Unique)	Single-run	Multi-run	(Unique)
eclipse6	230	(8)	244	190	(8)
hsqldb6	10	(0)	57	57	(0)
lusearch6	1	(0)	1	1	(0)
xalan6	57	(0)	69	54	(0)
avrora9	23	(0)	25	18	(0)
jython9	0	(0)	0	0	(0)
luindex9	0	(0)	0	0	(0)
lusearch9	41	(1)	40	38	(0)
pmd9	0	(0)	0	0	(0)
sunflow9	13	(1)	13	13	(0)
xalan9	78	(0)	82	69	(0)
elevator	2	(0)	2	2	(0)
hedc	3	(1)	3	2	(0)
philo	0	(0)	0	0	(0)
sor	0	(0)	0	0	(0)
tsp	7	(0)	7	7	(0)
moldyn	0	(0)	0	0	(0)
montecarlo	2	(0)	2	2	(0)
raytracer	0	(0)	0	0	(0)
Total	467	(11)	545	453	(8)

Table 2. Static atomicity violations reported by our implementations of Velodrome and DoubleChecker. For Velodrome and multi-run mode, *Unique* counts how many violations were *not* reported by single-run mode.

We use iterative refinement in two ways. First, we use it to evaluate the soundness of DoubleChecker's single- and multi-run modes by comparing the set of atomicity violations reported by Velodrome and DoubleChecker's single- and multi-run modes (Section 5.2). For each of the three configurations, we perform iterative refinement to completion and collect all methods blamed as non-atomic along the way.

Second, we use iterative refinement to determine the *final* specifications, i.e., specifications that lead to few or no atomicity violations, in order to evaluate performance (Section 5.3). To prepare the final specification for each program, we take the intersection of the finalized specifications (no more violations reported in 10 trials) for both Velodrome and DoubleChecker (single-run mode, since it is fully sound by design), to avoid any bias toward one approach.

We adjust the specifications in a few cases because of out-of-memory errors. raytracer and sunflow9 have one and two long-running transactions, respectively, that execute atomically and that ICD passes to PCD, causing PCD to run out of memory, so we exclude the corresponding methods from the specifications. On the flip side, refining the specification of xalan6 leads to so many transactions being created that DoubleChecker run out of memory, so we use an intermediate (not fully refined) specification for xalan6.

5.2 Soundness

DoubleChecker's *single-run* mode is sound and precise by design. By comparing it to Velodrome, we sanity-check our implementations while also observing the effect of timing differences between the two algorithms. *Multi-run* mode is not fully sound, so by comparing it to Velodrome and single-run mode, we evaluate how sound it is in practice. A caveat of this section's comparison is that the first and second runs use the same program inputs, thus representing an upper bound on soundness guarantees.

Table 2 shows, for each atomicity checker, the total number of violations reported during all steps of iterative refinement. Each violation in Table 2 represents a method identified by blame assignment at least once during this process. At a given step of iterative refinement, a violation reported in one trial may not always be reported in other trials, due to nondeterministic thread

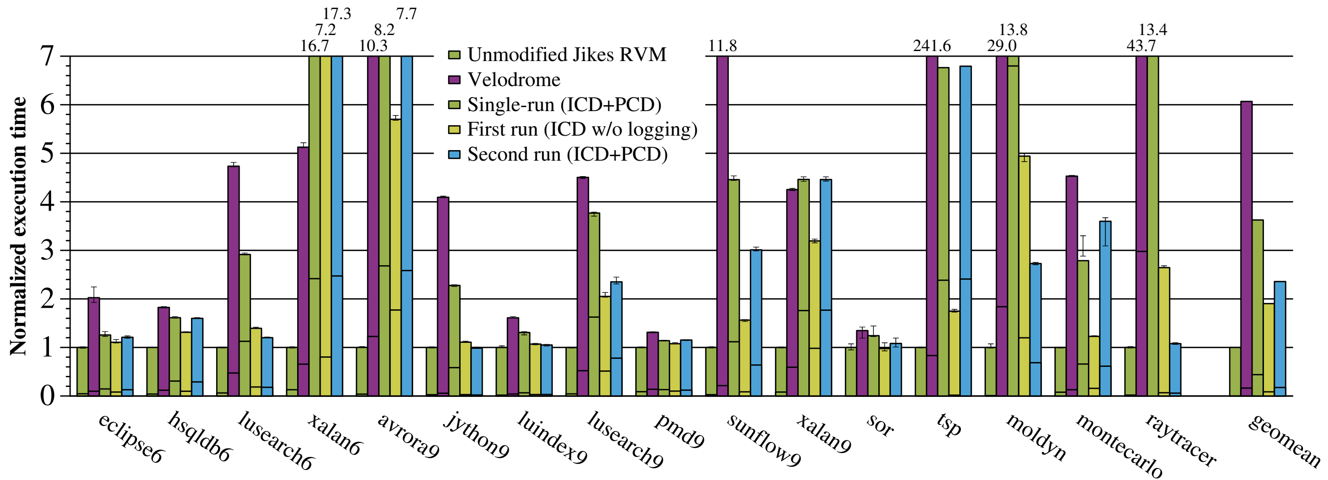

Figure 7. Run-time performance comparisons of Velodrome, DoubleChecker's single-run mode, and the first and second runs of DoubleChecker's multi-run mode. The sub-bars show GC time. The geomean GC time excludes short-running sor, which never triggers GC.

interleavings. Overall, the violations reported by Velodrome and DoubleChecker's single-run mode match closely. Both implementations are sound and precise, so (assuming correct implementations) differences are due to different thread interleavings caused by run-to-run nondeterminism and timing differences between the two analyses. The *Unique* value in parentheses counts violations reported by Velodrome but not by single-run mode; it is nonzero for just four programs, indicating single-run mode finds nearly all violations found by Velodrome. Single-run also mode finds several violations not found by Velodrome. We investigated the program with the greatest discrepancy, hsqldb6. By inserting random timing delays in Velodrome, we were able to reproduce six violations reported by DoubleChecker, providing some evidence that differences are due to timing effects.

Not surprisingly, multi-run mode does not detect quite as many violations as the sound single-run mode. Overall, multi-run modes detects 83% of all violations detected by single-run mode. Normalizing the detection rate across programs with at least one violation, multi-run mode detects 90% of a program's violations on average, which may be worthwhile in exchange for multi-run mode's lower run-time overhead (discussed next). Multi-run mode finds violations not found by single-run mode only for eclipse6; some but not all of these are the same violations found by Velodrome but not single-run mode.

5.3 Performance

This section compares the performance of Velodrome, DoubleChecker's single-run mode, and the first and second runs of DoubleChecker's multi-run mode. We use the final specifications for our performance experiments (Section 5.1), and exclude elevator, hedc, and philo, since they are not compute bound [14].

Figure 7 shows the execution time of Jikes RVM running various configurations of the Velodrome and DoubleChecker implementations. Execution times are normalized to the first configuration, *Unmodified Jikes RVM*. Each bar is the median of 25 trials to minimize effects of any machine noise. We also show the mean as the center of 95% confidence intervals. Sub-bars show the fraction of time taken by GC.

Velodrome. Our implementation of *Velodrome* slows programs by 6.1X on average. This result corresponds with the 12.7X slowdown reported in prior work [14], although it is hard to compare results since we implement Velodrome in a JVM and use an overlapping but different set of benchmarks. Comparing results for the benchmarks evaluated by prior work, we find that our implementation adds substantially more overhead in several cases. In particular, the Velodrome paper reports 71.7X, 4.5X, and 9.2X slowdowns for tsp, moldyn, and raytracer, respectively [14]. It is somewhat surprising that our implementation in a JVM would add more overhead than the dynamic bytecode instrumentation approach by the Velodrome authors [12, 14]. By running various partial configurations, we find that 82% of this overhead comes from synchronization required to provide analysis–access atomicity, which is unsurprising since atomic operations can lead to remote cache misses on otherwise mostly-read-shared accesses.

According to the Velodrome authors [16], their implementation eschews synchronization when metadata does not actually need to change, i.e., the current transaction is already the last writer or reader. We have implemented and evaluated this variant, which is unsound and can miss dependences in the presence of concurrent accesses, and in fact it crashes on avrora9 due to races accessing metadata. This unsound variant slows programs by 4.1X on average, providing the most help to the programs afflicted most. DoubleChecker still outperforms this unsound variant of Velodrome.

DoubleChecker's single-run mode. The remaining configurations in Figure 7 are for DoubleChecker. *Single-run (ICD+PCD)* shows the time incurred to run ICD and PCD in the same execution. This configuration slows programs by 3.6X (260% overhead) on average. Using partial configurations, we find that about two-fifths of this overhead comes from Octet, building the IDG, and detecting IDG cycles. (This partial configuration is similar to the first run of multi-run mode, presented next.) Logging read and write accesses as part of ICD accounts for nearly all of the remaining overhead. Less than one-tenth of the overhead on average comes from PCD, since ICD filters out most transactions. Single-run mode spends a high fraction of time in GC for several programs, largely because of the memory footprint of long-lived read/write logs. Overall, DoubleChecker's single-run mode avoids much of Velodrome's costs and adds 1.9 times less overhead than Velodrome.

Velodrome outperforms DoubleChecker's single-run mode for one program, xalan6. When executing xalan6, ICD detects many imprecise dependences, triggering SCC detection frequently, and SCC detection finds many imprecise SCCs, leading to high PCD overhead. ICD detects SCCs serially, and PCD detects cycles serially; making them parallel could alleviate this bottleneck.

	Single-run mode (or first run of multi-run mode)					Second run of multi-run mode				
	# Instrumented					# Instrumented				
Name	Regular transactions	Regular accesses	Non-trans. accesses	# IDG edges	# ICD SCCs	Regular transactions	Regular accesses	Non-trans. accesses	# IDG edges	# ICD SCCs
eclipse6	793,000	137,000,000	6,610,000	68,400	124	617,000	46,400,000	7,100,000	38,900	80
hsqldb6	87,000	13,400,000	147,000	26,400	76	86,400	10,100,000	148,000	26,200	75
lusearch6	95,700	143,000,000	1,440,000	17	0	0	0	0	0	0
xalan6	1,140,000	70,400,000	17,500,000	211,000	15,500	1,170,000	70,900,000	16,900,000	211,000	15,700
avrora9	22,100,000	264,000,000	362,000,000	2,310,000	854	9,260,000	122,000,000	363,000,000	2,340,000	932
jython9	8	53,200,000	29	0	0	0	0	0	0	0
luindex9	7	8,610,000	25	0	0	0	0	0	0	0
lusearch9	813,000	115,000,000	27,100,000	141	6	64,900	13,500,000	0	142	8
pmd9	7	2,650,000	25	0	0	0	0	0	0	0
sunflow9	35,000	263,000,000	129,000	1,080	25	10,600	176,000,000	129,000	1,020	24
xalan9	1,580,000	67,000,000	14,500,000	66,500	444	1,480,000	66,500,000	15,100,000	67,000	457
elevator	3,200	17,000	5,590	419	24	3,180	16,100	5,590	427	23
hedc	79	38,400	114	83	3	25	37,200	114	85	3
philo	6	16	458	144	0	0	0	0	0	0
sor	2	16	18,700	189	0	0	0	0	0	0
tsp	12,000	386,000	694,000,000	14,100	0	1,340	6,650	691,000,000	11,500	0
moldyn	573,000	194,000,000	50,500,000	38	0	0	0	0	0	0
montecarlo	102,000	179,000,000	93,300,000	30,600	2,860	89,700	145,000,000	108,000,000	30,800	2,730
raytracer	25,700	890,000,000	108,000,000	215	1	4	113	0	9	1

Table 3. Run-time characteristics of DoubleChecker for the single-run and the second run in the multi-run mode. Each average is rounded to a whole number with at most three significant digits.

DoubleChecker's multi-run mode. *First run (ICD w/o logging)* executes only ICD, without logging accesses. Its functionality is similar to a subset of single-run mode evaluated above, and its overhead is unsurprising: it slows programs by 1.9X (90% overhead) on average. The first run of multi-run mode is significantly faster than single-run mode because the former avoids logging.

Second run (ICD+PCD) executes both ICD and PCD (similar to single-run mode), except it only instruments transactions statically identified by the first run, and it instruments non-transactional accesses if and only if the first run identified any non-transactional accesses involved in cycles. The second run slows programs by 2.4X (140% overhead) on average.

We also evaluate a configuration of the second run that always instruments non-transactional accesses—regardless of whether the first run detected any cycles involving unary transactions. Overhead increases to 169%, justifying conditional instrumentation of non-transactional accesses during the second run.

Since the first run detects few imprecise cycles, one might expect the second run would have little work to do. However, the first run identifies transactions *statically* by method signature, leading to many more (dynamic) instrumented accesses in the second run than the total number of (dynamic) accesses identified as involved in cycles in the first run. The filter for non-transactional accesses is even coarser; the second run must instrument all non-transactional accesses in many cases. For this reason, DoubleChecker's ICD and PCD analyses perform better than using Velodrome for the second run, i.e., ICD is still effective as a dynamic transaction filter in the second run. Using Velodrome for the second run (i.e., instrumenting only the transactions statically identified by the first run) slows programs by 2.9X.

A promising direction for future work is to devise an effective way for the first run to more precisely communicate potentially imprecise cycles to the second run.

Summary. Overall, DoubleChecker substantially outperforms prior art. The single-run mode is a fully sound and precise atomicity checker that adds 1.9 times less overhead than Velodrome. Multi-run mode does not guarantee soundness, since atomicity checking is split between two runs, but it avoids the need for logging all program accesses (which the single-run mode needs in order to perform a precise analysis). The first and second runs of the multi-run mode add 5.6 and 3.7 times less overhead than Velodrome, respectively, providing a performance–accuracy tradeoff that is likely worthwhile in practice for providing more acceptable overhead for checking atomicity. DoubleChecker's significant performance benefits justify our design's key insights: it is indeed cheaper to track cross-thread dependences imprecisely in order to filter most of a program's execution from processing by a precise analysis.

5.4 Other Performance Investigations

The performance results so far use refined atomicity specifications that lead to no reported atomicity violations. Here we measure the performance of DoubleChecker during iterative refinement. At the beginning of iterative refinement (i.e., the strictest specification), DoubleChecker's single-run mode slows execution by 3.4X. Halfway through iterative refinement (after the first half of the non-atomic methods have been removed from the specification), single-run mode's slowdown is 3.6X. These slowdowns compare closely with the slowdown after full refinement (3.6X), suggesting that performance during iterative refinement is similarly reasonable.

The experiments so far evaluate implementations of DoubleChecker and Velodrome that do not instrument array accesses, since the Velodrome paper also omits this instrumentation [14]. Here we evaluate the additional overhead from array instrumentation. For implementation simplicity, DoubleChecker and Velodrome conflate all elements of an array by using array-level metadata instead of element-level metadata. This makes not only ICD but also Velodrome imprecise, so we disable cycle detection for both analyses. DoubleChecker's single-run mode then runs out of memory for xalan6 and xalan9, so we exclude these programs. DoubleChecker's single-run mode's average slowdown increases from 3.1X (without array instrumentation) to 3.7X (with array instrumentation), and Velodrome's slowdown increases from 6.3X to 7.3X. Note that all four slowdowns skip cycle detection and exclude xalan6 and xalan9.

Finally, to check whether ICD is beneficial as a first-pass filter, we use a "PCD-only" variant of single-run mode in which PCD processes *every* executed transaction, not just transactions identified by ICD's imprecise cycle detection. The PCD-only variant—which is something of a straw man since PCD essentially implements a less-efficient version of Velodrome's algorithm—increases

the slowdown of the single-run mode from 3.1X to 16.6X on average (both results exclude eclipse6, xalan6, avrora9, and xalan9, since the PCD-only variant runs out of memory when running them). This result confirms that ICD is essential as a first-pass filter for PCD.

5.5 Run-Time Characteristics

Table 3 shows execution characteristics of ICD in single-run mode (the first run of multi-run mode should yield similar results) and the second run of multi-run mode. Each value is the mean of 10 trials of a special statistics-gathering configuration of DoubleChecker; otherwise methodology is the same as Figure 7. For each of the two configurations, the table reports the number of regular (non-unary) transactions and accesses instrumented in both regular and unary transactions, and the number of cross-thread edges and SCCs in the IDG. Single-run mode instruments everything, while the second run instruments a subset of transactions. For several programs, the second run avoids instrumenting any accesses because the first run reports no SCCs. For lusearch9 and raytracer, the second run avoids instrumenting any non-transactional accesses since no first-run SCC contained a unary transaction, but non-transactional accesses are instrumented for all the other benchmarks. For programs where the second run instruments (nearly) all transactions and accesses (e.g., xalan6 and avrora9), there is little benefit from multi-run mode's optimization. Even when they should be the same, the counts sometimes differ across modes due to run-to-run nondeterminism.

Compared to how many memory accesses execute, there are few ICD edges, justifying ICD's approach that optimistically assumes accesses are not involved in cross-thread dependences. There are few ICD SCCs in most cases, justifying DoubleChecker's dual-analysis approach and explaining why PCD adds low overhead (except for xalan6; Section 5.3).

6. Related Work

This section details other static and dynamic analyses besides the most closely related work, Velodrome [14] (Section 2).

Checking conflict serializability. Farzan and Parthasarathy introduce a dynamic serializability-checking analysis based on finding cycles among transactions [9]. Their analysis bounds space overhead optimally so space is *not* proportional to the length of the run, by summarizing the dependence graph as transactions finish. DoubleChecker and Velodrome do not summarize the dependence graph, but they do rely on garbage collection (GC) to collect transactions not transitively reachable from any thread's current transaction ("last access" references from objects are treated as weak references), which reduces space overhead in practice. Still, DoubleChecker can add substantial space overhead, especially to maintain single-run mode's read/write logs. Future work might be able to apply summarization to DoubleChecker to reduce space overhead.

While Velodrome and DoubleChecker detect cycles online as the program executes, Farzan and Parthasarathy's analysis detects cycles offline, i.e., after the execution finishes. They compare their summarized dependence graph to an *unsummarized* dependence graph—but this unsummarized graph does not use GC, so its space overhead is unavoidably proportional to the length of the run.

Farzan and Parthasarathy's analysis does *not* track synchronization edges [9]. In contrast, DoubleChecker, which follows Velodrome [14], tracks synchronization edges—but tracking synchronization edges can lead to false positives when checking conflict serializability.

Other dynamic analyses. Some dynamic approaches are *predictive*, so they detect atomicity violations not only for the current execution's thread interleavings, but also for other interleavings that could have executed [29, 31]. Predictive analyses tend to be considerably more expensive than non-predictive analysis, particularly the more sound and/or precise they are.

Wang and Stoller propose dynamic analyses for checking atomicity based on detecting unserializable patterns [35, 36]. These approaches are predictive since they aim to find potential violations in other executions, but this process is inherently imprecise, so they may report false positives.

Atomizer is a dynamic atomicity checker that uses a variation of the lockset algorithm [28] to determine shared variables that can have racy accesses, and monitors those variables for potential atomicity violations. Atomizer reports false positives since it relies on the locket algorithm.

Static analyses. Static approaches can check all inputs soundly, but they are imprecise, and in practice they do not scale well to large programs nor to dynamic language features such as dynamic class loading. Type systems can help check atomicity but require a combination of type inference and programmer annotations [13, 15]. Model checking does not scale well to large programs because of state space explosion [8, 10, 19]. Static approaches are well suited to analyzing critical sections but not wait–notify synchronization.

Alternatives. *HAVE* combines static and dynamic analysis to obtain benefits of both approaches [4]. Because HAVE speculates about untaken branches, it can report false positives.

Several approaches infer an atomicity specification automatically [6, 18, 24, 32, 37]. However, these approaches are inherently unsound and imprecise. Furthermore, many of these approaches add high overhead to track cross-thread dependences, e.g., AVIO slows programs by more than an order of magnitude [24].

Prior work *exposes* atomicity violations by making them more likely to occur [26, 27] and thus more likely for a non-predictive analysis to detect. Exposing atomicity violations is complementary to checking atomicity.

Transactional memory enforces programmer-specified atomicity annotations by speculatively executing atomic regions as *transactions*, which are rolled back if a region conflict occurs [20]. *Atom-Aid* relies on custom hardware to execute regions atomically and to detect some atomicity violations [25].

Static analysis can infer needed locks automatically from an atomicity specification (e.g., [5]). The inferred locks are inherently imprecise, leading to overly conservative locking.

In summary, most prior work gives up on soundness or precision or both, and dynamic approaches slow programs substantially. In contrast, DoubleChecker checks conflict serializability soundly and precisely and significantly outperforms other dynamic approaches.

7. Conclusion

This paper presents a new direction for dynamic sound and precise atomicity checking that divides work into two cooperating analyses: a lightweight analysis that detects cross-thread dependences, and thus atomicity violations, soundly but imprecisely; and a precise analysis that focuses on potential cycles and rules out false positives. These cooperating analyses can execute in a single run, or the imprecise analysis can run alone and inform a second run, providing a performance–soundness tradeoff. DoubleChecker outperforms existing sound and precise atomicity checking, suggesting that its new direction has the potential to enable more widespread use of atomicity checking in the real world.

Acknowledgments

We thank Azadeh Farzan, Madhusudan Parthasarathy, Sanket Tavarageri, Harry Xu, Jaeheon Yi, and the anonymous reviewers for feedback on the text; and Man Cao, Cormac Flanagan, Stephen

Freund, Milind Kulkarni, Jaeheon Yi, and Minjia Zhang for discussions and advice.

References

[1] B. Alpern, S. Augart, S. M. Blackburn, M. Butrico, A. Cocchi, P. Cheng, J. Dolby, S. Fink, D. Grove, M. Hind, K. S. McKinley, M. Mergen, J. E. B. Moss, T. Ngo, and V. Sarkar. The Jikes Research Virtual Machine Project: Building an Open-Source Research Community. *IBM Systems Journal*, 44:399–417, 2005.

[2] S. M. Blackburn, R. Garner, C. Hoffman, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S. Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J. E. B. Moss, A. Phansalkar, D. Stefanović, T. VanDrunen, D. von Dincklage, and B. Wiedermann. The DaCapo Benchmarks: Java Benchmarking Development and Analysis. In *OOPSLA*, pages 169–190, 2006.

[3] M. D. Bond, M. Kulkarni, M. Cao, M. Zhang, M. Fathi Salmi, S. Biswas, A. Sengupta, and J. Huang. Octet: Capturing and Controlling Cross-Thread Dependences Efficiently. In *OOPSLA*, pages 693–712, 2013.

[4] Q. Chen, L. Wang, Z. Yang, and S. D. Stoller. HAVE: Detecting Atomicity Violations via Integrated Dynamic and Static Analysis. In *FASE*, pages 425–439, 2009.

[5] S. Cherem, T. Chilimbi, and S. Gulwani. Inferring Locks for Atomic Sections. In *PLDI*, pages 304–315, 2008.

[6] L. Chew and D. Lie. Kivati: Fast Detection and Prevention of Atomicity Violations. In *EuroSys*, pages 307–320, 2010.

[7] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*, chapter 11. The MIT Press, McGraw-Hill Book Company, 2nd edition, 2001.

[8] A. Farzan and P. Madhusudan. Causal Atomicity. In *CAV*, pages 315–328, 2006.

[9] A. Farzan and P. Madhusudan. Monitoring Atomicity in Concurrent Programs. In *CAV*, pages 52–65, 2008.

[10] C. Flanagan. Verifying Commit-Atomicity Using Model-Checking. In *SPIN*, pages 252–266, 2004.

[11] C. Flanagan and S. N. Freund. Atomizer: A Dynamic Atomicity Checker for Multithreaded Programs. *SCP*, 71(2):89–109, 2008.

[12] C. Flanagan and S. N. Freund. The RoadRunner Dynamic Analysis Framework for Concurrent Programs. In *PASTE*, pages 1–8, 2010.

[13] C. Flanagan, S. N. Freund, M. Lifshin, and S. Qadeer. Types for Atomicity: Static Checking and Inference for Java. *TOPLAS*, 30(4):20:1–20:53, 2008.

[14] C. Flanagan, S. N. Freund, and J. Yi. Velodrome: A Sound and Complete Dynamic Atomicity Checker for Multithreaded Programs. In *PLDI*, pages 293–303, 2008.

[15] C. Flanagan and S. Qadeer. A Type and Effect System for Atomicity. In *PLDI*, pages 338–349, 2003.

[16] S. Freund, 2013. Personal communication.

[17] J. L. Greathouse, Z. Ma, M. I. Frank, R. Peri, and T. Austin. Demand-Driven Software Race Detection using Hardware Performance Counters. In *ISCA*, pages 165–176, 2011.

[18] C. Hammer, J. Dolby, M. Vaziri, and F. Tip. Dynamic Detection of Atomic-Set-Serializability Violations. In *ICSE*, pages 231–240, 2008.

[19] J. Hatcliff, Robby, and M. B. Dwyer. Verifying Atomicity Specifications for Concurrent Object-Oriented Software using Model-Checking. In *VMCAI*, pages 175–190, 2004.

[20] M. Herlihy and J. E. B. Moss. Transactional Memory: Architectural Support for Lock-Free Data Structures. In *ISCA*, pages 289–300, 1993.

[21] L. Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. *CACM*, 21(7):558–565, 1978.

[22] B. R. Liblit. *Cooperative Bug Isolation*. PhD thesis, University of California at Berkeley, 2004.

[23] S. Lu, S. Park, E. Seo, and Y. Zhou. Learning from Mistakes: A Comprehensive Study on Real World Concurrency Bug Characteristics. In *ASPLOS*, pages 329–339, 2008.

[24] S. Lu, J. Tucek, F. Qin, and Y. Zhou. AVIO: Detecting Atomicity Violations via Access-Interleaving Invariants. In *ASPLOS*, pages 37–48, 2006.

[25] B. Lucia, J. Devietti, K. Strauss, and L. Ceze. Atom-Aid: Detecting and Surviving Atomicity Violations. In *ISCA*, pages 277–288, 2008.

[26] C.-S. Park and K. Sen. Randomized Active Atomicity Violation Detection in Concurrent Programs. In *FSE*, pages 135–145, 2008.

[27] S. Park, S. Lu, and Y. Zhou. CTrigger: Exposing Atomicity Violation Bugs from Their Hiding Places. In *ASPLOS*, pages 25–36, 2009.

[28] S. Savage, M. Burrows, G. Nelson, P. Sobalvarro, and T. Anderson. Eraser: A Dynamic Data Race Detector for Multi-Threaded Programs. In *SOSP*, pages 27–37, 1997.

[29] A. Sinha, S. Malik, C. Wang, and A. Gupta. Predictive Analysis for Detecting Serializability Violations through Trace Segmentation. In *MEMOCODE*, pages 99–108, 2011.

[30] L. A. Smith, J. M. Bull, and J. Obdrzálek. A Parallel Java Grande Benchmark Suite. In *SC*, pages 8–8, 2001.

[31] F. Sorrentino, A. Farzan, and P. Madhusudan. PENELOPE: Weaving Threads to Expose Atomicity Violations. In *FSE*, pages 37–46, 2010.

[32] W. N. Sumner, C. Hammer, and J. Dolby. Marathon: Detecting Atomic-Set Serializability Violations with Conflict Graphs. In *RV*, pages 161–176, 2012.

[33] U.S.–Canada Power System Outage Task Force. Final Report on the August 14th Blackout in the United States and Canada. Technical report, Department of Energy, 2004.

[34] C. von Praun and T. R. Gross. Static Conflict Analysis for Multi-Threaded Object-Oriented Programs. In *PLDI*, pages 115–128, 2003.

[35] L. Wang and S. D. Stoller. Accurate and Efficient Runtime Detection of Atomicity Errors in Concurrent Programs. In *PPoPP*, pages 137–146, 2006.

[36] L. Wang and S. D. Stoller. Runtime Analysis of Atomicity for Multi-threaded Programs. *IEEE TSE*, 32:93–110, 2006.

[37] M. Xu, R. Bodík, and M. D. Hill. A Serializability Violation Detector for Shared-Memory Server Programs. In *PLDI*, pages 1–14, 2005.

[38] X. Yang, S. M. Blackburn, D. Frampton, and A. L. Hosking. Barriers Reconsidered, Friendlier Still! In *ISMM*, pages 37–48, 2012.

[39] X. Yang, S. M. Blackburn, D. Frampton, J. B. Sartor, and K. S. McKinley. Why Nothing Matters: The Impact of Zeroing. In *OOPSLA*, pages 307–324, 2011.