

Automatic Synchronization Correction for Atomicity

Cormac Flanagan
Department of Computer Science
University of California, Santa Cruz
Santa Cruz, CA 95064

Stephen N. Freund
Department of Computer Science
Williams College
Williamstown, MA 01267

ABSTRACT

Multithreaded programs are notoriously prone to synchronization errors. Much prior work has tackled the problem of detecting such errors. This paper focuses on the subsequent problem of *synchronization correction*. We present a constraint-based analysis that, given an erroneous program, automatically infers (where possible) what additional locking operations should be inserted in order to yield a correctly-synchronized program. For performance reasons, our algorithm also attempts to minimize the number of additional lock acquires and the duration for which the acquired locks are held. We present experimental results that validate this approach on a number of standard Java library classes.

1. INTRODUCTION

Multithreaded programs are notoriously prone to errors due to incorrect synchronization. Earlier work in this area focused on detecting synchronization errors that cause race conditions [24, 23, 28, 29, 1, 4, 9, 11], atomicity violations [14, 8, 10, 12, 19, 31], and other consistency violations [2, 30]. Of course, *detecting* defects is useful only if it is followed by a second step of *defect correction*. In this work, we focus on this subsequent defect correction phase, and in particular on the problem of providing automated support for correcting synchronization errors.

We start with programs in which some methods have been documented as `atomic`, which means that they should include sufficient synchronization so that their execution is *serializable*, i.e., they can be considered to execute without interleaved actions of concurrent threads. If such a program fails to satisfy its atomicity specification, our tool attempts to modify the program (by introducing additional synchronization operations) so that it does satisfy its specification. In addition, we try to minimize the number of additional locks acquired (to reduce synchronization overhead) and to only hold these locks for short durations (to reduce lock contention [3]).

While our approach will not fix all concurrency errors, it does work well on a large and important class of defects. Our experience with atomicity checking tools has demonstrated that, in large programs, a small number of synchronization errors can lead to a large number of atomicity errors, possibly in many different modules. Identifying the root causes of these atomicity violations can be time consuming and may

require subtle knowledge about program’s intended synchronization structure.

This paper presents a *synchronization correction algorithm* that, in addition to identifying atomicity violations, also identifies likely causes for these atomicity violations and (where possible) suggests additional synchronization operations that are sufficient to correct these violations. This correction algorithm performs a deep analysis on the program’s synchronization structure. In previous work, we (1) developed *Rcc/Sat*, an analysis for inferring protecting locks for each field [11], and (2) a constraint-based framework for inferring the most precise atomicity for each method [12]. In this paper, we extend this line of research to tackle the more difficult problem of synchronization correction.

Inserting additional synchronization can introduce the potential for deadlock. We do not explicitly reason about deadlock, and instead assume the programmer ensures deadlock-freedom by code inspection or via some other static or dynamic analysis.

We present preliminary results that validate this analysis on three standard Java library classes that have synchronization errors. For each class, our analysis can automatically infer the additional synchronization operations that are necessary to correct these defects. We also performed defect-injection experiments, which demonstrated that the analysis is capable of correcting the vast majority of randomly-inserted synchronization defects in Java library classes.

2. MOTIVATING EXAMPLE

We illustrate the behavior of our analysis on the example shown in Figure 1, in which a `Stack` is represented as a linked `List` of `Elements`. The underlined annotations specify the program’s locking structure. For example, the field `List.elems` is guarded by `this`, the implicit lock of the `List` object. This lock also guards the fields `List.elems.num` and `List.elems.next`, since `elems` has type `Elements(this)`, and in `Elements` the ghost parameter `x` guards the `num` and `next` fields. These underlined type annotations can be automatically inferred by the *Rcc/Sat* type inference algorithm [11].

The class `Stack` is intended to be thread-safe, and so its `push` and `dup` methods are declared as `atomic`. However, the given program violates this atomicity specification. For example, `push` and `dup` call `add` without acquiring any locks, resulting in race conditions. Synchronizing on `this` inside `add` corrects this race condition and ensures that `push` is atomic, but `dup` is still incorrect, since a concurrent threads could modify the stack between `dup`’s two calls to `add`. Thus,

Figure 2: Class Stack with Atomicity Annotations and Corrected Synchronization

a) Program with Atomicity Annotations	b) Program with Correct Synchronization
<pre> class Elem(ghost x) { int num guarded_by x; Elem(x) next guarded_by x; } class List { Elem(this) elems guarded_by this; <u>α₁</u> void add(int v) { <u>sync?t1 (this)</u> this.elems = new Elem(this)(v,this.elems); } <u>α₂</u> int removeFirst() { <u>sync?t2 (this)</u> sync(this) { <u>sync?t3 (this)</u> let int x = this.elems.num in { <u>sync?t4 (this)</u> { <u>sync?t5 (this)</u> this.elems = this.elems.next; <u>sync?t6 (this)</u> return x; } } } } } class Stack { final List data; atomic void push(int x) { <u>sync?t7 (this)</u> <u>sync?t8 (this.data)</u> this.data.add(x); } atomic int dup() { <u>sync?t9 (this)</u> <u>sync?t10 (this.data)</u> let int x = this.data.removeFirst() in { <u>sync?t11 (this)</u> <u>sync?t12 (this.data)</u> { <u>sync?t13 (this)</u> <u>sync?t14 (this.data)</u> this.data.add(x); <u>sync?t15 (this)</u> <u>sync?t16 (this.data)</u> this.data.add(x); } } } } </pre>	<pre> class Elem(ghost x) { int num guarded_by x; Elem(x) next guarded_by x; } class List { Elem(this) elems guarded_by this; <u>this?mover:atomic</u> void add(int v) { <u>sync(this)</u> this.elems = new Elem(this)(v,this.elems); } <u>this?mover:atomic</u> int removeFirst() { sync(this) { let int x = this.elems.num in { this.elems = this.elems.next; return x; } } } } class Stack { final List data; atomic void push(int x) { this.data.add(x); } atomic int dup() { <u>sync(this.data)</u> let int x = this.data.removeFirst() in { this.data.add(x); this.data.add(x); } } } </pre>

inferring the necessary synchronization operations requires reasoning about both race conditions and atomicity violations.

Given this incorrect program, our algorithm automatically generates the corrected program of Figure 2(b), where the methods `add` and `dup` contain additional synchronization. (Alternatively, inserting synchronization into `push` and `dup` would also suffice.) Our algorithm also infers the most precise atomicity for each unannotated method. For example, the method `add` is assigned the conditional atomicity

“`this?mover:atomic`”, which states that if the lock `this` is already held, then `add` is a `mover` (that is, its execution commutes with steps of concurrent threads); if `this` is not held, then `add` is `atomic` (that is, it can be assumed to execute in a serial manner, without interleaved steps of concurrent threads).

Our analysis works in four phases.

1. The first phase encloses each program statement e in

Figure 1: Class Stack and Inferred Locking Annotations

```

class Elem(ghost x) {
  int num guarded_by x;
  Elem(x) next guarded_by x;
}

class List {
  Elem(this) elems guarded_by this;

  void add(int v) {
    this.elems = new Elem(this)(v, this.elems);
  }

  int removeFirst() {
    sync(this) {
      let int x = this.elems.num in {
        this.elems = this.elems.next;
        return x;
      }
    }
  }
}

class Stack {
  final List data;

  atomic void push(int x) {
    this.data.add(x);
  }

  atomic int dup() {
    let int x = this.data.removeFirst() in {
      this.data.add(x);
      this.data.add(x);
    }
  }
}

```

the tagged synchronization operations:

$$\text{sync?}t_1(l_1) \{ \dots \{ \text{sync?}t_n(l_n) \{ e \} \} \dots \}$$

where l_1, \dots, l_n are all candidate locks in scope, and the t_i are tags that uniquely identify each inserted synchronization operation. The candidate locks for e are all expressions up to a certain fixed size that are valid constant expressions in the static scope of e . This phase also inserts atomicity variables α_i for methods without declared atomicities, such as `add`, producing the program shown in Figure 2(a).

The goal of our analysis is then to determine a set T of tags denoting which of these tagged synchronization operations are necessary and sufficient to yield a correctly synchronized program. As part of its reasoning, the analysis also needs to infer an *assignment* A mapping atomicity variables α_i to atomicities.

2. The second phase of the analysis translates the program with tagged synchronization operations (as in Figure 2(a)) into a collection of constraints \bar{C} over the tag set T and assignment A .
3. The third phase solves these constraints using an iterative least fixed point algorithm, to yield the program of Figure 2(a) with the grayed-out synchronization operations removed. This program is correctly

synchronized, but still contains some unnecessary synchronization operations.

4. The fourth phase then identifies and removes these redundant synchronization operations, yielding the final program of Figure 2(b).

This constraint-based approach extends our earlier work on detecting synchronization errors [12], but for synchronization correction the addition of tagged synchronization operations requires an extended constraint language and new constraint solving algorithms.

The presentation of our results proceeds as follows. Section 3 describes the idealized Java subset that we used to formalize our analysis. Section 4 presents our constraint language. Section 5 describes how to generate constraints, which are then solved by the algorithm of Section 6. This solution is then optimized in Section 7. Our implementation and experiments are described in Section 8 and Section 9. Section 10 discusses related work, and Section 11 concludes.

3. THE SOURCE LANGUAGE AJC

We formalize our ideas in terms of the idealized language AJC (AtomJava with Correction), whose syntax is summarized in Figure 3. For simplicity, the idealized language AJC does not support subclassing, although it is supported in our implementation [12].

As in Java, each field declaration includes the name and type of the field. Additionally, in AJC, each field also has an associated *guard* g , which states that the field is either (1) **final**, (2) **unguarded**, or (3) **guarded_by** some lock l , which must be held at each access (*i.e.*, read or write) to that field. Sometimes the fields of a class need to be protected by a lock external to the class. For this purpose, each AJC class declaration includes a binding for a sequence of *ghost variables* denoting locks that can be used to protect fields of the class. Ghost variables are only used during type checking and do not exist at run time. A class type $cn\langle l^* \rangle$ includes a class name cn and an appropriate number of lock parameters for that class. The *Rcc/Sat* type inference algorithm can infer appropriate guards and lock parameters for unannotated programs [11].

A method declaration can also include a number of ghost variable bindings, for which corresponding lock expressions must be provided at call sites. Each method declaration also includes an *atomicity specification* s , such as **atomic** or **this?mover:atomic**, as described in Section 4.4.

AJC expressions include object allocation, field access and update, method invocation, variable binding and reference, conditionals, and while loops. Object allocation `newy c(e*)` includes a sequence of expressions used to initialize the object fields. For technical reasons, the `new` keyword is subscripted by y , which is a ghost variable bound to the object being created while evaluating the field initialization expressions. This enables the types of the initialization expressions to refer to the new object. We omit the binding from examples when it is not needed.

The language supports multiple threads of control via the construct `e.fork`. Here, e should evaluate to an object with a nullary `run` method, which is called by a newly-spawned thread.

Synchronization between threads is achieved via the construct `sync l e`, which executes by first evaluating l to yield

Figure 3: AJC Syntax

$P ::= \text{defn}^* e$	(program)
$\text{defn} ::= \text{class } cn \langle \text{ghost } x^* \rangle \text{ body}$	(class decl.)
$\text{body} ::= \{ \text{field}^* \text{meth}^* \}$	(class body)
$\text{field} ::= c \text{fn } g$	(field decl.)
$g ::= \text{final} \mid \text{guarded_by } l$	
$\mid \text{unguarded}$	(field guards)
$\text{meth} ::= s \text{c mn} \langle \text{ghost } x^* \rangle (\text{arg}^*) \{ e \}$	(method decl.)
$\text{arg} ::= c \text{ } x$	(arg. decl.)
$c ::= cn \langle l^* \rangle$	(class type)
$l ::= e$	(lock expr.)
$e ::= x \mid \text{null} \mid \text{new}_y c(e^*)$	(expressions)
$\mid e.f d \mid e.f d = e \mid e.mn \langle l^* \rangle (e^*)$	
$\mid \text{let } c \text{ } x = e \text{ in } e \mid \text{while } e \text{ } e \mid \text{if } e \text{ } e \text{ } e$	
$\mid \text{sync } l \text{ } e \mid e.\text{fork} \mid \text{sync}?t \text{ } l \text{ } e$	
$x, y \in \text{Var}$	$t \in \text{Tag}$
$cn \in \text{ClassName}$	$fn \in \text{FieldName}$
$mn \in \text{MethodName}$	

an object reference; the implicit lock associated with that object is then acquired; the expression e is then evaluated, and finally the lock is released. AJC also contains a tagged synchronization construct $\text{sync}?t \text{ } l \text{ } e$, where t is a unique tag identifying that operation.

Although omitted from the formal system for simplicity, our examples use integers and sequential composition, which we treat in the expected fashion.

4. ATOMICITY CONSTRAINTS

4.1 Basic Atomicities

Our approach to verifying atomicity is based on classifying expressions according to what actions they may perform, and in particular how these actions may interact with actions of concurrent threads. For this purpose, we introduce the following five *basic atomicities*:

$b ::= \text{const} \mid \text{mover} \mid \text{atomic} \mid \text{cmpd} \mid \text{error}$

The informal meaning of these basic atomicities is as follows:

- **const** expressions do not access any mutable state, and hence always yield the same result when evaluated in the same environment. Such expressions may include calls to **const** methods, etc.
- **mover** expressions can access mutable state, but only if that mutable state is either local to the current thread or protected by a lock held by the current thread. In addition, **mover** expressions cannot acquire or release locks. Thus, **mover** expressions commute with actions of concurrent threads.
- **atomic** expressions, in addition to accessing thread-local or protected state, may also acquire and release locks according to the two-phase locking discipline. That is, nested synchronization, as in:

$\text{sync } l_1 \{ \dots \text{sync } l_2 \{ \dots \} \dots \}$

is **atomic**, but the sequential composition of synchronization operations, as in:

$\text{sync } l_1 \{ \dots \}; \text{sync } l_2 \{ \dots \}$

is not **atomic**. By Lipton's theory of reduction [20], **atomic** expressions are serializable, and are therefore amenable to sequential reasoning techniques, which significantly facilitates subsequent formal and informal reasoning [14]. Our previous investigation showed that the vast majority of methods in multithreaded Java programs are atomic [10, 12].

- **cmpd** expressions do not follow the two-phase locking discipline, but which do hold the correct protecting lock when accessing any guarded field. Such expressions are not necessarily serializable and are therefore not amenable to sequential reasoning.
- **error** expressions violate the program's synchronization discipline by accessing a field without holding the guarding lock. Programs with **error** expressions do not type check.

Basic atomicities are ordered by the subatomicity relation \sqsubseteq_b :

$\text{const} \sqsubseteq_b \text{mover} \sqsubseteq_b \text{atomic} \sqsubseteq_b \text{cmpd} \sqsubseteq_b \text{error}$

Let \sqcup_b denote the corresponding join operator for basic atomicities. Suppose that the basic atomicities b_1 and b_2 reflect the behavior of e_1 and e_2 respectively. Then:

- The atomicity $b_1 \sqcup_b b_2$ reflects the non-deterministic choice between executing either e_1 or e_2 .
- The *sequential composition* $b_1; b_2$ reflects the behavior of executing $e_1; e_2$, and is defined as:

$$b_1; b_2 = \begin{cases} \text{cmpd} & \text{if } b_1 = b_2 = \text{atomic} \\ b_1 \sqcup_b b_2 & \text{otherwise} \end{cases}$$

- The *iterative closure* b_1^* reflects the behavior of executing e_1 an arbitrary number of times, and is defined as:

$$b_1^* = \begin{cases} \text{cmpd} & \text{if } b_1 = \text{atomic} \\ b_1 & \text{otherwise} \end{cases}$$

LEMMA 1. *The operations $b_1 \sqcup_b b_2$ and $b_1; b_2$ and b_1^* are monotonic in both b_1 and b_2 .*

4.2 Atomicity Expressions

We follow a constraint-based approach to atomicity inference and synchronization correction. For each method body, we generate an *atomicity expression* d that encodes the various operations performed by that method body. The syntax of atomicity expressions (see Figure 4) includes the following constructs:

- Basic atomicities.
- Atomicity variables α , which support atomicity inference.

An *assignment* maps atomicity variables to closed atomicity expressions (that is, to atomicity expressions that do not contain atomicity variables):

$\alpha \in \text{AtomVar}$
 $A \in \text{Assignment} = \text{AtomVar} \rightarrow \text{ClosedAtomExp}$

Figure 4: Atomicity Expressions

$d ::= b \mid \alpha \mid d;d \mid d \sqcup d \mid d^* \mid l?d:d \quad (\text{AtomExp})$ $\begin{array}{l} \mid \mathcal{S}(l,d) \mid \mathcal{R}(t,l,d) \\ \mid d \cdot \theta \mid wfa(P,E,d) \end{array}$	(AtomExp)
$\theta ::= [\vec{x} := \vec{l}] \quad (\text{substitution})$	(substitution)

- Sequential composition, join, and iterative closure of atomicity expressions (which correspond to sequential composition, branching, and looping operations in the original code).
- The *conditional atomicity* $l?d_1:d_2$, which is equal to d_1 if the lock l is held, and is equal to d_2 otherwise. For example, the atomicity of an access to a field protected by lock l is $l? \text{mover} : \text{error}$, formalizing that an access to the field has atomicity **mover** if l is held, and has atomicity **error** otherwise.
- The atomicity expression $\mathcal{S}(l,d)$ yields the atomicity for a synchronized expression **sync** $l e$, where d is the atomicity expression for e .
- The construct $\mathcal{R}(t,l,d)$ is generated for each tagged synchronization operation **sync?** $t l e$. If the inferred tag set T includes the tag t , then this synchronization operation is chosen for insertion, and $\mathcal{R}(t,l,d)$ is equivalent to $\mathcal{S}(l,d)$. If the inferred tag set does not include t , then $\mathcal{R}(t,l,d)$ is equivalent to d . More formally, the application of a tag set T to an atomicity expression is the compatible closure of the following function:

$$T(\mathcal{R}(t,l,d)) = \begin{cases} \mathcal{S}(l,d) & \text{if } t \in T \\ d & \text{otherwise} \end{cases}$$

- The *delayed substitution* operation $d \cdot \theta$ is used when changing scopes, where the substitution θ is a finite map from variables to lock expressions. For example, if a method's atomicity refers to a formal method parameter, then this formal parameter must be replaced by the corresponding actual parameter to derive the correct atomicity for a call to that method.
- The construct $wfa(P,E,d)$ yields the smallest (*i.e.*, most precise) atomicity that is at least as large as d and that only depends on locks that are in scope in the environment E . This construct is used when d may refer to a variable that is going out of scope, such as at the end of a **let** construct.

An atomicity expression is *closed* if it does not contain atomicity variables. An atomicity expression is *tag-free* if it does not contain the construct $\mathcal{R}(t,l,d)$.

4.3 Atomicities

We formalize the meaning of a closed atomicity expression as a map from the set of lock held by the current thread to a basic atomicity, and we refer to this map as an *atomicity*.

$$\begin{aligned} L \in \text{LockSet} &= {}_2\text{Lock} \\ a \in \text{Atomicity} &= \text{LockSet} \rightarrow \text{BasicAtomicity} \end{aligned}$$

The function $\llbracket \cdot \rrbracket$, defined in Figure 5, maps closed, tag-free atomicity expressions to atomicities. (For clarity, we write " $\llbracket d \rrbracket_L = \dots$ " to abbreviate " $\llbracket d \rrbracket = \lambda L. \dots$ ".)

Figure 5: Atomicity Meaning Function

$\llbracket \cdot \rrbracket$	$: \text{ClosedAtomExp} \rightarrow \text{Atomicity}$
$\llbracket b \rrbracket_L$	$= b$
$\llbracket d_1;d_2 \rrbracket_L$	$= \llbracket d_1 \rrbracket_L; \llbracket d_2 \rrbracket_L$
$\llbracket d_1 \sqcup d_2 \rrbracket_L$	$= \llbracket d_1 \rrbracket_L \sqcup_a \llbracket d_2 \rrbracket_L$
$\llbracket d^* \rrbracket_L$	$= (\llbracket d \rrbracket_L)^*$
$\llbracket l?d_1:d_2 \rrbracket_L$	$= \begin{cases} \llbracket d_1 \rrbracket_L & \text{if } l \in L \\ \llbracket d_2 \rrbracket_L & \text{otherwise} \end{cases}$
$\llbracket \mathcal{S}(l,d) \rrbracket_L$	$= \begin{cases} \llbracket d \rrbracket_L & \text{if } l \in L \\ \llbracket d \rrbracket_{L \cup \{l\}} & \text{if } \text{atomic} \sqsubseteq_a \llbracket d \rrbracket_{L \cup \{l\}} \\ \text{atomic} & \text{otherwise} \end{cases}$
$\llbracket d \cdot \theta \rrbracket_L$	$= \llbracket d \rrbracket_{L'}$ where $L' = \{x \mid \theta(x) \in L\}$
$\llbracket wfa(P,E,d) \rrbracket_L$	$=$ (see Section 5)

We order atomicities according to the point-wise extension \sqsubseteq_a of the ordering relation \sqsubseteq_b on basic atomicities, with corresponding minimal element \perp_a and join operation \sqcup_a . We also extend the sequential composition and iterative closure operations to atomicities in a point-wise manner.

$$\begin{aligned} a_1 \sqsubseteq_a a_2 &\text{ iff } \forall L \subseteq \text{Lock}. (a_1(L) \sqsubseteq_b a_2(L)) \\ \perp_a &\stackrel{\text{def}}{=} \lambda L. \text{const} \\ a_1 \sqcup_a a_2 &\stackrel{\text{def}}{=} \lambda L. (a_1(L) \sqcup_b a_2(L)) \\ a^* &\stackrel{\text{def}}{=} \lambda L. (a(L))^* \\ a_1; a_2 &\stackrel{\text{def}}{=} \lambda L. (a_1(L); a_2(L)) \end{aligned}$$

We order assignments according to the point-wise extension \sqsubseteq_A of the ordering relation \sqsubseteq_a on atomicities, with corresponding minimal element \perp_A and join operation \sqcup_A :

$$\begin{aligned} A_1 \sqsubseteq_A A_2 &\text{ iff } \forall \alpha. (\llbracket A_1(\alpha) \rrbracket \sqsubseteq_a \llbracket A_2(\alpha) \rrbracket) \\ A_1 =_A A_2 &\text{ iff } \forall \alpha. (\llbracket A_1(\alpha) \rrbracket = \llbracket A_2(\alpha) \rrbracket) \\ \perp_A &\stackrel{\text{def}}{=} \lambda \alpha. \text{const} \\ A_1 \sqcup_A A_2 &\stackrel{\text{def}}{=} \lambda \alpha. (A_1(\alpha) \sqcup_A A_2(\alpha)) \end{aligned}$$

We extend assignments in a compatible manner to be maps from atomicity expressions to atomicity expressions. Given an A and a tag-free atomicity expression d , the atomicity expression $A(d)$ is then closed and tag-free, and its meaning is the atomicity $\llbracket A(d) \rrbracket$. Furthermore, this meaning function is monotonic in A (a necessary prerequisite for performing our least fixpoint analysis).

LEMMA 2 (MONOTONICITY). *For all tag-free atomicity expressions d , the function $\lambda A. \llbracket A(d) \rrbracket$ is monotonic.*

PROOF. By induction on the structure of d . \square

4.4 Atomicity Constraints

A *syntactic atomicity* s is either an atomicity variable or a closed, tag-free atomicity expression. Each method declaration includes a corresponding syntactic atomicity. For example, a programmer could specify a precise atomicity for a method, such as **atomic**. More commonly, the programmer might omit this specification, in which case the type checker uses a fresh atomicity variable as the method's declared atomicity.

Given a program P , for each method with declared atomicity s and method body e , we generate an atomicity expression d for the method body e (as described in the following section) and produce the *constraint* $d \sqsubseteq s$. Thus, a

constraint is a subatomicity relation between an atomicity expression and a syntactic atomicity. An assignment A satisfies a constraint $d \sqsubseteq s$ with respect to a tag set T (written $A; T \models C$) if

$$\llbracket A(T(d)) \rrbracket \sqsubseteq_a \llbracket A(s) \rrbracket$$

Generating a constraint for each method yields a constraint set \bar{C} for the entire program. The assignment A satisfies a set \bar{C} with respect to T (written $A; T \models \bar{C}$) if A satisfies each constraint in \bar{C} . Thus, our goal is to find an assignment A and a tag set T such that A satisfies \bar{C} with respect to T . The tag set T then specifies which tagged synchronization operations need to be inserted into the program in order to yield a corrected program that satisfies the desired atomicity specifications.

Since the type checker introduces a fresh atomicity variable for each unannotated method, each atomicity variable α annotates exactly one method, and so α has a unique lower bound d such that the constraint $d \sqsubseteq \alpha$ occurs in \bar{C} . We use the notation $\bar{C}(\alpha)$ to refer to this lower bound d . We refer to such constraints with a variable as an upper bound as *propagation constraints*, and we refer to all other constraints (with a closed atomicity expression as upper bound) as *checking constraints*.

5. PHASE 2: CONSTRAINT GENERATION

We express the algorithm for converting a program with previously-inserted tagged synchronization operations into a collection of constraints \bar{C} as a set of rules for reasoning about the judgment:

$$P; E \vdash e : c \cdot d \cdot \bar{C}$$

Here, c is the type inferred for the expression e ; d is the atomicity expression generated for e ; and \bar{C} is the set of constraints generated from this expression. The program P is included to provide access to class declarations, and E is an environment providing types for the free program and ghost variables of the expression e :

$$E ::= \epsilon \mid E, c \ x \mid E, \mathbf{ghost} \ x$$

The complete set of type judgments and rules is contained in Figure 6. We briefly describe some of the more important rules.

[LOCK EXP] The judgment $P; E \vdash_{\text{lock}} l : \bar{C}$ checks that l is a well-formed lock expression in environment E . The lock expression l can be either a ghost variable or a program expression e . In the latter case, e must denote a fixed lock throughout the execution of the program to ensure soundness. Thus, we require that e has atomicity **const**.

In addition, we require the size $|e|$ of the lock expression to be bounded by the constant $MaxLockSize$. This requirement ensures that there is only a finite number of valid lock expressions at any program point, which in turn bounds the size of conditional atomicities and number of possible tagged synchronization statements. This ensures termination of our correction algorithm.

[EXP VAR] A variable access has **const** atomicity, since all variables are immutable in AJC.

[EXP IF] The atomicity of a conditional expression is the atomicity of the *test* subexpression, sequentially composed with the join of the atomicities of the *then* and *else* branches.

[EXP LET] This rule for **let** $x = e_1$ **in** e_2 infers atomicity expressions d_1 and d_2 for e_1 and e_2 , respectively. Since the atomicity expression d_2 may refer to the let-bound variable x , we apply the substitution $\theta = [x := e_1]$ to yield a corresponding atomicity that does not mention x . Several complications arise here.

First, since d_2 may include an atomicity variable α , we cannot apply the substitution θ immediately because α may later resolve to x . Instead, we use the *delayed substitution form* $d_2 \cdot \theta$ to delay this substitution until after atomicity variables are resolved.

Second, e_1 may not be **const** (in general, we cannot determine which expressions are **const** until after type inference), in which case $d_2 \cdot \theta$ may not be a valid atomicity. Therefore, we use the well-formed atomicity construct $wfa(P, E, d_2 \cdot \theta)$ to yield a valid atomicity for e_2 that is well-formed in environment E . The meaning of this construct is defined via:

$$\llbracket wfa(P, E, d) \rrbracket_L = \llbracket d \rrbracket_{L'} \quad \text{where } L' = \left\{ l \in L \mid \begin{array}{l} P; E \vdash_{\text{lock}} l : \bar{C} \\ \text{and } \perp_A; \emptyset \models \bar{C} \end{array} \right\}$$

As described above, the judgment $P; E \vdash_{\text{lock}} l$ checks that l is a well-formed lock expression in an environment E and program P provided the constraints \bar{C} holds. Since the meaning function $\llbracket \cdot \rrbracket_L$ is defined only on closed, tag-free atomicity expressions, \bar{C} will also be closed and tag-free, and so we check that is satisfiable via $\perp_A; \emptyset \models \bar{C}$. Thus, L' contains only locks that are held and that are valid in the environment E , and d is evaluated in the context of these held and valid locks.

[EXP REF] The rule for a field access $e.fn$ first checks that e is of some type $cn\langle l_{1..n} \rangle$, and that cn is a class parameterized by n ghost variables, say $x_{1..n}$, that declares a field fn of some type t . The type t may refer to the variables **this** and $x_{1..n}$ in scope at the field declaration. Since these variables are not in scope at the field access, the type rule introduces a substitution θ that replaces them with the corresponding expressions e and $l_{1..n}$, and ensures that $\theta(t)$ is a well-formed type.

The [EXP REF] rule performs a case analysis on the field's guard. If the field is **final**, then the read operation has atomicity **const**, since there can be no concurrent writes. If the field is **unguarded**, then the read operation is **atomic**, since it may not commute with concurrent writes. If the field is **guarded_by** l , then the lock $\theta(l)$ must be held and the read operation is a **mover**.

[EXP SYNC] The rule for the synchronized statement **sync** $l \ e$ checks that l has atomicity **const**, and so always denotes the same lock. The rule then yields the atomicity expression $\mathcal{S}(l, d)$, where d is the atomicity of e . The meaning $\llbracket \mathcal{S}(l, d) \rrbracket_L$ of this atomicity expression is either (1) $\llbracket d \rrbracket_L$, if the lock is already held; (2) $\llbracket d \rrbracket_{L \cup \{l\}}$, if d is non-atomic; or (3) **atomic** otherwise.

[EXP SYNC-OPT] The rule for the tagged synchronized statement **sync?** $t \ l \ e$ is similar, and yields the atomicity expression $\mathcal{R}(t, l, d)$, which either behaves like $\mathcal{S}(l, d)$ or d , depending on whether this synchronization statement is enabled or disabled by the tag set T .

Figure 7: Constraints for Stack Program

$\mathcal{R}(t1, \text{this}, (\text{const}; (\text{const}; (\text{const}; \text{wfa}(P, E_1, \text{this?mover: error}))); \text{wfa}(P, E_1, \text{this?mover: error})))$	$\sqsubseteq \alpha_1$
$\mathcal{R}(t2, \text{this}, \mathcal{S}(\text{this}, \mathcal{R}(t3, \text{this}, \text{const}; \text{wfa}(P, E_2, \text{this?mover: error}); \text{wfa}(P, E_2, \text{this?mover: error});$ $\text{wfa}(P, E_2, (\mathcal{R}(t4, \text{this}, \mathcal{R}(t5, \text{this}, \text{const}; \text{wfa}(P, E_3, \text{this?mover: error});$ $\text{wfa}(P, E_3, \text{this?mover: error});$ $\text{wfa}(P, E_3, \text{this?mover: error})));$	
$\mathcal{R}(t6, \text{this}, \text{const})) \cdot \theta_1))))$	$\sqsubseteq \alpha_2$
$\mathcal{R}(t7, \text{this}, \mathcal{R}(t8, \text{this.data}, (\text{const}; \text{const}); \text{const}; \text{wfa}(P, E_4, \alpha_1 \cdot \theta_2)))$	$\sqsubseteq \text{atomic}$
$\mathcal{R}(t9, \text{this}, \mathcal{R}(t10, \text{this.data}, (\text{const}; \text{const}; \text{wfa}(P, E_5, \alpha_2 \cdot \theta_3));$ $\text{wfa}(P, E_5, \mathcal{R}(t11, \text{this}, \mathcal{R}(t12, \text{this.data},$ $(\mathcal{R}(t13, \text{this}, \mathcal{R}(t14, \text{this.data}, (\text{const}; \text{const}); \text{const}; \text{wfa}(P, E_6, \alpha_1 \cdot \theta_2)));$ $\mathcal{R}(t15, \text{this}, \mathcal{R}(t16, \text{this.data}, (\text{const}; \text{const}); \text{const}; \text{wfa}(P, E_6, \alpha_1 \cdot \theta_2)))))) \cdot \theta_4))))$	$\sqsubseteq \text{atomic}$
$E_1 = \text{List this, int v}$	
$E_2 = \text{List this}$	
$E_3 = \text{List this, int x}$	
$E_4 = \text{Stack this, int v}$	
$E_5 = \text{Stack this}$	
$E_6 = \text{Stack this, int x}$	
$\theta_1 = [x := \text{this.elems.num}]$	
$\theta_2 = [\text{this} := \text{this.data}, v := x]$	
$\theta_3 = [\text{this} := \text{this.data}]$	
$\theta_4 = [x := \text{this.data.removeFirst()}]$	

[EXP FORK] This rule for `e.fork` ensures that `e` contains a `run` method with the signature:

$$s \text{ c' run(ghost t11)() } \{ \dots \}$$

This signature contains a special ghost parameter named `t11`. This lock is held throughout the entire lifetime of the new thread, and so may be used to protect data local to this thread. A `run` method need not be atomic, but if `t11` is held when `run` is invoked, the method should not violate the locking discipline. Thus, we have the constraint $s \sqsubseteq (\text{t11?cmpd: error})$.

[PROG] This rule defines the top-level judgment $P \vdash \bar{C}$, where \bar{C} is the generated set of constraints for the program P . This rule uses three predicates defined as follows. (See [15] for their precise definition.)

- *ClassOnce*(P): no class is declared twice in P .
- *FieldsOnce*(P): no field name is declared twice in a class.
- *MethodsOnce*(P): no method name is declared twice in a class.

The constraints for our example program are shown in Figure 7. (For simplicity, we omit several trivial checking constraints.) The first constraint is generated for the method `add`. The various `const` atomicities in that constraint correspond to variable accesses; the two conditional atomicities `this?mover: error` correspond to the access and update of `this.elems`, which is guarded by `this`; the enclosing $\mathcal{R}(t1, \text{this}, \dots)$ corresponds to the tagged synchronization construct `sync?t1 (this) ...`; and the upper bound α_1 is the atomicity specification for `add`. The third constraint is for `push`, and includes a reference to the atomicity specification α_1 of the callee `add`, with a substitution θ_2 that maps formal to actual parameters for this call site. The remaining constraints are similar, although more verbose.

6. PHASE 3: CONSTRAINT SOLVING

Having generated a constraint set \bar{C} over the tags and atomicity variables in the program, the next step is to solve these constraints. If the program does not contain any tagged synchronization operations, then the generated constraints are tag-free, and our earlier work presents an algorithm for solving such constraints [12]. In this paper, we

now tackle the harder problem of solving a constraint set \bar{C} in the case where some constraints may include the tagged synchronization construct $\mathcal{R}(t, l, d)$, which corresponds to an automatically-inserted tagged synchronization operation `sync?t l e` in the source program.

The behavior of these tagged synchronization constructs depends on the chosen tag set T . If $t \in T$, then a real synchronization operation is inserted into the original program at this point, and $\mathcal{R}(t, l, d)$ behaves exactly like $\mathcal{S}(l, d)$. If $t \notin T$, then this potential synchronization point is ignored, and $\mathcal{R}(t, l, d)$ behave like d .

We wish to find a suitable choice of tag set T and assignment A such that $A; T \models \bar{C}$. We use an iterative least fixed point algorithm to compute the minimal assignment A that satisfies \bar{C} , but this assignment will of course depend on the chosen tag set. Hence, on each iteration of the algorithm, we choose the tag set that yields the smallest possible assignment for use in the next iteration.

Choosing a tag set in this manner is possible if atomicity expressions are *well-formed*. An atomicity expression d is well-formed if (1) each tag occurs at most once in d , and (2) whenever d contains a conditional atomicity $l?d_1:d_2$ then

1. $\llbracket d_1 \rrbracket \sqsubseteq \llbracket d_2 \rrbracket$,
2. `atomic` $\sqsubseteq \llbracket d_2 \rrbracket$, and
3. d_1 and d_2 mention the same lock expressions and atomicity variables.

The first and second requirements provide a necessary monotonicity property for our optimization algorithm, namely that the atomicity of `sync l e` must be smaller than the atomicity of e (i.e., removing a synchronized statement cannot make an expression's atomicity become smaller). To motivate the second requirement, suppose e has atomicity $l?const:mover$. The atomicity of `sync l e` would be the larger atomicity $l?const:atomic$, and the monotonicity property would not hold. The third property exists for similar, but more subtle reasons.

The constraint generation rules only generate well-formed atomicity expressions, and well-formedness is preserved by the various operations we perform on atomicity expressions.

The function $\text{min}(d)$, defined in Figure 8 returns a tag set T that minimizes $\llbracket d \rrbracket$, where d is a closed, well-formed atomicity expression.

Figure 8: Tag Minimization Function

min	:	$ClosedAtomExpr \rightarrow 2^{Tag}$
$min(b)$	=	\emptyset
$min(d_1; d_2)$	=	$min(d_1) \cup min(d_2)$
$min(d_1 \sqcup d_2)$	=	$min(d_1) \cup min(d_2)$
$min(d^*)$	=	$min(d)$
$min(l? d_1 : d_2)$	=	$min(d_1) \cup min(d_2)$
$min(d \cdot \theta)$	=	$min(d)$
$min(S(l, d))$	=	$min(d)$
$min(wfa(P, E, d))$	=	$min(d)$
$min(\mathcal{R}(t, l, d))$	=	$\begin{cases} T \cup \{t\} & \text{if } [T(d)] \text{ depends on } l \\ T & \text{otherwise} \end{cases}$
		where $T = min(d)$

LEMMA 3. *Suppose that d is closed and well-formed, and let $T = min(d)$. Then T only contains tags that occur in d , and for all T' , $[T(d)] \sqsubseteq_a [T'(d)]$.*

PROOF. By structural induction on d . \square

The following function $f_{\bar{C}}$ describes each iteration of our algorithm. For each variable α , the function $f_{\bar{C}}(A)$ computes the closed atomicity expression $d = A(\bar{C}(\alpha))$, computes the tag set T that minimizes d , and then returns this minimal atomicity $[T(d)]$:

$$f_{\bar{C}} : Assignment \rightarrow Assignment$$

$$f_{\bar{C}}(A) = \lambda \alpha. [T(d)], \text{ where } d = A(\bar{C}(\alpha)) \text{ and } T = min(d)$$

Suppose A is the least fixpoint of $f_{\bar{C}}$, that is, $A = fix(f_{\bar{C}}, \perp_A)$, where we define the fixpoint operator to terminate once it reaches assignments that are semantically equivalent with respect to $=_A$:

$$fix(F, X) \stackrel{\text{def}}{=} \text{if } X =_A F(X) \text{ then } X \text{ else } fix(F, F(X))$$

The tag set T defined by

$$T = \cup \{min(A(d)) \mid (d \sqsubseteq s) \in \bar{C}\}$$

minimizes all atomicity expressions in \bar{C} . Now consider any propagation constraint $d \sqsubseteq \alpha$ in \bar{C} . We have that

$$A(\alpha) = f_{\bar{C}}(A)(\alpha) = [T(A(\bar{C}(\alpha)))] = [T(A(d))]$$

and so $A; T \models d \sqsubseteq \alpha$. Thus, A satisfies all propagation constraints in \bar{C} with respect to T , and it simply remains to check if A also satisfies the checking constraints in \bar{C} . The following function *solve* performs this analysis:

$$solve(\bar{C}) = \begin{cases} \langle A, T \rangle & \text{if } A = fix(f_{\bar{C}}, \perp_A) \text{ and } A; T \models \bar{C} \\ & \text{and } T = \cup \{min(A(d)) \mid (d \sqsubseteq s) \in \bar{C}\} \\ \text{undef} & \text{otherwise} \end{cases}$$

LEMMA 4.

1. *If $solve(\bar{C}) = \text{undef}$ then \bar{C} is unsatisfiable.*
2. *If $solve(\bar{C}) = \langle A, T \rangle$ then $A; T \models \bar{C}$.*

Proving that the *solve* algorithm terminates is non-trivial, because delayed substitutions could lead to arbitrarily large lock expressions and infinite ascending chains of atomicities and assignments. We bound the size of lock expressions to exclude this possibility. A lock expression l is *bounded* if $|l| <$

MaxLockSize. Similarly, an atomicity is *bounded* if it only contains bounded lock expressions, and an assignment is *bounded* if it only yields bounded atomicities. An atomicity expression or constraint is *bounded* if it is only conditional on bounded lock expressions, and every delayed substitution occurs inside the construct $wfa(P, E, \cdot)$. The *solve* algorithm terminates on the bounded constraint systems produced by the constraint generation rules.

THEOREM 5 (TERMINATION). *The constraint solving algorithm terminates on bounded constraint systems.*

PROOF. See [12]. \square

The algorithm computes the following solution for our Stack example:

$$A(\alpha_1) = A(\alpha_2) = \text{this?mover:atomic}$$

$$T = \{\mathbf{t1}, \mathbf{t2}, \mathbf{t3}, \mathbf{t4}, \mathbf{t5}, \mathbf{t8}, \mathbf{t10}, \mathbf{t12}, \mathbf{t14}, \mathbf{t16}\}$$

In particular, the algorithm concludes that the tagged synchronization operations $\mathbf{t6}$, $\mathbf{t7}$, $\mathbf{t9}$, $\mathbf{t11}$, $\mathbf{t13}$, and $\mathbf{t15}$ do not decrease the atomicity of their containing methods, and so do not contribute to yielding a correctly-synchronized program. These omitted tagged synchronization operations are grayed-out in Figure 2(a). The remaining tagged synchronization operations in T yield a correctly synchronized program, but incur an unnecessarily-large synchronization overhead, since many of them are redundant. The next section describes how to eliminate the redundant operations.

7. PHASE 4: SYNCHRONIZATION OPTIMIZATION

Having computed a tag set T and assignment A such that $\langle A, T \rangle = solve(\bar{C})$, it remains to relax the synchronization (by reducing T) while preserving the satisfiability of the constraints. For this purpose, we use a greedy algorithm, where the tags T are ordered by some heuristic into a worklist Z and are iteratively removed:

Figure 9: Phase 4 (version 1)

```

Z := T;
foreach z ∈ Z do
  T' := T \ {z};
  if T'(\bar{C}) is satisfiable then T := T';
end foreach

```

The correctness of this basic algorithm is fairly straightforward to verify. However, the algorithm performs a lot of redundant computation in checking the satisfiability of \bar{C} for various tag sets. As a first step towards optimizing this algorithm, we inline some of the operations being performed by this routine:

Figure 10: Phase 4 (version 2)

```

Z := T;
foreach z ∈ Z do
  T' := T \ {z};
  \bar{D} := T'(\bar{C});
  A' := fix(f_{\bar{D}}, \perp_A);
  if \forall (d \sqsubseteq s) \in \bar{D}. [A'(d)] \sqsubseteq_a [s] then T := T';
end foreach

```

It is now clear that the main performance overhead is in the repeated iterative fixpoint computation of $fix(f_{\bar{D}}, \perp_A)$,

which always begins with the least assignment \perp_A . However, there is a key monotonicity property that we can exploit. By maintaining an assignment A such that $A = \text{fix}(f_{T(\bar{C})}, \perp_A)$, we can more efficiently compute $\text{fix}(f_{T'(\bar{C})}, \perp_A)$, where $T' \subseteq T$, by starting from the current assignment A rather than the minimal assignment \perp_A . The following lemma ensures the correctness of this optimization.

LEMMA 6. *Suppose*

$$\begin{aligned} \langle A_0, T_0 \rangle &= \text{solve}(\bar{C}) & T' &= T \setminus \{z\} \\ A_0 &\sqsubseteq_A A & A &= \text{fix}(f_{T(\bar{C})}, \perp_A) \\ T &\subseteq T_0 & \bar{D} &= T'(\bar{C}) \end{aligned}$$

Then $A \sqsubseteq_A \text{fix}(f_{\bar{D}}, \perp_A)$ and hence $\text{fix}(f_{\bar{D}}, \perp_A) =_A \text{fix}(f_{\bar{D}}, A)$.

Our optimized algorithm is then:

Figure 11: Phase 4 (version 3)

```

Z := T;
foreach z ∈ Z do
  // invariant: A = fix(f_{T(\bar{C})}, \perp_A);
  T' := T \ {z};
  \bar{D} := T'(\bar{C});
  A' := fix(f_{\bar{D}}, A);
  if \forall (d \sqsubseteq a) \in \bar{D}. \llbracket A'(d) \rrbracket \sqsubseteq_a \llbracket a \rrbracket then
    A := A'; T := T';
  end if
end foreach
```

This algorithm computes the following assignment and tag set for the Stack program.

$$\begin{aligned} A(\alpha_1) = A(\alpha_2) &= \text{this?mover:atomic} \\ T &= \{\mathbf{t1}, \mathbf{t10}\} \end{aligned}$$

All redundant synchronization operations are now eliminated, and Figure 2(b) shows the code with the two synchronization operations that fix the program. Different orderings for the work list may yield different results. For example, the following is also a solution:

$$\begin{aligned} A(\alpha_1) &= \text{this?mover:error} \\ A(\alpha_2) &= \text{this?mover:atomic} \\ T &= \{\mathbf{t8}, \mathbf{t10}\} \end{aligned}$$

One interesting aspect of this algorithm is that we can order the work list to achieve specific effects. Trying to remove tags for outermost synchronization statements first may help reduce the size of critical sections, whereas removing tags for innermost statements first may improve performance by reducing the number of synchronization operations.

8. IMPLEMENTATION

We have implemented synchronization correction in *Bohr*, our tool for detecting and correcting errors in Java programs [12]. *Bohr* takes as input source code that may optionally contain atomicity specifications in stylized comments starting with “#”, as in “/*# mover */”. *Bohr* runs in two phases. The first phase uses the *Rcc/Sat* tool to infer appropriate guards for each field and appropriate formal and actual ghost parameters for class and method declarations and uses.

Rcc/Sat is somewhat resilient to errors and will infer the most likely synchronization information, even if a small number of race conditions do exist. In such cases, the original *Rcc/Sat* tool [11] would report potential race conditions in places where the inferred annotations are not satisfied. These warnings are not reported by *Bohr*, unless the race conditions also lead to atomicity violations.

If a program contains many synchronization errors, *Rcc/Sat* will not be able to infer reasonable protecting locks for fields, resulting in a significant number of methods becoming non-atomic. These atomicity warnings may not be fixable by *Bohr*, because it will not know which locks to acquire to provide safe access to fields lacking `guarded.by` annotations. However, this situation often indicates that more fundamental flaws exist in the code that may require design changes or restructuring to fix. The exact point at which *Rcc/Sat* fails to provide the necessary annotations is somewhat dependent on how the tool is configured. For more details on *Rcc/Sat*, we refer the interested reader to our earlier paper [11].

Our present work is included in the second phase of *Bohr*, which computes any unspecified atomicities and inserts synchronization operations to correct these errors. *Bohr* outputs a fully annotated version of the source code, including any synchronization statements necessary to satisfy the specified atomicity requirements. If errors cannot be fixed by our algorithm, the checker prints warning messages for each atomicity violation identified.

9. EVALUATION

We applied *Bohr* to three standard Java 1.4.2 library classes to validate its effectiveness at fixing defects. These three classes are designed to be thread-safe, meaning that all public methods should be `atomic`. The following table summarizes the results of running *Bohr* with and without synchronization correction:

Class	Lines	Atomicity Warnings	
		No Correction	With Correction
<code>String</code>	2,307	1	0
<code>StringBuffer</code>	1,276	1	0
<code>Vector</code>	3,546	3	0

The “Lines” column includes the size of the class of interest and all superclasses. We annotated all referenced library classes with appropriate atomicities. For simplicity, we assumed that all subclasses of `Collection` are internally synchronized in these experiments. No fundamental problems arise are preventing both internally and externally-synchronized subclasses of a single class, but the analysis becomes more complex [12].

Bohr successfully corrected for the warnings reported in our previous work [12]. The `StringBuffer` and `Vector` warnings are real defects. The `String` warning is caused by benign races and is spurious, although *Bohr* did add synchronization to remove the races. *Bohr*’s correction to `StringBuffer` did introduce a potential deadlock, but this cannot be avoided without restructuring the class. No other potential deadlocks were introduced during these experiments, or those reported below.

To further assess the effectiveness of *Bohr*, we performed an experiment on Java 1.4.2 library classes in which we added atomicity specifications, randomly removed a small

number of synchronization operations, and measured how many operations *Bohr* would correctly reinsert. Specifically, we randomly selected and removed one, two, or five synchronization operations and then counted the number of defects that *Bohr* identified and fixed, either by reinserting the removed synchronization statements or, in some cases, inserting different but operationally equivalent synchronization. We verified these results by manual inspection of the *Bohr* output.

We repeated each scenario 15 times, and the following table summarizes the percentage of inserted defects that were, on average, identified and fixed for each configuration:

Class	Lines	Time (sec)	Correction Rate per Number of Defects		
			1	2	5
Observable	198	0.73	100%	100%	100%
Inflater	319	0.60	100%	97%	91%
Deflater	384	0.90	100%	95%	88%
Zipfile	498	25.0	100%	100%	90%
StringBuffer	1,276	41.8	100%	88%	49%
String	2,307	26.2	100%	100%	—
Vector	3,456	49.6	100%	100%	85%
SynchronizedList	3,837	13.8	94%	85%	85%

No result is reported for `String` with five defects since that class has fewer than five synchronization operations. *Bohr* performed well when a small number of defects were introduced. For the larger numbers of defects, the precision declined because *Rcc/Sat* did not always infer the appropriate locking discipline, due to the large number of data races introduced on specific object fields. *Rcc/Sat* can be adjusted to withstand a higher number but, as mentioned above, the mere fact that a large number of races exists often indicates a fundamental problem with the design of the code (as opposed to more local programming errors). Interestingly, *Bohr* determined that several synchronization operations in these classes are redundant and unnecessary.

Results from applying *Bohr* to small, complete programs are similarly promising. The application of our current implementation to significantly larger programs is limited by the lack of precise atomicity specifications for libraries, and by incomplete support in our current implementation for some synchronization idioms such as protecting locks [12].

10. RELATED WORK

Since an atomicity annotation describes aspects of the behavior or effect of an expression, we are essentially performing a form of effect reconstruction [27, 26]. Our work differs from most of the work on effect systems and dependent types [6] by investigating automated techniques for correcting errors. Sasturkar *et al* [22] have also developed a type inference algorithm for atomicity. Unlike *Bohr*, their system includes a notion of object ownership [4] and uses a dynamic analysis to infer race condition information. They have not explored synchronization correction.

Lipton [20] first proposed reduction as a way to reason about deadlocks without considering all possible interleavings. Partial-order reduction techniques are based on similar ideas [16]. Several papers have used Lipton’s theory of reduction to improve the efficiency of model checking [5, 25, 13].

The use of model checking for verifying atomicity is being explored by Hatcliff *et al* [19]. This model checking approach is more expressive than our type-based analysis, but it is vulnerable to state-space explosion. Their results suggest that verifying atomicity via model-checking is feasible for unit-testing. A more general (but more expensive) technique for verifying atomicity during model checking is *commit-atomicity* [8]. Several tools have explored verifying atomicity dynamically [10, 31], but these tools are sensitive to test case coverage.

View consistency is another approach to preventing threads interference [2, 30]. 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. We believe view consistency could be extended with an analogous idea of synchronization correction.

Recent approaches to supporting atomicity include lightweight transactions [18, 32, 21, 17] and automatic generation of synchronization code from high-level specifications [7]. Lightweight transactions in particular seem like a promising, and complementary, approach to providing atomicity guarantees. An interesting avenue of future work is to explore how to best merge synchronization-based and transaction-based approaches.

11. CONCLUSIONS

Synchronization errors are a common source of defects in software systems. Our synchronization correction algorithm enables us to not only identify concurrency errors statically, but also correct them in many situations. Preliminary experiments demonstrate the effectiveness of our approach at correcting existing and randomly-inserted defects. We hope to further validate our approach on larger programs. This would also give us the opportunity to explore different optimization approaches in the final step of the algorithm. In addition, we plan to integrate a deadlock detection analysis into *Bohr* so that the tool can avoid introducing any potential for deadlock when inserting synchronization operations.

12. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grants CCR-0341179 and CCR-0341387.

13. REFERENCES

- [1] R. Agarwal and S. D. Stoller. Type inference for parameterized race-free Java. In *Proceedings of the Conference on Verification, Model Checking, and Abstract Interpretation*, pages 149–160, 2004.
- [2] C. Artho, K. Havelund, and A. Biere. High-level data races. In *The First International Workshop on Verification and Validation of Enterprise Information Systems*, 2003.
- [3] A. D. Birrell. An introduction to programming with threads. Research Report 35, Digital Equipment Corporation Systems Research Center, 1989.
- [4] C. Boyapati and 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*, pages 56–69, 2001.

- [5] D. Bruening. Systematic testing of multithreaded Java programs. Master's thesis, Massachusetts Institute of Technology, 1999.
- [6] L. Cardelli. Typechecking dependent types and subtypes. In *Lecture notes in computer science on Foundations of logic and functional programming*, pages 45–57, 1988.
- [7] X. Deng, M. Dwyer, J. Hatcliff, and M. Mizuno. Invariant-based specification, synthesis, and verification of synchronization in concurrent programs. In *International Conference on Software Engineering*, pages 442–452, 2002.
- [8] C. Flanagan. Verifying commit-atomicity using model-checking. In *International SPIN Workshop on Model Checking of Software*, 2004.
- [9] C. Flanagan and S. N. Freund. Type-based race detection for Java. In *Proceedings of the ACM Conference on Programming Language Design and Implementation*, pages 219–232, 2000.
- [10] C. Flanagan and S. N. Freund. Atomizer: A dynamic atomicity checker for multithreaded programs. In *Proceedings of the ACM Symposium on Principles of Programming Languages*, pages 256–267, 2004.
- [11] C. Flanagan and S. N. Freund. Type inference against races. In *Proceedings of the Static Analysis Symposium*, pages 116–132, 2004.
- [12] C. Flanagan, S. N. Freund, and M. Lifshin. Type inference for atomicity. In *Proceedings of the ACM Workshop on Types in Language Design and Implementation*, pages 47–58, 2005.
- [13] C. Flanagan and S. Qadeer. Transactions for software model checking. In *Proceedings of the Workshop on Software Model Checking*, 2003.
- [14] C. Flanagan and S. Qadeer. A type and effect system for atomicity. In *Proceedings of the ACM Conference on Programming Language Design and Implementation*, pages 338–349, 2003.
- [15] M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *Proceedings of the ACM Symposium on Principles of Programming Languages*, pages 171–183, 1998.
- [16] P. Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems: An Approach to the State-Explosion Problem*. Lecture Notes in Computer Science 1032. Springer-Verlag, 1996.
- [17] T. Harris, S. Marlow, S. Peyton-Jones, and M. Herlihy. Composable memory transactions. In *Proceedings of the ACM Symposium on Principles and Practice of Parallel Programming*, pages 48–60, 2005.
- [18] T. L. Harris and K. Fraser. Language support for lightweight transactions. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages and Applications*, 2003.
- [19] J. Hatcliff, Robby, and 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.
- [20] R. J. Lipton. Reduction: A method of proving properties of parallel programs. *Communications of the ACM*, 18(12):717–721, 1975.
- [21] M. F. Ringenbun and D. Grossman. AtomCaml: First-class atomicity via rollback. In *Proceedings of the ACM International Conference on Functional Programming*, 2005.
- [22] A. Sasturkar, R. Agarwal, L. Wang, and 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*, pages 83–94, 2005.
- [23] S. Savage, M. Burrows, G. Nelson, P. Sobalvarro, and T. E. Anderson. Eraser: A dynamic data race detector for multi-threaded programs. *ACM Transactions on Computer Systems*, 15(4):391–411, 1997.
- [24] N. Sterling. Warlock: A static data race analysis tool. In *USENIX Winter Technical Conference*, pages 97–106, 1993.
- [25] S. D. Stoller. Model-checking multi-threaded distributed Java programs. In *Workshop on Model Checking and Software Verification*, volume 1885 of *Lecture Notes in Computer Science*, pages 224–244. Springer-Verlag, 2000.
- [26] J.-P. Talpin and P. Jouvelot. Polymorphic type, region and effect inference. *Journal of Functional Programming*, 2(3):245–271, 1992.
- [27] M. Tofte and J.-P. Talpin. Implementation of the typed call-by-value lambda-calculus using a stack of regions. In *Proceedings of the ACM Symposium on Principles of Programming Languages*, pages 188–201, 1994.
- [28] C. von Praun and T. Gross. Object race detection. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages and Applications*, pages 70–82, 2001.
- [29] C. von Praun and T. Gross. Static conflict analysis for multi-threaded object-oriented programs. In *Proceedings of the ACM Conference on Programming Language Design and Implementation*, pages 115–128, 2003.
- [30] C. von Praun and T. Gross. Static detection of atomicity violations in object-oriented programs. *Journal of Object Technology*, 3(6):103–122, 2004.
- [31] L. Wang and S. D. Stoller. Runtime analysis of atomicity for multi-threaded programs. Technical Report DAR 04-14, Computer Science Department, SUNY Stony Brook, July 2004. A preliminary version appeared in *Proc. Workshop on Runtime Verification*, 2003.
- [32] A. Welc, S. Jagannathan, and A. L. Hosking. Transactional monitors for concurrent objects. In *Proceedings of the European Conference on Object-Oriented Programming*, pages 519–542, 2004.