Exploiting Purity for Atomicity (Extended Version)

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

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

Shaz Qadeer
Microsoft Research
One Microsoft Way
Redmond, WA 98052

Abstract

The notion that certain procedures are *atomic* is a fundamental correctness property of many multithreaded software systems. A procedure is atomic if for every execution there is an equivalent serial execution in which the actions performed by any thread while executing the atomic procedure are not interleaved with actions of other threads. Several existing tools verify atomicity by using commutativity of actions to show that every execution *reduces* to a corresponding serial execution. However, experiments with these tools have highlighted a number of interesting procedures that, while intuitively atomic, are not reducible.

In this paper, we exploit the notion of *pure* code blocks to verify the atomicity of such irreducible procedures. If a pure block terminates normally, then its evaluation does not change the program state, and hence these evaluation steps can be removed from the program trace before reduction. We develop a static analysis for atomicity based on this insight, and we illustrate this analysis on a number of interesting examples that could not be verified using earlier tools based purely on reduction. The techniques developed in this paper may also be applicable in other approaches for verifying atomicity, such as model checking and dynamic analysis.

This is an extended and revised version of our ISSTA '04 paper.

Categories and Subject Descriptors: D.2.4 [Software Engineering]: Software/Program Verification–*reliability*; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs.

General Terms: Languages, Verification, Reliability.

Keywords: Atomicity, purity, reduction, concurrent programs.

1 Introduction

Multiple threads of control are widely used in software development because they help reduce latency and provide better utilization of multiprocessor machines. However, reasoning about the correctness of multithreaded code is complicated by the nondeterministic interleaving of threads and the potential for unexpected interference between concurrent threads. Since exploring all possible interleavings of the executions of the various threads is clearly impractical, techniques for specifying and controlling the interference between concurrent threads are crucial for the development of reliable multithreaded software.

A canonical and widely-applicable non-interference guarantee is *atomicity*. A procedure (or code block) is atomic if for every (arbitrarily interleaved) program execution, there is an equivalent execution with the same overall behavior where the atomic procedure is executed serially, that is, the procedure's execution is not interleaved with actions of other threads. The notion of atomicity provides multiple benefits.

- The non-interference guarantee provided by atomicity reduces the challenging problem of reasoning about an atomic procedure's behavior in a *multithreaded* context to the simpler problem of reasoning about the procedure's *sequential* behavior. The latter problem is significantly more amenable to standard techniques such as manual code inspection, dynamic testing, and static analysis.
- Atomicity is a natural methodology for multithreaded programming, and experimental results indicate that many existing procedures and library interfaces already follow this methodology [12].
- Many synchronization errors can be detected as violations of atomicity.

Recently, a number of analyses have been developed for verifying atomicity, using techniques such as theorem proving [18], static type and effect systems [16, 17], dynamic analysis [12, 38], and model checking [22]. All of these approaches use *reduction* [26, 32], which is based on commuting operations in an execution performed by different threads when they do not interfere with each other to obtain an equivalent *serial execution* (where the operations of each atomic procedure are performed contiguously). An expression is *reducible* if it consists of zero or more *right movers* (steps that right-commute with steps of other threads), followed by at most one atomic step (that need not commute with steps of other threads), followed by zero or more *left movers* (steps that left-commute with steps of other threads).

To illustrate this notion of reduction, consider the procedure `setxy` shown below. In this procedure, the operation `acquire(m)` is a right mover, and the operation `release(m)` is a left mover. Moreover, if all threads access `x` and `y` only while holding the lock `m`, then the writes to `x` and `y` are both right-movers and left-movers since no other thread can concurrently access these variables. Thus, as illustrated below, we can reduce any execution of `setxy` interleaved with arbitrary steps ("Z") from other threads into an equivalent serial execution.

Procedure `setxy`

```
atomic void setxy() {
    acquire(m); x=1; y=2; release(m);
}
```

Reduction

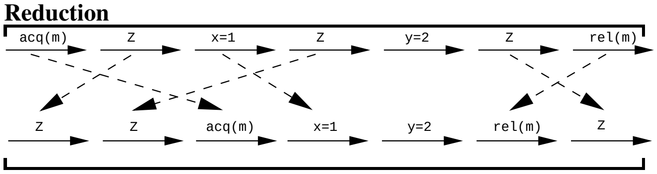

1.1 Purity

Reduction suffices to verify the atomicity of many procedures that use straightforward synchronization patterns, but it is often inadequate for procedures that use more subtle synchronization. A concrete example of this limitation is the procedure `busy_acquire` shown below, which uses a combination of busy-waiting and a compare-and-swap (`CAS`) operation to acquire a mutually-exclusive lock `m` (represented as a boolean).

Procedure `busy_acquire`

```
atomic void busy_acquire() {
     while (true) {
        if (CAS(m,0,1)) break;
     }
}
```

The operation `CAS(m,0,1)` has no effect and returns `false` if $m \neq 0$. However, if $m = 0$, then the operation `CAS(m,0,1)` sets `m` to `1` and returns `true`. This `CAS` operation does not commute with operations of concurrent threads, since it inspects and potentially updates the shared variable `m`. Hence, any execution of `busy_acquire` where the loop iterates multiple times cannot be *reduced* to a serial execution, and previous tools based purely on reduction cannot verify the atomicity of `busy_acquire`. In particular, our previous type and effect system for atomicity [17] cannot verify the atomicity of irreducible procedures like `busy_acquire`. The model checking approach described in [10] can verify the atomicity of `busy_acquire` but is limited by the state-explosion problem.

In this paper, we present a lightweight and scalable static analysis for verifying the atomicity of irreducible procedures such as `busy_acquire`. We present our analysis as an effect system (essentially, a collection of syntax-directed rules). This effect system is analogous to traditional type systems, except that it reasons about effects (which describe computations) as opposed to types (which describe values).

A key novelty of our analysis is the exploitation of *purity* when reasoning about atomicity. Essentially, a code block is pure if, whenever its evaluation terminates normally, it does not change the program state. This restriction does not apply when the block terminates *exceptionally*, for example, via a `break` or `return` statement. The body of the while loop in `busy_acquire` is pure, since if it updates `m` it immediately terminates exceptionally via the `break` statement. Otherwise, control is returned to the loop head without changing the program store. We introduce the `pure-while` statement to indicate a pure loop and rewrite the `busy_acquire` procedure as follows:

Procedure `busy_acquire` **with a** `pure` **loop**

```
atomic void busy_acquire() {
     pure-while (true) {
        if (CAS(m,0,1)) break;
     }
}
```

The intuition behind the reasoning of our analysis for `busy_acquire` is shown in the following figure. The figure shows an execution of `busy_acquire` consisting of three normally-terminating loop iterations in which the `CAS` fails, followed by an exceptionally-terminating iteration in which the `CAS` operation succeeds.

Execution trace for `busy_acquire`

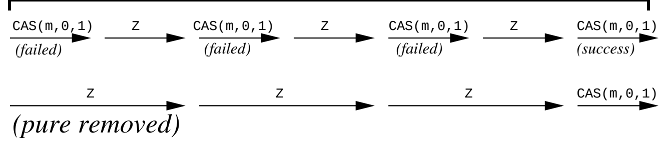

Since the normally-terminating iterations do not change the program state, our verification technique essentially removes them from the execution trace to yielding a trace containing a single loop iteration in which the `CAS` operation succeeds. Since every execution of `busy_acquire` is serializable in this manner, our analysis can conclude that `busy_acquire`, although irreducible, is atomic.

1.2 Abstraction via purity

A more interesting example of our analysis technique is the following procedure `alloc`, which searches for a free disk block. The flag `free[i]` indicates whether the `i`-th disk block is currently unused, and this flag is protected by the mutually-exclusive lock `l[i]`. When `alloc` identifies a free block, it allocates the block by setting the appropriate bit to false and returns the index of that block. The `alloc` procedure returns -1 if it fails to find a free block.

Procedure `alloc`

```
atomic int alloc() {
    int i = 0;
    int r = -1;
    while (i < max) {
        pure {
            acquire(l[i]);
            if (free[i]) {
                free[i] = false;
                release(l[i]);
                r = i;
                break;
            }
            release(l[i]);
        }
        i++;
    }
    return r;
}
```

Execution trace for `alloc`

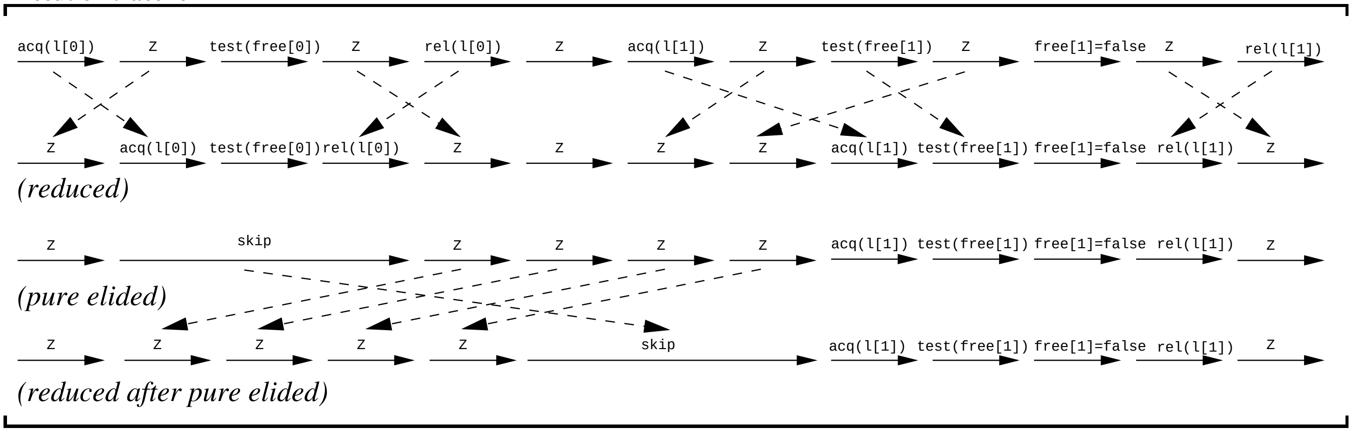

This procedure is not actually serializable, since there exist some (non-serial) executions of this procedure that are not equivalent to any serial executions. In particular, a concurrent thread could ensure that there is always at least one free block at any point in time, yet the sequential search performed by `alloc` could still fail to find a free block. Thus the concrete implementation of `alloc` is not atomic, and this lack of atomicity significantly complicates reasoning about the behavior or correctness of `alloc`. However, `alloc` *is* atomic in an abstract sense because any execution performs the atomic action of either allocating a block or returning -1 if no free block was found.

In order to facilitate sequential reasoning for non-atomic procedures such as `alloc`, this paper introduces the notion of *abstract atomicity* and shows that the procedure `alloc` is atomic under a more permissive or *abstract* semantics. In this abstract semantics, the execution of the `pure` block in `alloc` is *optional*, and may or may not be executed on any loop iteration. Thus, the abstract semantics introduces additional nondeterminism and admits additional execution sequences for `alloc`. Despite this additional nondeterminism, every serial execution of `alloc` under the abstract semantics satisfies the correctness specification of returning either -1 or the index of a free block. This correctness property can be verified using sequential reasoning techniques.

Given that every serial execution of `alloc` is correct, the contribution of our analysis is to verify that every possible *interleaved* execution of `alloc` yields the same behavior as some serial execution, thus allowing us to conclude that all interleaved executions of `alloc` are also correct. Thus, our analysis enables reasoning about behavior and correctness of abstractly atomic (though not atomic) procedures such as `alloc` using sequential reasoning. In contrast, previous techniques only achieved this goal (of enabling sequential reasoning) for procedures that are both atomic and reducible.

The central intuition behind the reasoning performed by our analysis to verify abstract atomicity is illustrated graphically in the execution traces for `alloc` shown above. The first trace contains an execution of `alloc` that succeeds on the second loop iteration, interleaved with arbitrary actions "Z" of concurrent threads. We only show steps of `alloc` that modify shared variables. By reduction, we can prove that each individual loop iteration is reducible. Since the first execution of the `pure` block is normally-terminating and hence effect-free, we replace it with `skip` in the execution sequence, at which point applying reduction a second time yields an equivalent serial execution.

The Calvin-R checker [18] can verify similar atomicity properties. However, that tool focuses on checking more complete functional specifications of concurrent programs and has a higher annotation overhead and analysis complexity than the technique in this paper.

1.3 Abstraction via instability

Our analysis also supports *unstable* variables, such as performance counters, which do not affect program correctness. These variables are typically not protected by locks and have race conditions on them. Consequently, accesses to these variables do not commute. Our analysis verifies atomicity with respect to an abstract semantics in which every write access to an unstable variable writes a nondeterministic value and every read access reads a nondeterministic value. Under such an abstract semantics, read and write accesses to unstable variables both right and left commute. For example, a program may use an unstable `packetCount` variable to record the number of packets received for tracking performance. Operations on that variable do not affect the atomicity of the code in which they appear. We present a complete example in Section 4.4.

Outline. The presentation of our results proceeds as follows. The following section introduces an idealized language that we use for studying atomicity. Section 3 presents the effect system for atomicity, and Section 4 illustrates this analysis on a number of example programs. Section 5 sketches an extension of our technique based on a more flexible notion of purity. We discuss related work in Section 6 and conclude with a discussion of future directions in Section 7.

2 The language CAP

We formalize our ideas in terms of CAP, a small, imperative, multithreaded language with higher-order functions and dynamic thread creation. In essence, CAP is a restricted subset of C, extended with facilities for reasoning about atomicity and purity.

CAP expressions include values, variable reference and assignment, primitive and function applications, conditionals, and `let`-expressions. The `fork` e expression creates a new thread for the evaluation of e. Values are constants and function definitions. Constants must include integer constants but are otherwise unspecified. The definition $f(\overline{x})\ e$ introduces a function named f. The formal parameters $\overline{x}$ are bound within the body e, and they may be α-renamed in the usual fashion. For generality, we leave the set of primitives unspecified, but they might include, for example, synchronization primitives such as lock creation, acquire, and release operations for mutual exclusion locks. We assume the set of primitives also include arithmetic operations and `assert`.

In addition to terminating normally and yielding a resulting value, the evaluation of a CAP expression can also terminate *exceptionally* via the `break` construct, which transfers control from the current expression to the end of the closest dynamically-enclosing `block` construct. The construct `loop` e repeatedly evaluates e until e `break`'s to an enclosing `block`.

To facilitate our atomicity analysis, expressions can be annotated with the keyword `pure`. The evaluation of an expression `pure` e is optional, and may be skipped. The keyword `pure` states that, when the expression e is evaluated and terminates normally, that evaluation does not change the program state. (Only exceptionally terminating evaluations of a pure expression are allowed to change the program state.) If a `pure` expression temporarily changes the program state, for example, by acquiring a lock, then it must restore the state by releasing the lock before terminating normally. The language CAP supports unstable variables, so the set of variable names is divided into stable and unstable variables. By convention, unstable variable names begin with "_".

Syntax

$$
\begin{array}{lclcl}
e & \in & \mathit{Expr} & ::= & v \mid x_r \mid x_r := e \mid p(\overline{e}) \mid e^F(\overline{e}) \\
 & & & & \mid \texttt{if}\ e\ e\ e \mid \texttt{loop}\ e \mid \texttt{block}\ e \mid \texttt{break} \\
 & & & & \mid \texttt{let}\ x = e\ \texttt{in}\ e \mid \texttt{fork}\ e \\
 & & & & \mid \texttt{atomic}\ e \mid \texttt{pure}\ e \\
 & & & & \\
v & \in & \mathit{Value} & ::= & c \mid f(\overline{x})\ e \\
r & \in & \mathit{Tag} & ::= & \bullet \mid \epsilon \\
 & & & & \\
x & \in & \mathit{Var} & = & \mathit{StableVar} \uplus \mathit{UnstableVar} \\
f & \in & \mathit{FnName} & & \\
F & \in & 2^{\mathit{FnName}} & & \\
p & \in & \mathit{Prim} & & \\
c & \in & \mathit{Const} & &
\end{array}
$$

We introduce syntactic sugar for some common constructs.

$$
\begin{array}{rcll}
e_1; e_2 & \equiv & \texttt{let}\ x = e_1\ \texttt{in}\ e_2 & \text{for } x \text{ not free in } e_2 \\
\texttt{while}\ e_1\ e_2 & \equiv & \texttt{block loop}\ \{\ \texttt{if}\ e_1\ e_2\ \texttt{break}\ \} & \\
\texttt{pure-while}\ e_1\ e_2 & \equiv & \texttt{block loop pure}\ \{\ \texttt{if}\ e_1\ e_2\ \texttt{break}\ \} &
\end{array}
$$

Note that if e_1 and e_2 are pure, then `while` e_1 e_2 and `pure-while` e_1 e_2 are semantically equivalent (that is, replacing `pure-while` e_1 e_2 in a program with `while` e_1 e_2 does not change the observable behaviors of the program).

To simplify our presentation, the CAP effect system does not reason about race conditions, control flow, or purity, since these topics can be addressed by other analyses. Instead, we assume the program has already been annotated and checked by alternative analyses as follows:

1. Each variable access (read or write) has a *conflict tag*, which is $\bullet$ if that access may be involved in a race condition on a stable variable, and is ϵ otherwise. Thus, all accesses to unstable variables or correctly synchronized stable variables will have conflict tag ϵ. Existing analysis techniques [4, 11, 20, 36, 14] can be used to infer these conflict tags.
2. Each function call $e^F(\overline{e})$ has a *call tag* F denoting the set of functions that may be invoked by that call. These call tags can be computed by a standard flow analysis.
3. Each `pure` e expression is side-effect free when e evaluates normally. We present an effect system to check purity in the Appendix. Nielson, Nielson, and Hankin [30] provide a general overview of other effect-based techniques for tracking side effects, and these may be extended for our purposes as well.

We also assume programs being checked have passed a conventional type checker to catch basic type errors, such as performing an arithmetic operation on non-numeric arguments. Factoring these other issues enables us to focus on the key aspects of this work without the added complexity of these other analyses. The core focus of our analysis is on verifying that every expression or procedure that is annotated as `atomic` is, in fact, serializable.

3 Effect system

We formalize our static analysis for abstract atomicity as an effect system. Previous type and effect systems [17, 16] could only verify the atomicity of procedures that are reducible. By introducing optionally-executed pure blocks and unstable variables, our effect system can also verify many interesting irreducible procedures, such as those in Sections 1 and 4, are still atomic. (This abstract semantics is defined in Section B.)

Each expression in our language can terminate either normally (by evaluating to a value) or exceptionally (via `break`). For each termination mode, our effect system assigns to each expression an *atomicity* from the following set:

$$a, b, c \in \mathit{Atomicity} = \{R, L, B, \bot, A, \top\}$$

This atomicity identifies whether the evaluation of the expression

- right-commutes with operations of other threads (R);
- left-commutes with operations of other threads (L);
- both right- and left-commutes (B);
- cannot terminate in that mode ($\bot$);
- can be viewed as a single atomic action (A); or
- exhibits none of these properties ($\top$).

Atomicities are partially ordered by the relation $\sqsubseteq$, as follows:

$$
\begin{tikzcd}[column sep=small]
 & \top \arrow[d, no head] & \\
 & A \arrow[dl, no head] \arrow[dr, no head] & \\
L \arrow[dr, no head] & & R \arrow[dl, no head] \\
 & B \arrow[d, no head] & \\
 & \bot &
\end{tikzcd}
$$

Let $\sqcup$ denote the join operator based on this ordering. If atomicities a_1 and a_2 reflect the normal-termination behavior of expressions e_1 and e_2 respectively, then the *sequential composition* $a_1; a_2$ reflects the normal-termination behavior of $e_1; e_2$, and is defined by the following table.

Effect System

$$\boxed{\Gamma \vdash e : a \uparrow b}$$

[EXP CONST]
$$\frac{}{\Gamma \vdash c : B \uparrow \bot}$$

[EXP FUN]
$$\frac{\Gamma(f) = \langle a, b\rangle \quad \Gamma \vdash e : a \uparrow b}{\Gamma \vdash f(\overline{x})\, e : B \uparrow \bot}$$

[EXP PRIM]
$$\frac{\Gamma \vdash \overline{e} : a \uparrow b}{\Gamma \vdash p(\overline{e}) : (a; \Gamma(p)) \uparrow b}$$

[EXP READ]
$$\frac{}{\Gamma \vdash x_\epsilon : B \uparrow \bot}$$

[EXP READ RACE]
$$\frac{}{\Gamma \vdash x_\bullet : A \uparrow \bot}$$

[EXP ASSIGN]
$$\frac{\Gamma \vdash e : a \uparrow b}{\Gamma \vdash x_\epsilon := e : (a; B) \uparrow b}$$

[EXP ASSIGN RACE]
$$\frac{\Gamma \vdash e : a \uparrow b}{\Gamma \vdash x_\bullet := e : (a; A) \uparrow b}$$

[EXP LET]
$$\frac{\Gamma \vdash e_1 : a_1 \uparrow b_1 \quad \Gamma \vdash e_2 : a_2 \uparrow b_2}{\Gamma \vdash \texttt{let}\ x = e_1\ \texttt{in}\ e_2 : (a_1; a_2) \uparrow (b_1 \sqcup (a_1; b_2))}$$

[EXP IF]
$$\frac{\Gamma \vdash e : a \uparrow b \qquad \Gamma \vdash e_i : a_i \uparrow b_i}{\Gamma \vdash \texttt{if}\ e\ e_1\ e_2 : (a; (a_1 \sqcup a_2)) \uparrow (b \sqcup (a; (b_1 \sqcup b_2)))}$$

[EXP LOOP]
$$\frac{\Gamma \vdash e : a \uparrow b}{\Gamma \vdash \texttt{loop}\ e : \bot \uparrow (a^*; b)}$$

[EXP FORK]
$$\frac{\Gamma \vdash e : a \uparrow b}{\Gamma \vdash \texttt{fork}\ e : A \uparrow \bot}$$

[EXP INVOKE]
$$\frac{\begin{array}{c}\Gamma \vdash e : a \uparrow b \qquad \Gamma \vdash \bar{e} : a' \uparrow b' \\ a'' = \sqcup\{a \mid f \in F \wedge \Gamma(f) = \langle a, b\rangle\} \\ b'' = \sqcup\{b \mid f \in F \wedge \Gamma(f) = \langle a, b\rangle\}\end{array}}{\Gamma \vdash e^F(\overline{e}) : (a; a'; a'') \uparrow (b \sqcup (a; b') \sqcup (a; a'; b''))}$$

[EXP BLOCK]
$$\frac{\Gamma \vdash e : a \uparrow b}{\Gamma \vdash \texttt{block}\ e : (a \sqcup b) \uparrow \bot}$$

[EXP BREAK]
$$\frac{}{\Gamma \vdash \texttt{break} : \bot \uparrow B}$$

[EXP ATOMIC]
$$\frac{\Gamma \vdash e : a \uparrow b \qquad a, b \sqsubseteq A}{\Gamma \vdash \texttt{atomic}\ e : a \uparrow b}$$

[EXP PURE]
$$\frac{\Gamma \vdash e : a \uparrow b \qquad a \sqsubseteq A}{\Gamma \vdash \texttt{pure}\ e : B \uparrow b}$$

$$\boxed{\Gamma \vdash \bar{e} : a \uparrow b}$$

[EXP EMPTY SEQ]
$$\frac{}{\Gamma \vdash \epsilon : B \uparrow \bot}$$

[EXP SEQ]
$$\frac{\Gamma \vdash \bar{e} : a \uparrow b \qquad \Gamma \vdash e : a' \uparrow b'}{\Gamma \vdash (\bar{e}.e) : (a; a') \uparrow (b \sqcup (a; b'))}$$

;	$\bot$	B	L	R	A	$\top$
$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
B	$\bot$	B	L	R	A	$\top$
L	$\bot$	L	L	$\top$	$\top$	$\top$
R	$\bot$	R	A	R	A	$\top$
A	$\bot$	A	A	$\top$	$\top$	$\top$
$\top$	$\bot$	$\top$	$\top$	$\top$	$\top$	$\top$

Similarly, if atomicity a reflects the normal-termination behavior of e, then the *iterative closure* a^* reflects the normal-termination behavior of executing e zero or more times, and is defined by

$$\begin{array}{rcl}\bot^* & = & B \\ A^* & = & \top \\ a^* & = & a \text{ for } a \in \{B, L, R, \top\}\end{array}$$

Note that

1. sequential composition is associative and B is the left and right identity of this operation,
2. iterative closure is idempotent, and
3. sequential composition distributes over joins.

An effect environment Γ maps each function name to a pair of atomicities $\langle a, b\rangle$ that describe the function's behavior under normal and exceptional termination. In addition, Γ also maps each primitive operation to a corresponding atomicity (note that primitives never terminate exceptionally):

$$\begin{array}{rl}\Gamma & : \quad (\mathit{FnName} \rightarrow \mathit{Atomicity} \times \mathit{Atomicity}) \\ & \cup \quad (\mathit{Prim} \rightarrow \mathit{Atomicity})\end{array}$$

The atomicity of some common primitives are:

$$\begin{array}{rcl@{\qquad}rcl}\Gamma(\texttt{assert}) & = & B & \Gamma(\texttt{new\_lock}) & = & B \\ \Gamma(\texttt{CAS}) & = & A & \Gamma(\texttt{acquire}) & = & R \\ \Gamma(\texttt{+}) & = & B & \Gamma(\texttt{release}) & = & L\end{array}$$

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

$$\Gamma \vdash e : a \uparrow b$$

This judgment states that the expression e has atomicity a under normal termination, and atomicity b under exceptional termination. The rules defining these judgments are mostly straightforward. For example, the "evaluation" of a constant terminates normally, does not interfere with other threads, and cannot terminate exceptionally.

[EXP CONST]
$$\frac{}{\Gamma \vdash c : B \uparrow \bot}$$

The rule [EXP LET] states that the normal atomicity of a let expression $\texttt{let}\ x = e_1\ \texttt{in}\ e_2$ is the sequential composition $a_1; a_2$ of the normal atomicities of e_1 and e_2. The exceptional atomicity of a let expression reflects the places where the let expression could `break`.

[EXP LET]
$$\frac{\Gamma \vdash e_1 : a_1 \uparrow b_1 \qquad \Gamma \vdash e_2 : a_2 \uparrow b_2}{\Gamma \vdash \texttt{let}\ x = e_1\ \texttt{in}\ e_2 : (a_1; a_2) \uparrow (b_1 \sqcup (a_1; b_2))}$$

Another example of how the exceptional atomicity reflects all the ways in which an expression may `break` is the [EXP IF] rule.

[EXP IF]
$$\frac{\Gamma \vdash e : a \uparrow b \qquad \Gamma \vdash e_i : a_i \uparrow b_i}{\Gamma \vdash \texttt{if}\ e\ e_1\ e_2 : (a; (a_1 \sqcup a_2)) \uparrow (b \sqcup (a; (b_1 \sqcup b_2)))}$$

The rule [EXP LOOP] states that the normal atomicity of the loop is $\bot$, since it never terminates normally. The exceptional atomicity for a loop reflects the fact that the loop body could terminate normally many times before terminating exceptionally.

[EXP LOOP]
$$\frac{\Gamma \vdash e : a \uparrow b}{\Gamma \vdash \texttt{loop}\ e : \bot \uparrow (a^*; b)}$$

The atomicity of a variable read x_r depends on the conflict tag r. If $r = \epsilon$, then this read commutes with steps of other threads, and so has normal atomicity B. If $r = \bullet$, then this read has normal atomicity A, indicating that it is an atomic action that may not commute with steps of other threads. The rules for variable writes are similar.

$$[\text{EXP READ}] \quad \frac{}{\Gamma \vdash x_\epsilon : B \uparrow \bot} \qquad [\text{EXP READ RACE}] \quad \frac{}{\Gamma \vdash x_\bullet : A \uparrow \bot}$$

A `block` e expression never terminates exceptionally. Either the body e terminates normally, or it executes a `break` expression that terminates e early. In the latter case, we still consider `block` e to exit normally. A `break` expression only terminates exceptionally and is a both mover.

$$[\text{EXP BLOCK}] \quad \frac{\Gamma \vdash e : a \uparrow b}{\Gamma \vdash \texttt{block}\ e : (a \sqcup b) \uparrow \bot} \qquad [\text{EXP BREAK}] \quad \frac{}{\Gamma \vdash \texttt{break} : \bot \uparrow B}$$

A key innovation of our effect system is our treatment of pure blocks. The rule [EXP PURE] for `pure` e states that the normal atomicity of the body of a pure block must be at most A. This requirement ensures that any side effects during the evaluation of e are not visible to other threads. Since, under normal termination, the pure block has no observable effect, our effect system "optimizes" the normal atomicity of pure block to a both-mover B.

$$[\text{EXP PURE}] \quad \frac{\Gamma \vdash e : a \uparrow b \qquad a \sqsubseteq A}{\Gamma \vdash \texttt{pure}\ e : B \uparrow b}$$

Finally, the normal and exceptional atomicities of the body of an atomic construct are required to be at most A.

$$[\text{EXP ATOMIC}] \quad \frac{\Gamma \vdash e : a \uparrow b \qquad a, b \sqsubseteq A}{\Gamma \vdash \texttt{atomic}\ e : a \uparrow b}$$

Our effect system is sound in the sense that any execution trace of a well-typed program is equivalent to a serial execution of that program. In this serial execution, the steps of each `atomic` block are executed sequentially, without steps interleaved from other threads. To verify this serializability property, we first reduce each normally-terminating `pure` block into a sequence of contiguous steps and we replace that sequence with a single `skip` step. We then reduce `atomic` blocks in the modified execution to obtain an equivalent serial execution. We refer the interested reader to an extended version of this paper for a full proof of this result [13].

4 Applications

In this section, we present several examples to illustrate the expressiveness of our effect system for atomicity. In these examples, we sometimes enclose expressions in parentheses or braces and use additional constructs such as `return` for clarity. We start by considering the double-checked initialization pattern, commonly used to ensure that a shared variable is initialized exactly once [35].[1]

[1] Our analysis assumes a sequentially consistent memory model. Double-checked initialization may not work correctly under other memory models.

4.1 Double-checked initialization

To avoid excessive synchronization overhead, the variable `x` below is initially tested without holding its protecting lock `l`. If the first test fails, the lock is acquired, and if `x` is still `null`, then it is initialized. Note that the read $\texttt{x}_\epsilon$ is not a conflicting access, since it commutes with concurrent reads, but the write $\texttt{x}_\bullet$ may conflict with reads of other threads. Since the procedure consists of an atomic operation (the first read of `x`) followed by a right-mover operation (`acquire(l)`), the procedure is not reducible and cannot be verified as atomic using previous analyses.

Double-Checked Locking

```
atomic void init() {
    block {
        pure { if (x• != null) break; }
        acquire(l);
        if (xϵ == null) x• = new();
        release(l);
    }
}
```

Our approach exploits the fact that the first test of `x` is both pure and optional; omitting this test does not affect the correctness of the program, only its performance, and thus we can enclose this test in a `pure` construct. If the first test succeeds, the procedure returns via a reducible trace. If the first test fails, then that test has no effect on the program store and we replace it by skip in the trace (just as for `alloc`), yielding a reducible trace through the function `init`. By this reasoning, our effect system verifies each possible execution of `init` has an equivalent serial execution, and hence `init` is atomic.

4.2 Caching

In the next example, the function `compute` constructs the value for a given key but is an expensive operation, so we wish to cache previously-computed results. We assume the cache operations `cachePut` and `cacheGet` are atomic (for example, because they acquire the lock protecting the cache); `cacheGet` is a pure (side-effect free) function; and that `compute` is a both-mover. We would like to verify that `lookup` is atomic, to ensure that it still behaves correctly even if when concurrently invoked by multiple threads.

Caching

```
atomic void cachePut(String k, Object val) { ...}
atomic pure Object cacheGet(String k) { ... }

// expensive operation
both-mover Object compute(String k) { ... }

atomic Object lookup(String k) {
    pure {
        Object r = cacheGet(kϵ);
        if (rϵ != null) return rϵ;
    }
    Object r = compute(kϵ);
    cachePut(kϵ,rϵ);
    return rϵ;
}
```

The function `lookup` is irreducible, since it contains sequentially composed atomic operations, `cacheGet` and `cachePut`. Note that the alternative implementation of holding the cache lock throughout `lookup` would introduce undesirable contention, since `compute` is a long-running operation. However, the cache lookup is clearly an optimization and can be omitted without affecting program correctness. We exploit this fact by enclosing the cache lookup in a `pure` construct. If the cache lookup is successful, the function `lookup` immediately returns via a reducible trace. If the cache lookup fails, it has no effect on the program store. Our analysis leverages this information (documented by the `pure` keyword) to essentially "remove" the cache lookup from the trace by replacing it with skip and to produce an equivalent, reducible trace. Thus, all executions through the function `lookup` have an equivalent serial execution, and so the function `lookup` is atomic.

4.3 Wait and notify

The `wait` and `notify` routines facilitate notification between concurrent threads. The routine `wait(l)` should only be called if the lock `l` is held; this routine then releases l, blocks until a concurrent thread calls `notify(l)`, and then returns after re-acquiring l. Typically, the routine `wait(l)` is called inside a loop that iterates until a desired condition holds, and concurrent threads call `notify(l)` whenever a state change may affect the desired condition. We model `wait(l)` and `notify(l)` as {`release(l);acquire(l)`} and `skip`, respectively. This model captures the essence that other threads may acquire `l` during the execution of `wait(l)`. In other words, `wait` is not atomic.

The following code fragment illustrates the use of `wait` to iterate until the variable `x` (protected by lock `l`) is `false`, and we assume that *body* is atomic.

Wait example

```
acquire(l);
while x {
    wait(l);
}
body;
release(l);
```

For this example, even though `wait(l)` is not atomic, our type system can verify that the entire code fragment, although irreducible, is still atomic. Before applying our type system, we first need to refactor this code using the following equivalence rules for program expressions:

Equivalence rules

$$\begin{array}{rcll} e; \texttt{block}\ e' & = & \texttt{block}\ e; e' & \text{if } e \text{ cannot } \texttt{break} \\ e; \texttt{loop}\ \{e'; e\} & = & \texttt{loop}\ \{e; e'\} & \\ \texttt{break} & = & \texttt{break}; e & \\ \texttt{if}\ e_1\ \{e_2; e\}\ \{e_3; e\} & = & \{\texttt{if}\ e_1\ e_2\ e_3\}; e & \end{array}$$

Applying these rules to the above code fragment in the appropriate manner yields the following refactored code that has equivalent behavior, but where the body of the loop is now pure. (Note that not all uses of `wait` can be refactored in this manner.)

Refactored wait example

```
block loop pure {
    acquire(l);
    if x release(l) break;
}
body;
release(l);
```

The purity of the refactored loop allows our effect system to verify that each loop iteration except the last has no side-effect and can be elided from the execution sequence. The resulting execution sequence acquires the lock, checks that `x` is `false`, executes *body*, and releases the lock. This sequence is both atomic and reducible. Since every possible execution of the original code fragment is equivalent to such an atomic execution, the original code fragment is atomic.

4.4 Packet counter

The following example counts the number of packets received in a program with the `_packetCount` variable, which is used only for monitoring or performance purposes. To avoid synchronization overhead, the program accesses `_packetCount` without synchronization, with the expectation that the resulting race conditions will not cause the resulting count to be substantially incorrect. By marking `_packetCount` as unstable, we can still consider procedures like `receive` to be atomic, despite the presence of race conditions. (We do need to check the sequential correctness of `receive` under the abstract semantics where `_packetCount` may change nondeterministically.)

Packet counter

```
int _packetCount;
Queue packets;

atomic void enqueue(Queue q, Packet p) { ... }

atomic void receive(Packet p) {
    _packetCount_ε++;
    enqueue(packets_ε, p_ε);
}
```

5 Abstraction via weak purity

The technique of optionally-executed pure blocks is sufficient to handle many examples, such as those described in Section 4. In this section, we sketch a more general notion of purity that yields a more expressive effect system for atomicity.

Consider the following function which models optimistic concurrency control based on transaction retry. We have a shared data variable `z` that we wish to update according to `z = f(z)`. However, the function `f` is a long-running operation, so we do not wish to hold `z`'s protecting lock `m` when computing `f`. Instead, we record a local copy `x` of `z`, compute `f(x)`, and then update `z` if the value of `z` has not changed. If `z` has changed, then we retry the transaction. This technique ensures that the update of `z` to `f(z)` is serialized with respect to other updates to `z`, without requiring the lock guarding `z` to be held while computing `f(z)`.

Transaction retry

```
atomic void apply_f() {
    int x, fx;
    weak-pure {
        acquire(m);
        xε = zε;
        release(m);
    }
    weak-pure-while (true) {
        fxε = f(xε);
        acquire(m);
        if (xε == zε) {
            zε = fxε;
            release(m);
            break;
        }
        xε = zε;
        release(m);
    }
}
```

The code block before the loop is not pure because it modifies the local variable `x`. The body of the while loop is also not pure because it modifies the local variables `x` and `fx`. To deal with this prototypical example, we introduce a weaker notion of purity that allows us to prove the atomicity of `apply_f`. We first classify variables as either *thread-local* or *shared*, depending on whether the variable may be accessed by one or multiple threads, respectively. A code block can be annotated as `weak-pure` if it is atomic and does not modify any shared variables under normal termination. Its evaluation may modify thread-local variables, making `weak-pure` less restrictive than `pure` (which may not modify either thread-local or shared variables under normal termination).

The construct `weak-pure-while` e_1 e_2 is desugared in a fashion similar to `pure-while`:

$$\texttt{weak-pure-while}\ e_1\ e_2 \equiv \\ \qquad \texttt{block loop weak-pure}\ \{\ \texttt{if}\ e_1\ e_2\ \texttt{break}\ \}$$

and is semantically equivalent to `while` e_1 e_2 (provided $\{$`if` e_1 e_2 `break`$\}$ is weakly-pure).

The abstract semantics executes `weak-pure` e as normal, except that if e accesses a shared variable x, then an arbitrary value is returned. This arbitrary value is consistent for all accesses to x during a single execution of e. Thus the abstract semantics for `weak-pure` introduces additional execution traces and we need to check (formally or informally) the sequential correctness of `apply_f` under this abstract semantics. The choice of abstract semantics ensures that `weak-pure` blocks neither read nor modify values of shared variables, and enables us to treat atomic `weak-pure` blocks as both-movers.

The typing rule for `weak-pure` is identical to the rule used to reason about pure blocks. The rule requires that the normal atomicity of the body of a weak-pure block must be at most A and "optimizes" the normal atomicity of the block to a both-mover B.

[EXP WEAKPURE]

$$\frac{\Gamma \vdash e : a \uparrow b \qquad a \sqsubseteq A}{\Gamma \vdash \texttt{weak-pure}\ e : B \uparrow b}$$

6 Related work

Lipton [26] first proposed reduction as a way to reason about deadlocks in concurrent programs without considering all possible interleavings. Reduction has subsequently been extended to support proofs of general safety and liveness properties [8, 3, 25, 6, 29]. Bruening [5] and Stoller [37] have used reduction to improve the efficiency of model checking. Flanagan and Qadeer have pursued a similar approach [15], and Qadeer *et al* [33] have used reduction to infer procedure summaries in concurrent programs.

We previously applied reduction to verify atomicity in a static type and effect system for Java programs [17, 16]. This paper improves on that approach by enabling us to reason about the atomicity of code that is not immediately reducible.

The Calvin-R checker for multithreaded code relates procedure implementations to their functional specifications with an abstraction relation based on both reduction and simulation [18]. While capable of checking the atomicity of the examples in this paper, the overhead of that approach, in terms of annotation size and analysis complexity, is much greater. In contrast, the approach presented in this paper is more scalable, intuitive, and easier to use for checking atomicity properties.

Wang and Stoller [38] have developed a dynamic algorithm that can verify the atomicity of some irreducible code sequences. Their approach constructs the feasible interleavings of steps from two blocks of code and then determines whether all such interleavings are serializable. Unlike our approach, that algorithm does not require abstraction or auxiliary analysis to recognize pure blocks, and it is in some sense a complementary approach to ours.

The Atomizer is another dynamic analysis tool for detecting atomicity violations [12]. Our experience with the Atomizer, which uses reduction, suggests that the techniques developed in this paper could eliminate a nontrivial number of spurious warnings in reduction-based atomicity checkers.

The use of model checking for verifying atomicity is being explored by Hatcliff *et al* [22], and they present two approaches, based on Lipton's theory of reduction and partial-order reductions [19], respectively. Model checking offers several advantages over our effect system. For example, it requires many fewer programmer-inserted annotations and can accommodate complex synchronization disciplines more easily. Their experimental results suggest that verifying atomicity via model-checking is feasible for unit-testing. Their approach currently only verifies the atomicity of reducible procedures, but we believe that integrating our notions of abstraction and atomicity into their system could yield many of the benefits of both approaches.

In related work, Robby *et al* [34] demonstrate how to refactor code in order to extract some reducible code blocks embedded inside irreducible functions. This technique could, for example, refactor `alloc` to utilize an auxiliary (and reducible) method that contains a variant of the code inside the body of the `for` loop. In this way, one could check the atomicity of the auxiliary method, and possibly specify its behavior with standard pre- and post-conditions. However, the entire `alloc` function could not be shown to be atomic in an abstract sense, without performing an analysis like the one outlined in this paper.

A number of tools have been developed for detecting race conditions, both statically and dynamically. The Race Condition

Checker [11] uses a type system to catch race conditions in Java programs. This approach has been extended [4] and adapted to other languages [20]. Other static race detection tools include Warlock [36], for ANSI C programs, and ESC/Java [14], which catches a variety of software defects in addition to race conditions.

Atomicity is a semantic correctness condition for multithreaded software. It is related to strict serializability [31], a correctness condition for database transactions, and linearizability [23], a correctness condition for concurrent objects. It is possible that techniques for verifying atomicity can be leveraged to develop lightweight checking tools for related correctness conditions.

Other languages have included a notion of atomicity as a primitive operation. Hoare [24] and Lomet [28] first proposed the use of atomic blocks for synchronization, and the Argus [27] and Avalon [9] projects developed language support for implementing atomic objects. Persistent languages [1, 2] augment atomicity with data persistence in order to introduce transactions into programming languages. Other recent approaches to supporting atomicity include lightweight transactions [21, 39] and automatic generation of synchronization code from high-level specifications [7].

7 Conclusion

Atomicity is an important correctness property for multithreaded software. Current reduction-based tools can verify atomicity requirements in common cases, but they cannot handle situations in which code that is intuitively atomic is not immediately reducible. A number of frequently used programming idioms fall into this category.

This paper describes a static analysis technique capable of verifying the atomicity of many such problematic cases, by applying reduction to an abstraction of the program. The abstraction notions we have presented —based on optional execution, purity, and instability— are intuitive, and the correctness of abstractly atomic procedures under the serial abstract semantics can be verified using sequential reasoning techniques. Our static analysis then verifies that all interleaved executions of these abstractly atomic procedures are also correct.

Although we present our analysis as an effect system, these concepts may be applicable in other domains. For example, software model checkers (such as [22]) could identify and exploit pure code blocks while performing reduction. Dynamic analyses for atomicity [12] could potentially benefit from these ideas as well.

Acknowledgments

We thank Leslie Lamport and Chandu Thekkath for useful discussions. Cormac Flanagan was partly supported by the National Science Foundation under Grant CCR-0341179 and by faculty research funds granted by the University of California at Santa Cruz. Stephen Freund was partly supported by the National Science Foundation under Grant CCR-0341387 and by faculty research funds granted by Williams College.

8 References

[1] M. P. Atkinson, K. J. Chisholm, and W. P. Cockshott. PS-Algol: an Algol with a persistent heap. *ACM SIGPLAN Notices*, 17(7):24–31, 1981.

[2] M. P. Atkinson and D. Morrison. Procedures as persistent data objects. *ACM Transactions on Programming Languages and Systems*, 7(4):539–559, 1985.

[3] R.-J. Back. A method for refining atomicity in parallel algorithms. In *PARLE 89: Parallel Architectures and Languages Europe*, volume 366 of *Lecture Notes in Computer Science*, pages 199–216. Springer-Verlag, 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] E. Cohen and L. Lamport. Reduction in TLA. In *Proceedings of the International Conference on Concurrency Theory*, volume 1466 of *Lecture Notes in Computer Science*, pages 317–331. Springer-Verlag, 1998.

[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] T. W. Doeppner, Jr. Parallel program correctness through refinement. In *Proceedings of the ACM Symposium on the Principles of Programming Languages*, pages 155–169, 1977.

[9] J. L. Eppinger, L. B. Mummert, and A. Z. Spector. *Camelot and Avalon: A Distributed Transaction Facility*. Morgan Kaufmann, 1991.

[10] C. Flanagan. Verifying commit-atomicity using model-checking. In *SPIN 2004: 11th International SPIN Workshop on Model Checking of Software*, pages 252–266, 2004.

[11] 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.

[12] C. Flanagan and S. N. Freund. Atomizer: A dynamic atomicity checker for multithreaded programs. In *Proceedings of the ACM Symposium on the Principles of Programming Languages*, pages 256–267, 2004.

[13] C. Flanagan, S. N. Freund, and S. Qadeer. Exploiting purity for atomicity. Technical Note 04-05, Williams College, 2004.

[14] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for Java. In *Proceedings of the ACM Conference on Programming Language Design and Implementation*, pages 234–245, 2002.

[15] C. Flanagan and S. Qadeer. Transactions for software model checking. In *Proceedings of the Workshop on Software Model Checking*, 2003.

[16] 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.

[17] C. Flanagan and S. Qadeer. Types for atomicity. In *Proceedings of the ACM Workshop on Types in Language Design and Implementation*, pages 1–12, 2003.

[18] S. N. Freund and S. Qadeer. Checking concise specifications for multithreaded software. In *Journal of Object Technology*, 2004. (to appear). A preliminary version appeared in *Workshop on Formal Techniques for Java-like Programs*, 2003.

[19] P. Godefroid. Using partial orders to improve automatic verification methods. In *Proceedings of the IEEE Conference on Computer Aided Verification*, pages 176–185, 1991.

[20] D. Grossman. Type-safe multithreading in Cyclone. In *Proceedings of the ACM Workshop on Types in Language Design and Implementation*, pages 13–25, 2003.

[21] 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, pages 388–402, 2003.

[22] 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*, pages 175–190, 2004.

[23] M. P. Herlihy and J. M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, 1990.

[24] C. A. R. Hoare. Towards a theory of parallel programming. In *Operating Systems Techniques*, volume 9 of *A.P.I.C. Studies in Data Processing*, pages 61–71. Academic Press, 1972.

[25] L. Lamport and F. B. Schneider. Pretending atomicity. Research Report 44, DEC Systems Research Center, 1989.

[26] R. J. Lipton. Reduction: A method of proving properties of parallel programs. *Communications of the ACM*, 18(12):717–721, 1975.

[27] B. Liskov, D. Curtis, P. Johnson, and R. Scheifler. Implementation of Argus. In *Proceedings of the Symposium on Operating Systems Principles*, pages 111–122, 1987.

[28] D. B. Lomet. Process structuring, synchronization, and recovery using atomic actions. *Language Design for Reliable Software*, pages 128–137, 1977.

[29] J. Misra. *A Discipline of Multiprogramming: Programming Theory for Distributed Applications*. Springer-Verlag, 2001.

[30] F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer, 1999.

[31] C. Papadimitriou. *The theory of database concurrency control*. Computer Science Press, 1986.

[32] D. Peled. Combining partial order reductions with on-the-fly model checking. In D. Dill, editor, *CAV 94: Computer Aided Verification*, Lecture Notes in Computer Science 818, pages 377–390. Springer-Verlag, 1994.

[33] S. Qadeer, S. K. Rajamani, and J. Rehof. Summarizing procedures in concurrent programs. In *Proceedings of the ACM Symposium on the Principles of Programming Languages*, pages 245–255, 2004.

[34] Robby, E. Rodriguez, M. B. Dwyer, and J. Hatcliff. Checking strong specifications using an extensible software model checking framework. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 404–420, 2004.

[35] D. C. Schmidt and T. H. Harrison. Double-checked locking – a optimization pattern for efficiently initializing and accessing thread-safe objects. In R. Martin, F. Buschmann, and D. Riehle, editors, *Pattern Languages of Program Design 3*. Addison-Wesley, 1997.

[36] N. Sterling. WARLOCK — a static data race analysis tool. In *USENIX Technical Conference Proceedings*, pages 97–106, Winter 1993.

[37] S. D. Stoller. Model-checking multi-threaded distributed Java programs. In *Workshop on Model Checking and Software Verification*, pages 224–244, 2000.

[38] L. Wang and S. D. Stoller. Run-time analysis for atomicity. In *Proceedings of the Workshop on Runtime Verification*, volume 89(2) of *Electronic Notes in Computer Science*. Elsevier, 2003.

[39] A. Welc, S. Jagannathan, and A. L. Hosking. Transactional monitors for concurrent objects. In *Proceedings of the 18th European Conference on Object-Oriented Programming*, 2004.

A Effect system for purity

We present in this appendix an effect system to check that all normally-terminating `pure` expressions in a program are side-effect-free. This effect system is relatively simple but sufficient to check all examples in this paper. The effect system essentially tracks all locks acquired by each pure expression to ensure that these locks are released before termination. More complex analyses could improve precision by, for example, tracking more precise control-flow and data-flow information.

The effect system reasons about the judgment

$$\Pi, X \vdash_p e : L \to L'$$

where Π is the set of functions that are side-effect-free under normal-termination, and X is the set of variables that may change during evaluation of e. The set L is the set of locks held at the beginning of evaluation of e, and L' is the set of locks held after e terminates normally.

Most rules are straightforward. Any variable may be read, but only variables not appearing in X may be modified.

[PURE READ]

$$\frac{}{\Pi, X \vdash_p x_r : L \to L}$$

[PURE ASSIGN]

$$\frac{\Pi, X \vdash_p e : L_1 \to L_2 \qquad x \in X}{\Pi, X \vdash_p x_r := e : L_1 \to L_2}$$

The rules typically construct the set of locks held after evaluation by "threading" the lockset through each subexpression, as demonstrated by the rule for let expressions:

[PURE LET]

$$\frac{\Pi, X \vdash_p e_1 : L_1 \to L_2 \qquad \Pi, X \cup \{x\} \vdash_p e_2 : L_2 \to L_3}{\Pi, X \vdash_p \mathtt{let}\ x = e_1\ \mathtt{in}\ e_2 : L_1 \to L_3}$$

We introduce specific rules for the primitive operations that acquire and release mutual exclusion locks, as well as for the idiom of breaking when a `CAS` operation succeeds. Additional rules could model other synchronization primitives, as necessary.

[PURE ACQ]

$$\frac{x \notin X \quad x \notin L}{\Pi, X \vdash_p \mathtt{acquire}(x) : L \to L \cup \{x\}}$$

[PURE REL]

$$\frac{x \notin X \quad x \in L}{\Pi, X \vdash_p \mathtt{release}(x) : L \to L \setminus \{x\}}$$

[PURE IF CAS]

$$\frac{\Pi, X \vdash_p e_i : L \to L}{\Pi, X \vdash_p \mathtt{if}\ \mathtt{CAS}(e_1, e_2, e_3)\ \mathtt{break}\ e_4 : L \to L}$$

The top-level judgment

$$\Pi \vdash_p P$$

states that the annotation `pure` e is valid if

$$\Pi, \mathit{UnstableVar} \vdash_p e : \emptyset \to \emptyset$$

That is, a pure block may not change any stable variables or terminate with a different set of locks held than when evaluation started: see [PURE PROG]. This rule also requires that every function in Π is pure.

Purity Effect System

$$\boxed{\Pi, X \vdash_p e : L_1 \to L_2}$$

[PURE CONST]
$$\frac{}{\Pi, X \vdash_p v : L \to L}$$

[PURE WRONG]
$$\frac{}{\Pi, X \vdash_p \mathbf{wrong} : L \to L}$$

[PURE PRIM]
$$\frac{\Pi, X \vdash_p \overline{e} : L_1 \to L_2 \quad p \text{ is effect-free}}{\Pi, X \vdash_p p(\overline{e}) : L_1 \to L_2}$$

[PURE READ]
$$\frac{}{\Pi, X \vdash_p x_r : L \to L}$$

[PURE ASSIGN]
$$\frac{\Pi, X \vdash_p e : L_1 \to L_2 \quad x \in X}{\Pi, X \vdash_p x_r := e : L_1 \to L_2}$$

[PURE ACQ]
$$\frac{x \notin X \quad x \notin L}{\Pi, X \vdash_p \texttt{acquire}(x) : L \to L \cup \{x\}}$$

[PURE REL]
$$\frac{x \notin X \quad x \in L}{\Pi, X \vdash_p \texttt{release}(x) : L \to L \setminus \{x\}}$$

[PURE IF CAS]
$$\frac{\Pi, X \vdash_p e_i : L \to L}{\Pi, X \vdash_p \texttt{if CAS}(e_1, e_2, e_3)\ \texttt{break}\ e_4 : L \to L}$$

[PURE LOOP]
$$\frac{\Pi, X \vdash_p e : L \to L}{\Pi, X \vdash_p \texttt{loop}\ e : L \to L}$$

[PURE LET]
$$\frac{\Pi, X \vdash_p e_1 : L_1 \to L_2 \quad \Pi, X \cup \{x\} \vdash_p e_2 : L_2 \to L_3}{\Pi, X \vdash_p \texttt{let}\ x = e_1\ \texttt{in}\ e_2 : L_1 \to L_3}$$

[PURE IF]
$$\frac{\Pi, X \vdash_p e : L_1 \to L_2 \quad \Pi, X \vdash_p e_i : L_2 \to L_3}{\Pi, X \vdash_p \texttt{if}\ e\ e_1\ e_2 : L_1 \to L_3}$$

[PURE BREAK]
$$\frac{}{\Pi, X \vdash_p \texttt{break} : L \to L'}$$

[PURE INVOKE]
$$\frac{\Pi, X \vdash_p e : L_1 \to L_2 \quad \Pi, X \vdash_p \bar{e} : L_2 \to L_3 \quad F \subseteq \Pi}{\Pi, X \vdash_p e^F(\overline{e}) : L_1 \to L_3}$$

$$\boxed{\Pi, X \vdash_p \bar{e} : L_1 \to L_2}$$

[PURE EMPTY SEQ]
$$\frac{}{\Pi, X \vdash_p \epsilon : L \to L}$$

[PURE SEQ]
$$\frac{\Pi, X \vdash_p \bar{e} : L_1 \to L_2 \quad \Pi, X \vdash_p e : L_2 \to L_3}{\Pi, X \vdash_p \bar{e}, e : L_2 \to L_3}$$

$$\boxed{\Pi \vdash_p P}$$

[PURE PROG]
$$\frac{\begin{array}{rcl} P \text{ contains } f(\overline{x})\ e \text{ and } f \in \Pi & \Rightarrow & \Pi, \mathit{UnstableVar} \vdash_p e : \emptyset \to \emptyset \\ P \text{ contains } \texttt{pure}\ e & \Rightarrow & \Pi, \mathit{UnstableVar} \vdash_p e : \emptyset \to \emptyset \end{array}}{\Pi \vdash_p P}$$

A.1 Correctness proof for Purity

THEOREM 1 (CORRECTNESS OF THE PURITY EFFECT SYSTEM). *If there exists Π such that $\Pi \vdash_p P$ then the purity annotations in P are correct.*

PROOF. Suppose

$$\begin{array}{rl} P & \to^* \quad H_1, M_1, T.E[\texttt{pure}\ e].T' \\ & \to_i^* \quad H_2, M_2, T.E[v].T'.T'' \end{array}$$

Then we must have

$$\begin{array}{rl} P & \to^* \quad H_1, M_1, T.E[\texttt{pure}\ e].T' \\ & \to_i \quad H_1, M_1, T.E[\texttt{in-pure}\ e].T' \\ & \to_i^* \quad H_2, M_2, T.E[\texttt{in-pure}\ v].T'.T'' \\ & \to_i \quad H_2, M_2, T.E[v].T'.T'' \end{array}$$

Then by Theorem 2,

$$\Pi \vdash_p H_1, M_1, T.E[\texttt{pure}\ e].T'$$

and therefore

$$\Pi, \mathit{UnstableVar} \vdash_p e : \emptyset \to \emptyset$$

By Lemma 2 we have $T'' = \epsilon$ and H_1 and H_2 agree except on $\mathit{UnstableVar}$ and M_1 and M_2 agree on $\mathit{Dom}(M_1)$, and therefore

$$H_1, M_1, T.E[v].T' \to_i^* H_2, M_2, T.E[v].T'$$

as required. □

THEOREM 2 (PURITY SUBJECT REDUCTION). *If $\Pi \vdash_p P$ and $P \to P'$ then $\Pi \vdash_p P'$.*

PROOF. The definition of $\Pi \vdash_p P$ only concerns expressions that are within the body of a function definition or a `pure` construct. Such expressions are not in evaluation context positions and hence are not affected by any of the possible reduction rules for $P \to P'$. Thus $\Pi \vdash_p P'$. □

Let $\mathit{locks}(M, i) = \{m \mid M(m) = \langle \texttt{lock}, \mathit{tid} \rangle\}$. For $L \subseteq \mathit{locks}(M, i)$, let the synchronization store $\mathit{rel}(M, L)$ be defined by

$$\mathit{rel}(M, L)(m) = \begin{cases} \langle \texttt{lock}, 0 \rangle & \text{if } m \in L \\ M(m) & \text{otherwise} \end{cases}$$

LEMMA 1 (REPLACEMENT). *Suppose*

$$\Pi, X \vdash_p E[e] : L_1 \to L_2$$

where e is not $\texttt{CAS}(v_1, v_2, v_3)$. Then there exists L_3 such that

$$\Pi, X \vdash_p e : L_1 \to L_3$$

Furthermore, for any e' and L_4 such that

$$\Pi, X \vdash_p e' : L_4 \to L_3$$

we have

$$\Pi, X \vdash_p E[e'] : L_4 \to L_2$$

PROOF. By induction on the size and shape of E.

- For the base case, if $E = [\,]$ then $L_2 = L_3$ and this case clearly holds.

- Suppose $E = (x_r := E')$. Then

$$\begin{array}{ll} & \Pi, X \vdash_p x_r := E'[e] : L_1 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p E'[e] : L_1 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p e : L_1 \to L_3 \quad \text{by IH} \end{array}$$

  Suppose $\Pi, X \vdash_p e' : L_4 \to L_3$. Then

$$\begin{array}{lll} & \Pi, X \vdash_p e' : L_4 \to L_3 & \\ \Rightarrow & \Pi, X \vdash_p E'[e'] : L_4 \to L_2 & \text{by IH} \\ \Rightarrow & \Pi, X \vdash_p E[e'] : L_4 \to L_2 & \text{by [PURE ASSIGN]} \end{array}$$

- Suppose $E = p(\overline{v}, E', \overline{e})$. Use the following derived rule:

$$\text{[PURE SEQ 2]} \quad \frac{\Pi, X \vdash_p \overline{v} : L_1 \to L_1 \quad \Pi, X \vdash_p e : L_1 \to L_2 \quad \Pi, X \vdash_p \overline{e} : L_2 \to L_3}{\Pi, X \vdash_p \overline{v}, e, \overline{e} : L_1 \to L_3}$$

  Then

$$\begin{array}{ll} & \Pi, X \vdash_p p(\overline{v}, E'[e], \overline{e}) : L_1 \to L_3 \\ \Rightarrow & \Pi, X \vdash_p \overline{v}, E'[e], \overline{e} : L_1 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p \overline{v} : L_1 \to L_1 \\ \text{and} & \Pi, X \vdash_p E'[e] : L_1 \to L_2 \\ \text{and} & \Pi, X \vdash_p \overline{e} : L_2 \to L_3 \\ \Rightarrow & \Pi, X \vdash_p e : L_1 \to L_4 \quad \text{by IH} \end{array}$$

  Suppose $\Pi, X \vdash_p e' : L_5 \to L_4$. Then

$$\begin{array}{lll} & \Pi, X \vdash_p e' : L_5 \to L_4 & \\ \text{and} & \Pi, X \vdash_p \overline{v} : L_5 \to L_5 & \\ \Rightarrow & \Pi, X \vdash_p E'[e'] : L_5 \to L_2 & \text{by IH} \\ \Rightarrow & \Pi, X \vdash_p E[e'] : L_4 \to L_3 & \text{by [PURE SEQ 2]} \end{array}$$

- Suppose $E = E'^F(\overline{e})$. Then

$$\begin{array}{ll} & \Pi, X \vdash_p E'[e]^F(\overline{e}) : L_1 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p E'[e] : L_1 \to L_3 \\ \text{and} & \Pi, X \vdash_p \overline{e} : L_3 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p e : L_1 \to L_4 \quad \text{by IH} \end{array}$$

  Suppose $\Pi, X \vdash_p e' : L_5 \to L_4$. Then

$$\begin{array}{lll} & \Pi, X \vdash_p e' : L_5 \to L_4 & \\ \Rightarrow & \Pi, X \vdash_p E'[e'] : L_5 \to L_3 & \text{by IH} \\ \Rightarrow & \Pi, X \vdash_p E[e'] : L_5 \to L_2 & \text{by [PURE INVOKE]} \end{array}$$

- Suppose $E = v^F(\overline{v}, E', \overline{e})$. Similar to rule for p.

- Suppose $E = (\texttt{if}\ E'\ e_1\ e_2)$. Then

$$\begin{array}{ll} & \Pi, X \vdash_p \texttt{if}\ E'[e]\ e_1\ e_2 : L_1 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p E'[e] : L_1 \to L_3 \\ \text{and} & \Pi, X \vdash_p e_i : L_3 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p e : L_1 \to L_4 \quad \text{by IH} \end{array}$$

  Suppose $\Pi, X \vdash_p e' : L_5 \to L_4$. Then

$$\begin{array}{lll} & \Pi, X \vdash_p e' : L_5 \to L_4 & \\ \Rightarrow & \Pi, X \vdash_p E'[e'] : L_5 \to L_3 & \text{by IH} \\ \Rightarrow & \Pi, X \vdash_p E[e'] : L_5 \to L_2 & \text{by [PURE IF]} \end{array}$$

- Suppose $E = (\texttt{let}\ x = E'\ \texttt{in}\ e_2)$. Then

$$\begin{array}{ll} & \Pi, X \vdash_p \texttt{let}\ x = E'[e]\ \texttt{in}\ e_2 : L_1 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p E'[e] : L_1 \to L_3 \\ \text{and} & \Pi, X \cup \{x\} \vdash_p e_2 : L_3 \to L_2 \\ \Rightarrow & \Pi, X \vdash_p e : L_1 \to L_4 \quad \text{by IH} \end{array}$$

  Suppose $\Pi, X \vdash_p e' : L_5 \to L_4$. Then

$$\begin{array}{lll} & \Pi, X \vdash_p e' : L_5 \to L_4 & \\ \Rightarrow & \Pi, X \vdash_p E'[e'] : L_5 \to L_3 & \text{by IH} \\ \Rightarrow & \Pi, X \vdash_p E[e'] : L_5 \to L_2 & \text{by [PURE LET]} \end{array}$$

□

LEMMA 2. *Let $i = |T|$. If*

$$\begin{array}{l} \Pi, X \vdash_p e : L \to \emptyset \\ L \subseteq \mathit{locks}(M, i) \\ H, M, T.E[e].T' \to_i^* H', M', T.E[v].T'.T'' \end{array}$$

and in each intermediate state the i-th thread has outermost evaluation context E then

$$\begin{array}{l} T'' = \epsilon \\ H \text{ and } H' \text{ agree on } \mathit{Dom}(H) - X \\ \mathit{rel}(M, L) \text{ and } M' \text{ agree on } \mathit{Dom}(M) \end{array}$$

PROOF. The proof proceeds by induction on the length of the execution sequence.

- The base case clearly holds for execution sequences of length zero.

- For the inductive case, we have

$$\begin{array}{lll} H, M, T.E[e].T' & \to_i & H'', M'', T.E[e''].T'.T'' \\ & \to_i^* & H', M', T.E[v].T'.T''.T''' \end{array}$$

  - Suppose that the first evaluation step is not via [RED PRIM] for a `CAS` operation that succeeds. Then by Lemma 3 we have there exists L'' such that:

$$\begin{array}{l} T'' = \epsilon \\ \Pi, X \vdash_p e'' : L'' \to \emptyset \\ L'' \subseteq \mathit{locks}(M'', i) \\ H \text{ and } H'' \text{ agree on } \mathit{Dom}(H) \setminus X \\ \mathit{rel}(M, L) \text{ and } \mathit{rel}(M'', L'') \text{ agree on } \mathit{Dom}(M) \end{array}$$

    By induction we have that:

$$\begin{array}{l} T''' = \epsilon \\ H'' \text{ and } H' \text{ agree on } \mathit{Dom}(H'') - X \\ \mathit{rel}(M'', L'') \text{ and } M' \text{ agree on } \mathit{Dom}(M'') \end{array}$$

Since $Dom(H) \subset Dom(H'')$ and $Dom(M) \subset Dom(M'')$, we then have:

H and H' agree on $Dom(H) - X$
$rel(M, L)$ and M' agree on $Dom(M)$

Thus the lemma holds for this case.

– Suppose that the first evaluation step is via [RED PRIM] for a `CAS` operation that succeeds. Since the only rule for a `CAS` operation is [PURE IF CAS], we have

$$e = E'[\texttt{if CAS}(v_1, v_2, v_3)\ \texttt{break}\ e_4]$$

and the evaluation sequence

$$\begin{array}{ll} & H, M, T.E[e].T' \\ = & H, M, T.E[E'[\texttt{if CAS}(v_1, v_2, v_3)\ \texttt{break}\ e_4]].T' \\ \rightarrow_i & H', M'', T.E[E'[\texttt{if}\ \mathit{true}\ \texttt{break}\ e_4]].T' \\ \rightarrow_i & H', M'', T.E[E'[\texttt{break}]].T' \end{array}$$

Our purity effect system ensures that E' does not include `block` evaluation context, hence the `break` breaks to somewhere in E, which does not satisfy the assumptions of this lemma.

□

LEMMA 3. *Let $i = |T|$. If*

$\Pi, X \vdash_p e : L \rightarrow \emptyset$
$L \subseteq locks(M, i)$
$H, M, T.E[e].T' \rightarrow_i H', M', T.E[e'].T'.T''$

*and the evaluation step is not via [*RED PRIM*] for a* `CAS` *operation that succeeds, then there exists L' such that for all values v,*

$T'' = \epsilon$
$\Pi, X \vdash_p e' : L' \rightarrow \emptyset$
$L' \subseteq locks(M', i)$
H and H' agree on $Dom(H) - X$
$rel(M, L)$ and $rel(M', L')$ agree on $Dom(M)$

PROOF. The proof proceeds by case analysis on the reduction rule used.

- [RED READ] In this case

$$\begin{array}{rcl} e &=& E'[x_r] \\ e' &=& E'[H(x)] \\ T'' &=& \epsilon \\ H' &=& H \\ M' &=& M \end{array}$$

Clearly

$\Pi, X \vdash_p x_r : L \rightarrow L$
$\Pi, X \vdash_p H(x) : L \rightarrow L$

Taking $L' = L$, we have

$$\Pi, X \vdash_p e' : L' \rightarrow \emptyset$$

and the lemma holds for this case.

- [RED WRITE] In this case

$$\begin{array}{rcl} e &=& E'[x_r := v] \\ e' &=& E'[v] \\ T'' &=& \epsilon \\ H' &=& H[x := v] \\ M' &=& M \\ x &\in& X \end{array}$$

Clearly

$\Pi, X \vdash_p x_r := v : L \rightarrow L$
$\Pi, X \vdash_p v : L \rightarrow L$

Taking $L' = L$, we have

$$\Pi, X \vdash_p e' : L' \rightarrow \emptyset$$

and the lemma holds for this case.

- [RED PRIM] We consider three cases separately, according to the three purity rules for primitive applications.
  - Consider first the case where p is effect-free. In this case

$$\begin{array}{rcl} e &=& E'[p(\overline{v})] \\ e' &=& E'[v'] \\ T'' &=& \epsilon \\ H' &=& H \\ M' &=& M \quad \text{as } p \text{ is effect-free} \end{array}$$

Clearly

$\Pi, X \vdash_p p(\overline{v}) : L \rightarrow L$
$\Pi, X \vdash_p v' : L \rightarrow L$

Taking $L' = L$, we have

$$\Pi, X \vdash_p e' : L' \rightarrow \emptyset$$

and the lemma holds for this case.

  - Suppose $p = \texttt{acquire}$. In this case

$$\begin{array}{rcl} e &=& E'[\texttt{acquire}(m)] \\ e' &=& E'[m] \\ T'' &=& \epsilon \\ H' &=& H \\ M(m) &=& \langle \texttt{lock}, 0 \rangle \\ M' &=& M[m := \langle \texttt{lock}, \mathit{tid} \rangle] \end{array}$$

Clearly

$\Pi, X \vdash_p \texttt{acquire}(m) : L \rightarrow L \cup \{m\}$
$\Pi, X \vdash_p m : L \cup \{m\} \rightarrow L \cup \{m\}$

Taking $L' = L \cup \{m\}$, we have

$\Pi, X \vdash_p e' : L' \rightarrow \emptyset$
$rel(M, L) = rel(M', L')$

and the lemma holds for this case.

  - Suppose $p = \texttt{release}$. In this case

$$\begin{array}{rcl} e &=& E'[\texttt{release}(m)] \\ e' &=& E'[m] \\ T'' &=& \epsilon \\ H' &=& H \\ M(m) &=& \langle \texttt{lock}, \mathit{tid} \rangle \\ M' &=& M[m := \langle \texttt{lock}, 0 \rangle] \end{array}$$

Clearly

$$\Pi, X \vdash_p \texttt{acquire}(m) : L \to L \setminus \{m\}$$
$$\Pi, X \vdash_p v' : L \setminus \{m\} \to L \setminus \{m\}$$

Taking $L' = L \setminus \{m\}$, we have

$$\Pi, X \vdash_p e' : L' \to \emptyset$$
$$rel(M, L) = rel(M', L')$$

and the lemma holds for this case.

- [RED WRONG], [RED CALL], [RED IF1], [RED IF2], [RED LOOP], [RED LET], [ABS UNSTABLE READ], [ABS UNSTABLE WRITE], [ABS GARBAGE]: These cases hold by reasoning similar to the above.
- [RED FORK], [RED ATOMIC], [RED IN-ATOMIC], [RED PURE], [RED PURE EX], [RED IN-PURE], [RED BLOCK], [ABS PURE SKIP]: None of these cases are possible since the purity effect system does not support these constructs.

□

B Formal semantics of CAP

To reason about the correctness of our analysis, we formalize the run-time behavior of CAP programs using an operational semantics. For this purpose, we extend CAP with additional run-time expressions, which record that execution is proceeding inside an atomic block or inside a pure block. In addition, since non-terminating pure blocks should not change the program state, we distinguish evaluation of a pure block that will terminate normally (`in-pure`) from one that will terminate exceptionally (`in-pureX`). We also extend the set of values to include synchronization locations, which are manipulated by the synchronization primitives.

Run-time Syntax

$$\begin{array}{llll} e & \in & \mathit{Expr} & ::= \ \ldots \mid \texttt{in-atomic}\ e \\ & & & \quad\ \mid \texttt{in-pure}\ e \mid \texttt{in-pureX}\ e \\ v & \in & \mathit{Value} & ::= \ \ldots \mid m \\ m & \in & \mathit{SyncLoc} & \end{array}$$

A program state P is a three-tuple H, M, T. A heap H is a partial map from variables to values. A synchronization heap M maps synchronization locations to synchronization values, and T is a sequence of threads. A thread is either an expression or **wrong**.

Evaluation contexts define the order of evaluation for CAP. An evaluation context E is an expression with a "hole" at the location of the next sub-expression to be evaluated, and $E[e]$ denotes the operation of filling the hole in E with the expression e. The notation $H[x := v]$ denotes a new heap that is identical to H except that it maps x to v.

The transition relation $\to_i$ performs a single step of thread i. In the standard semantics, evaluation of `pure` e has two possible outcomes. For the case where e terminates normally, yielding some value v, we reduce `pure` e to `in-pure` e via [RED PURE], which in turn evaluates to `in-pure` v, which reduces to v via [RED IN-PURE]. Alternatively, if e terminates exceptionally (via `break`), we reduce `pure` e to `in-pureX` e via [RED PURE EX], which evaluates to `in-pureX` E[`break`] and the `break` propagates outside the `in-pureX` expression via [RED BREAK]. Thus, these rules characterize all possible executions of `pure` e, but also capture (via `in-pure`/`in-pureX`) the future termination mode of e, which is necessary for the correctness proof.

Note that this semantics requires a choice to be made about the behavior of e when the `pure` e block is entered. The wrong choice will simply cause evaluation to block when incorrect behavior is observed. For example, the [RED BREAK] rule is not applicable if a `break` causes an exceptional exit from an `in-pure` block. Similarly, `in-pureX` v cannot be further reduced.

The transition relation $\leadsto$ performs a single step of an arbitrarily chosen thread in the standard semantics. However, we verify atomicity with respect to an *abstract semantics* $\to$ that admits additional execution sequences. First, in the abstract semantics the evaluation of `pure` e may be skipped entirely and an arbitrary value returned via [ABS PURE SKIP]. The abstract semantics also includes the rules [ABS UNSTABLE READ] and [ABS UNSTABLE WRITE] for reading/writing arbitrary values from/to unstable variables. Since `pure` expressions may create garbage by allocating new temporaries, we also introduce the "garbage creation" rule [ABS GARBAGE]. This rule enables us to show that a trace including the normal evaluation of a `pure` expression is equivalent to a trace with its evaluation replaced by a [ABS PURE SKIP] step followed by a [ABS GARBAGE] step. While not strictly necessary, we permit that rule to change unstable variables to simplify some technical aspects of our correctness proof.

We use $\to^+$ and $\to^*$ to denote the transitive and reflexive-transitive closure of $\to$, respectively. The *serial abstract semantics* $\mapsto$ is similar to $\to$, with the additional restriction that a thread cannot perform a step if another thread is inside an `in-atomic` block. Thus $\mapsto$ does not interleave the execution of an `atomic` block with instructions from other threads.

B.1 Primitive operations

The meaning of a primitive operation p is given by the partial function $\mathcal{I}_p$, which takes a sequence of argument values, a synchronization heap, and an integer identifying the current thread, and returns a value and a (possibly modified) synchronization heap. If the primitive is applied to incorrect arguments, $\mathcal{I}_p$ returns $\mathbf{wrong}$.

$$\begin{array}{l} \mathcal{I}_p : (\mathit{Const} \cup \mathit{SyncLoc})^* \times \mathit{SyncHeap} \times \mathit{Int} \rightharpoonup \\ \quad ((\mathit{Const} \cup \mathit{SyncLoc}) \times \mathit{SyncHeap}) \cup \{\mathbf{wrong}\} \end{array}$$

The `assert` primitive operation goes wrong if its argument is 0, and otherwise terminates normally without modifying the heap. Addition is a pure primitive operation, and does not modify the synchronization heap. The `new_lock` operation returns a synchronization location m that refers to a newly-allocated lock containing 0, indicating that is not held by any thread. The `acquire` operation acquires a lock, provided the lock is not held by another thread. If the lock is held by some thread, then the `acquire` operation blocks, and execution can only proceed on the other threads. The `release` operation never blocks; if the current thread holds the lock then the lock is released, and otherwise the `release` operation goes wrong. We define the semantics of these primitive operations as follows (where ϵ is the empty sequence and $v.s$ is prepends v at the front of sequence s):

$$\mathcal{I}_{\texttt{assert}}(v, M, \mathit{tid}) = \begin{cases} \langle 0, M \rangle & \text{if } v \neq 0 \\ \mathbf{wrong} & \text{if } v = 0 \end{cases}$$

Semantics

State space

$$
\begin{array}{lcrcl}
P & \in & \mathit{State} & = & \mathit{Heap} \times \mathit{SyncHeap} \times \mathit{ThreadSeq} \\
H & \in & \mathit{Heap} & = & \mathit{Var} \rightarrow\!\!\!\!\!\!\circ\;\; \mathit{Value} \\
M & \in & \mathit{SyncHeap} & = & \mathit{SyncLoc} \rightarrow\!\!\!\!\!\!\circ\;\; \mathit{SyncValue} \\
T & \in & \mathit{ThreadSeq} & = & (\mathit{Expr} \cup \mathbf{wrong})^*
\end{array}
$$

Evaluation contexts

$$
\begin{array}{rcl}
E & ::= & [\,] \mid x_r := E \mid p(\overline{v}, E, \overline{e}) \mid E^F(\overline{e}) \mid v^F(\overline{v}, E, \overline{e}) \\
 & & \mid \texttt{if}\ E\ e\ e \mid \texttt{let}\ x = E\ \texttt{in}\ e \mid \texttt{block}\ E \\
 & & \mid \texttt{in-atomic}\ E \mid \texttt{in-pure}\ E \mid \texttt{in-pureX}\ E
\end{array}
$$

Transition relations $\quad \rightarrow_i, \rightarrow, \mapsto, \rightsquigarrow \quad \subseteq \quad \mathit{State} \times \mathit{State}$

Transition rules (where $i = |T| + 1$)

[RED READ]	$H, M, T.E[x_r].T'$	$\rightarrow_i$	$H, M, T.E[H(x)].T'$	
[RED WRITE]	$H, M, T.E[x_r := v].T'$	$\rightarrow_i$	$H[x := v], M, T.E[v].T'$	if $x \in \mathit{Dom}(H)$
[RED PRIM]	$H, M, T.E[p(\overline{v})].T'$	$\rightarrow_i$	$H, M', T.E[v'].T'$	if $\mathcal{I}_p(\overline{v}, M, i) = \langle v', M' \rangle$
[RED WRONG]	$H, M, T.E[p(\overline{v})].T'$	$\rightarrow_i$	$H, M, T.\mathbf{wrong}.T'$	if $\mathcal{I}_p(\overline{v}, M, i) = \mathbf{wrong}$
[RED CALL]	$H, M, T.E[(f(\overline{x})\ e)^F(\overline{v})].T'$	$\rightarrow_i$	$H[\overline{x} := \overline{v}], T.E[e].T'$	if $\overline{x} \cap \mathit{dom}(H) = \emptyset$
[RED IF1]	$H, M, T.E[\texttt{if}\ v\ e_1\ e_2].T'$	$\rightarrow_i$	$H, M, T.E[e_1].T'$	if $v \neq 0$
[RED IF2]	$H, M, T.E[\texttt{if}\ 0\ e_1\ e_2].T'$	$\rightarrow_i$	$H, M, T.E[e_2].T'$	
[RED LOOP]	$H, M, T.E[\texttt{loop}\ e].T'$	$\rightarrow_i$	$H, M, T.E[e; \texttt{loop}\ e].T'$	
[RED LET]	$H, M, T.E[\texttt{let}\ x = v\ \texttt{in}\ e].T'$	$\rightarrow_i$	$H[x := v], M, T.E[e].T'$	if $x \notin \mathit{dom}(H)$
[RED FORK]	$H, M, T.E[\texttt{fork}\ e].T'$	$\rightarrow_i$	$H, M, T.E[0].T'.e$	
[RED ATOMIC]	$H, M, T.E[\texttt{atomic}\ e].T'$	$\rightarrow_i$	$H, M, T.E[\texttt{in-atomic}\ e].T'$	
[RED IN-ATOMIC]	$H, M, T.E[\texttt{in-atomic}\ v].T'$	$\rightarrow_i$	$H, M, T.E[v].T'$	
[RED PURE]	$H, M, T.E[\texttt{pure}\ e].T'$	$\rightarrow_i$	$H, M, T.E[\texttt{in-pure}\ e].T'$	
[RED PURE EX]	$H, M, T.E[\texttt{pure}\ e].T'$	$\rightarrow_i$	$H, M, T.E[\texttt{in-pureX}\ e].T'$	
[RED IN-PURE]	$H, M, T.E[\texttt{in-pure}\ v].T'$	$\rightarrow_i$	$H, M, T.E[v].T$	
[RED BREAK]	$H, M, T.E[\texttt{block}\ E'[\texttt{break}]].T'$	$\rightarrow_i$	$H, M, T.E[0].T'$	if E' does not contain `block` [ ] or `in-pure` [ ]
[RED BLOCK]	$H, M, T.E[\texttt{block}\ v].T'$	$\rightarrow_i$	$H, M, T.E[v].T'$	
[ABS PURE SKIP]	$H, M, T.E[\texttt{pure}\ e].T'$	$\rightarrow_i$	$H, M, T.E[v].T'$	
[ABS UNSTABLE READ]	$H, M, T.E[x_r].T'$	$\rightarrow_i$	$H, M, T.E[v].T'$	if $x \in \mathit{Dom}(H) \cap \mathit{UnstableVar}$
[ABS UNSTABLE WRITE]	$H, M, T.E[x_r := v].T'$	$\rightarrow_i$	$H[x := v'], M, T.E[v].T'$	if $x \in \mathit{Dom}(H) \cap \mathit{UnstableVar}$
[ABS GARBAGE]	H, M, T $\mathit{Dom}(H) \cap \mathit{StableVar}$	$\rightarrow_i$	H', M', T	if H and H' agree on and M and M' agree on Dom(M)
(standard semantics)	H, M, T	$\rightsquigarrow$	P'	if $H, M, T \rightarrow_i P'$ by a rule other than [ABS PURE SKIP], [ABS GARBAGE], [ABS UNSTABLE WRITE], or [ABS UNSTABLE READ]
(abstract semantics)	H, M, T	$\rightarrow$	P'	if $H, M, T \rightarrow_i P'$
(serial abstract semantics)	H, M, T	$\mapsto$	P'	if $H, M, T \rightarrow_i P'$ and T_j does not contain `in-atomic` for $j \neq i$

$$
\begin{array}{l}
\mathcal{I}_+(m.n, M, \mathit{tid}) = \quad \langle m + n, M \rangle \\
\mathcal{I}_{\texttt{new\_lock}}(\epsilon, M, \mathit{tid}) = \\
\qquad \langle m, M[m := \langle \texttt{lock}, 0 \rangle] \rangle \qquad \text{if } m \notin \mathit{dom}(M) \\
\mathcal{I}_{\texttt{acquire}}(m, M[m := \langle \texttt{lock}, 0 \rangle], \mathit{tid}) = \\
\qquad \langle m, M[m := \langle \texttt{lock}, \mathit{tid} \rangle] \rangle \\
\mathcal{I}_{\texttt{release}}(m, M, \mathit{tid}) = \\
\qquad \begin{cases} \langle m, M[m := \langle \texttt{lock}, 0 \rangle] \rangle & \text{if } M(m) = \langle \texttt{lock}, \mathit{tid} \rangle \\ \mathbf{wrong} & \text{otherwise} \end{cases} \\
\mathcal{I}_{\texttt{CAS}}(m.c_1.c_2, M, \mathit{tid}) = \\
\qquad \begin{cases} \langle 0, M \rangle & \text{if } M(m) \neq c_1 \\ \langle 1, M[m := c_2] \rangle & \text{if } M(m) = c_1 \end{cases}
\end{array}
$$

We now formalize our assumption on primitive atomicities. A primitive p *right-commutes* with a primitive q if for all i and j such that $i \neq j$, the following conditions hold:

1. If $\mathcal{I}_p(\overline{v}, M, i) = \langle v', M' \rangle$ and $\mathcal{I}_q(\overline{u}, M', j) = \langle u', M'' \rangle$, then there is M''' such that $\mathcal{I}_q(\overline{u}, M, j) = \langle u', M''' \rangle$ and $\mathcal{I}_p(\overline{v}, M''', i) = \langle v', M'' \rangle$.
2. If $\mathcal{I}_p(\overline{v}, M, i) = \mathbf{wrong}$ and $\mathcal{I}_q(\overline{u}, M, j) = \langle u', M' \rangle$, then $\mathcal{I}_p(\overline{v}, M', i) = \mathbf{wrong}$.
3. If $\mathcal{I}_p(\overline{v}, M, i) = \langle v', M' \rangle$ and $\mathcal{I}_q(\overline{u}, M', j) = \mathbf{wrong}$, then $\mathcal{I}_q(\overline{u}, M, j) = \mathbf{wrong}$.

The primitive p *left-commutes* with the primitive q if q right-commutes with p. An effect environment is *valid* if the atomicity $\Gamma(p)$ of each primitive p is correct, in that the following properties hold:

1. If $\Gamma(p) \sqsubseteq R$, then p right-commutes with any primitive q.
2. If $\Gamma(p) \sqsubseteq L$, then p left-commutes with any primitive q.
3. $B \sqsubseteq \Gamma(p)$.

Note that since primitive operations only read or write the synchronization heap, and not the regular heap, they trivially commute with all non-primitive operations, which only access the regular heap.

B.2 Correctness of annotations

This section defines the meaning of the conflict annotations, call annotations, and purity in terms of the CAP semantics. An expression e is *about to read* x if $e \equiv E[x_r]$. Similarly, e is *about to write* x if $e \equiv E[x_r := v]$. An expression e is *about to access* x if e is about to read or write x. The following conditions characterize correct annotations. As mentioned earlier, auxiliary type and effect analyses may check conformance to these requirements.

1. The conflict annotations in a program P are *correct* if whenever $P \rightarrow^* H, M, T$:
   (a) if $T_i \equiv E[x_\epsilon]$ and $x \in \mathit{StableVar}$, then no other thread in T is about to write x, and
   (b) if $T_i \equiv E[x_\epsilon := v]$ and $x \in \mathit{StableVar}$, then no other thread in T is about to read or write x.
2. The call annotations in P are *correct* if whenever $P \rightarrow^* H, M, T$ and $T_i \equiv E[(f(\overline{x})\ e)^F(\overline{v})]$, then $f \in F$.
3. The purity annotations in P are *correct* if whenever $i = |T|$ and

$$\begin{array}{rll} P & \rightarrow^* & H_1, M_1, T.E[\texttt{pure}\ e].T' \\ & \rightarrow_i^* & H_2, M_2, T.E[v].T'.T'' \end{array}$$

   and in each intermediate state the i-th thread has outermost evaluation context E then $T'' = \epsilon$ and

$$H_1, M_1, T.E[v].T' \rightarrow_i^* H_2, M_2, T.E[v].T'$$

C Correctness proof

Our effect system guarantees a correspondence between non-serial executions and serial abstract executions of any well-formed program. In particular, if under the non-serial semantics ($\leadsto$) a well-formed program P can reach a state Q where no thread is executing an atomic block, then P can also reach Q under the serial abstract semantics ($\mapsto$). Hence, the serial abstract semantics, in which we do not have to consider interleaved actions of concurrent threads inside an atomic block, suffices for reasoning about the reachability of such states Q.

In the absence of `pure` blocks, we can prove this correspondence property via reduction [26], since each atomic block consists of a sequence of right-movers, followed by at most one atomic action, followed by a sequence of left-movers. However, `pure` blocks significantly complicate the correctness proof.

We begin by showing via subject reduction that every state reachable from an initial, well-formed state is well-formed. For this purpose, we extend the effect system to handle run-time expressions and states, as follows. The rule [PROG] ensures that run-time expressions only appear in evaluation contexts.

Effect System for Run-time States

$\boxed{\Gamma \vdash e : a \uparrow b}$

[EXP REF]
$$\frac{}{\Gamma \vdash m : B \uparrow \bot}$$

[EXP INATOMIC]
$$\frac{\Gamma \vdash e : a \uparrow b \quad a, b \sqsubseteq A}{\Gamma \vdash \texttt{in-atomic}\ e : a \uparrow b}$$

[EXP INPURE]
$$\frac{\Gamma \vdash e : a \uparrow b \quad a \sqsubseteq A}{\Gamma \vdash \texttt{in-pure}\ e : B \uparrow \bot}$$

[EXP INPURE EX]
$$\frac{\Gamma \vdash e : a \uparrow b \quad a \sqsubseteq A}{\Gamma \vdash \texttt{in-pureX}\ e : \bot \uparrow b}$$

[EXP WRONG]
$$\frac{}{\Gamma \vdash \textbf{wrong} : B \uparrow \bot}$$

$\boxed{\Gamma \vdash P}$

[PROG]
$$\frac{\begin{array}{c} \forall x \in \mathit{dom}(H).\ \Gamma \vdash H(x) : B \uparrow \bot \\ \Gamma \vdash T_i : a_i \uparrow b_i \qquad \forall i \in 1..n \\ \forall e, i.(T_i \text{ contains } \texttt{in-atomic}\ e \Rightarrow \exists E.\ e = E[\texttt{in-atomic}\ e]) \\ \forall e, i.(T_i \text{ contains } \texttt{in-pure}\ e \Rightarrow \exists E.\ e = E[\texttt{in-pure}\ e]) \\ \forall e, i.(T_i \text{ contains } \texttt{in-pureX}\ e \Rightarrow \exists E.\ e = E[\texttt{in-pureX}\ e]) \end{array}}{\Gamma \vdash H, M, T_1 \ldots T_n}$$

THEOREM 3 (SUBJECT REDUCTION). *If the call annotations in P are correct and $\Gamma \vdash P$ and $P \rightarrow P'$, then $\Gamma \vdash P'$.*

Let Γ be a fixed effect environment and let WT be the set of well-formed states under Γ.

$$WT \quad = \quad \{\langle H, M, T\rangle \mid \Gamma \vdash \langle H, M, T\rangle\}$$

If $\Gamma \vdash e : a \uparrow b$, then we define the normal atomicity $\alpha_n(e)$, exceptional atomicity $\alpha_x(e)$, and (combined) atomicity $\alpha(e)$ of the expression e as

$$\begin{array}{rcl} \alpha_n(e) & = & a \\ \alpha_x(e) & = & b \\ \alpha(e) & = & a \sqcup b \end{array}$$

An examination of the effect rules shows that these are three well-defined partial functions (see Lemma 9).

Each pure block that terminates normally evaluates via a sequence of `in-pure` states. Our correctness proof proceeds by induction on the nesting depth n of `in-pure` blocks. First, we reduce each `in-pure` block at depth n to a serial execution, in which all of the actions of the `in-pure` block occur contiguously. Then, by the definition of purity, we can replace that `in-pure` fragment of the execution sequence with a single [ABS PURE SKIP] step, followed by an [ABS GARBAGE] step to create any garbage generated by the removed `in-pure` fragment. This technique reduces the nesting depth of `in-pure` blocks by one. Proceeding in this manner eventually yields an execution sequence with no `in-pure` blocks, at which point we can apply reduction to serialize each atomic block and produce an execution of the serial semantics ($\mapsto$).

To assist in reasoning about the nesting depth of `in-pure` blocks, we define D^n to denote evaluation contexts with at least n nested `in-pure` blocks, so that $E = D^0 \supset D^1 \supset \cdots$.

$$\begin{array}{rcl} D^0 & = & E \\ D^{n+1} & = & \{E[\texttt{in-pure}\ E'] \mid E \in D^n\} \end{array}$$

For all thread identifiers i and integers n, we partition the set of well-formed states into four categories:

- R_i^n – states where thread i is in the right-mover part of an `in-pure` block at depth n;
- L_i^n – states where thread i is in the left-mover part of an `in-pure` block at depth n;
- U_i^n – states where thread i is in a `in-pure` block at depth n that has normal atomicity $\perp$, and hence diverges; and
- N_i^n – states where thread i does not contain `in-pure` block at depth n.

In addition, N^n contains states where no thread is in an `in-pure` block at depth n (and so N^1 describes states containing no `in-pure` blocks).

$$\begin{array}{rcl}
R_i^n & = & WT \cap \{\langle H, M, T\rangle \mid T_i \equiv E[\texttt{in-pure}\ e] \\
& & \qquad \wedge\ E \in D^{n-1}\ \wedge\ \alpha_n(e) \in \{R, A, \top\}\} \\
L_i^n & = & WT \cap \{\langle H, M, T\rangle \mid T_i \equiv E[\texttt{in-pure}\ e] \\
& & \qquad \wedge\ E \in D^{n-1}\ \wedge\ \alpha_n(e) \in \{B, L\}\} \\
U_i^n & = & WT \cap \{\langle H, M, T\rangle \mid T_i \equiv E[\texttt{in-pure}\ e] \\
& & \qquad \wedge\ E \in D^{n-1}\ \wedge\ \alpha_n(e) = \perp\} \\
N_i^n & = & WT \cap \{\langle H, M, T\rangle \mid \forall E \in D^{n-1}.\ T_i \not\equiv E[\texttt{in-pure}\ e]\} \\
N^n & = & \forall i.\ N_i^n
\end{array}$$

We next define two restrictions of the transition relation $\rightarrow_i$. The definition of these relations requires some additional notation. For any state predicate $X \subseteq \mathit{State}$ and transition relation $Y \subseteq \mathit{State} \times \mathit{State}$, by X/Y we mean the transition relation obtained by restricting Y to pairs whose first component is in X. Similarly, by $Y\backslash X$ we mean the restriction of Y to pairs whose second component is in X. The first restricted relation $\stackrel{n}{\Longrightarrow}_i$ (for $n \geq 0$) is a restriction of $\rightarrow_i$ to states where thread i has at most n nested `in-pure` blocks (at nesting depth 1 through n inclusive). The second relation $\stackrel{n}{\models\!\Longrightarrow}_i$ (for $n \geq 1$) is a serialized version of $\stackrel{n}{\Longrightarrow}_i$ where each `in-pure` block at depth n is executed serially.

$$\begin{array}{rcl}
\stackrel{n}{\Longrightarrow}_i & = & N^{n+1}/\rightarrow_i \\
\stackrel{n}{\models\!\Longrightarrow}_i & = & (N^{n+1} \wedge \forall j \neq i.\ N_j^n)/\rightarrow_i \\
\stackrel{n}{\Longrightarrow} & = & \exists i.\ \stackrel{n}{\Longrightarrow}_i \\
\stackrel{n}{\models\!\Longrightarrow} & = & \exists i.\ \stackrel{n}{\models\!\Longrightarrow}_i
\end{array}$$

The following properties are obvious consequences of these definitions.

$$\begin{array}{ccccc}
\stackrel{n}{\models\!\Longrightarrow}_i & \subset & \stackrel{n}{\Longrightarrow}_i & \subset & \rightarrow_i \\
\stackrel{n}{\models\!\Longrightarrow} & \subset & \stackrel{n}{\Longrightarrow} & \subset & \rightarrow
\end{array}$$

Using these relations, we show that any non-serial execution via $\stackrel{n}{\Longrightarrow}$ can be reduced to an equivalent serial execution via $\stackrel{n}{\models\!\Longrightarrow}$.

LEMMA 4 (PURE REDUCTION). *Let P be a well-formed program with correct conflict and call annotations. Suppose $N^n(P)$ and $N^n(Q)$ and $P \stackrel{n}{\Longrightarrow}^* Q$, where $n \geq 1$. Then $P \stackrel{n}{\models\!\Longrightarrow}^* Q$.*

The proof of this lemma and of other lemmas and theorems in this section is contained in the appendix.

We next prove that given any execution via $\stackrel{n}{\models\!\Longrightarrow}$ (in which `in-pure` blocks at depth n are executed serially), we can elide each of these serially-executed `in-pure` blocks at depth n, thus reducing the nesting depth of `in-pure` blocks and yielding an execution sequence under $\stackrel{n-1}{\Longrightarrow}$.

LEMMA 5 (PURE ELIDING). *If the pure annotations in P are correct and $P_0 \stackrel{n}{\models\!\Longrightarrow}^* P_k$ and $N^n(P_0)$ and $N^n(P_k)$, where $n \geq 1$, then $P_0 \stackrel{n-1}{\Longrightarrow}^* P_k$.*

Thus, given any execution sequence $P \stackrel{n}{\Longrightarrow}^* Q$, by repeatedly applying the Pure Reduction and Pure Eliding Lemmas, we eventually obtain an execution sequence $P \stackrel{0}{\Longrightarrow}^* Q$, where each intermediate state satisfies N^0 and thus does not contain `in-pure` blocks. We now leverage reduction one more time to reduce each atomic block in the execution sequence $P \stackrel{0}{\Longrightarrow}^* Q$. To assist in performing this reduction, we again partition the set of well-formed states into four categories:

- R_i – states where thread i is executing the right-mover part of an `in-atomic` block;
- L_i – states where thread i is executing the left-mover part of an `in-atomic` block;
- U_i – states where thread i is executing an `in-atomic` block that will never terminate; and
- N_i – states where thread i does not executing an `in-atomic` blocks.

We define N to contain states where no thread is executing an `in-atomic` block.

$$\begin{array}{rcl}
R_i & = & WT \cap \{\langle H, M, T\rangle \mid T_i \equiv E[\texttt{in-atomic}\ e]\ \wedge\ \alpha(e) \not\sqsubseteq L\} \\
L_i & = & WT \cap \{\langle H, M, T\rangle \mid T_i \equiv E[\texttt{in-atomic}\ e] \\
& & \qquad \wedge\ B \sqsubseteq \alpha(e) \sqsubseteq L\} \\
U_i & = & WT \cap \{\langle H, M, T\rangle \mid T_i \equiv E[\texttt{in-atomic}\ e]\ \wedge\ \alpha(e) = \perp\} \\
N_i & = & WT \cap \{\langle H, M, T\rangle \mid |T| < i\ \vee\ T_i \not\equiv E[\texttt{in-atomic}\ e]\} \\
N & = & \forall i.\ N_i
\end{array}$$

The following theorem then shows that we can reduce an execution $P \stackrel{0}{\Longrightarrow}^* Q$ to a serial execution $P \mapsto^* Q$.

LEMMA 6 (ATOMIC REDUCTION). *Let P be a program with correct conflict and call annotations and let Γ be a valid effect environment such that $\Gamma \vdash P$. Suppose $N(P)$ and $N(Q)$ and $P \stackrel{0}{\Longrightarrow}^* Q$. Then $P \mapsto^* Q$.*

The following Correctness Theorem states that given any non-serial execution $P \rightarrow^* Q$ with arbitrarily-nested `in-pure` and in-atomic blocks, we can always obtain an equivalent serial execution $P \mapsto^* Q$. To avoid the consideration of partially-executed atomic blocks, we require that no thread is executing an atomic block in either the initial state P or the final state Q. We conjecture that our results can be extended to cover some partially-executed blocks by following existing techniques [17].

THEOREM 4 (CORRECTNESS). *Let P be a program with correct annotations and let Γ be a valid effect environment such that $\Gamma \vdash P$. Suppose $N(P)$ and $N(Q)$ and $P \rightarrow^* Q$. Then $P \mapsto^* Q$.*

PROOF. This theorem follows from the following induction hypothesis, since there exists n such that $P \stackrel{n}{\Longrightarrow}^* Q$.

> *Induction Hypothesis:* Let P be a program with correct annotations and let Γ be a valid effect environment such

that $\Gamma \vdash P$. Suppose $N(P)$ and $N(Q)$ and $P \overset{n}{\Longrightarrow}^* Q$. Then $P \mapsto^* Q$.

The proof of this hypothesis is by induction on n. If $n = 0$, the conclusion follows by Lemma 6. Otherwise, if $P \overset{n}{\Longrightarrow}^* Q$, then

$$\begin{array}{ll} P \overset{n}{\models\!\Longrightarrow}^* Q & \text{by Lemma 4} \\ P \overset{n-1}{\Longrightarrow}^* Q & \text{by Lemma 5} \\ P \to^* Q & \text{by the inductive hypothesis} \end{array}$$

□

D Detailed Proofs

We first formalize the notion of commuting transitions as follows. The composition $Y \circ Z$ of two transition relations is the set

$$\{(p, r) \mid \exists q.\ (p, q) \in Y \text{ and } (q, r) \in Z\}$$

A transition relation Y *right-commutes* with a transition relation Z if $Y \circ Z \subseteq Z \circ Y$, and Y *left-commutes* with Z if $Z \circ Y \subseteq Y \circ Z$.

LEMMA 7. *If $H, M, T \ \to_i^* \ H', M', T'$ by any rule except [RED FORK] then T, T' agree except on i.*

PROOF. By definition. □

LEMMA 8 (CONTEXT).

1. *If $H, M, e \ \to^* \ H', M', e'$ then for all T, T', E we have $H, M, T.E[e].T' \to_i^* H', M', T, E[e'].T'$ where $i = |T| + 1$.*

2. *If $H, M, T.E[e].T' \ \to_i \ H', M', T, E[e'].T'$ where $i = |T| + 1$ then $H, M, e \to H', M', e'$.*

PROOF. By case analysis on the reduction rule in the hypothesis of each part of the lemma. □

LEMMA 9 (UNIQUE EFFECTS). *If $\Gamma \vdash P$ and P contains e then there exists unique a, b such that $\Gamma \vdash e : a \uparrow b$.*

LEMMA 10 (BREAKING). *If $\Gamma \vdash E[e] : a \uparrow b$ and $\Gamma \vdash e : a' \uparrow b'$ and E does not contain* `block` [ ] *or* `in-pure` [ ] *then $b' \sqsubseteq b$.*

PROOF. By induction on the structure of E. □

LEMMA 11 (CONTEXT REPLACEMENT). *If $\Gamma \vdash E[e_1] : a \uparrow b$ and $\Gamma \vdash e_1 : a_1 \uparrow b_1$ and $\Gamma \vdash e_2 : a_2 \uparrow b_2$ and $a_2 \sqsubseteq a_1$ and $b_2 \sqsubseteq b_1$, then $\Gamma \vdash E[e_2] : a' \uparrow b'$ where $a' \sqsubseteq a$ and $b' \sqsubseteq b$.*

PROOF. By induction on the structure of E. □

LEMMA 12 (ATOMICITY MONOTONICITY). *If $H, M, T \ \to_i \ H', M', T'$ and $\Gamma \vdash H, M, T$ and $T_i = E[e]$ and $T_i' = E[e']$ then $\alpha(e') \sqsubseteq \alpha(e)$ and $\alpha_n(e') \sqsubseteq \alpha_n(e)$ and $\alpha_x(e') \sqsubseteq \alpha_x(e)$.*

PROOF. By case analysis on $\to_i$ and Lemma 11. □

RESTATEMENT OF THEOREM 3 (SUBJECT REDUCTION). *If the call annotations in P are correct and $\Gamma \vdash P$ and $P \to P'$, then $\Gamma \vdash P'$.*

PROOF. Proof is by induction over execution sequences, using the previous three lemmas. □

LEMMA 13. *Suppose $P \to_i H', M', T'$ and $N^{n+1}(P)$ and $n \geq 1$ and*

$$\begin{array}{rcl} P &=& H, M, T \\ T_i &=& E[\texttt{in-pure}\ e] \\ E &\in& D^{n-1} \\ T_i' &=& E'[\texttt{in-pure}\ e'] \\ E' &\in& D^{n-1} \end{array}$$

Then $E = E'$.

LEMMA 14 (SEQUENTIALITY). *If E does not contain* `in-pure` *and $\Gamma \vdash e : A \uparrow \bot$ and $\alpha(E[e])$ is defined and $\alpha(E[v])$ is defined then*

$$\begin{array}{rcl} \alpha(E[e]) &=& A; \alpha(E[v]) \\ \alpha_n(E[e]) &=& A; \alpha_n(E[v]) \\ \alpha_x(E[e]) &=& A; \alpha_x(E[v]) \end{array}$$

PROOF. The first consequence follows directly from the last two. We prove the last two consequences by simultaneous structural induction on E. We show the proof for a few cases; the rest are similar.

Case 1 [ ]

$$\begin{array}{rl} & \alpha_n(([\,])[e]) \\ = & \alpha_n(e) \\ = & A \\ = & A; B \\ = & A; \alpha_n(v) \\ = & A; \alpha_n(([\,])[v]) \end{array}$$

and similarly for $\alpha_x(([\,])[e])$.

Case 2 $x_r := E$

$$\begin{array}{rll} & \alpha_n((x_r := E)[e]) & \\ = & \alpha_n(x_r := E[e]) & \\ = & \alpha_n(E[e]); a & \text{by [EXP ASSIGN],} \\ & & \text{for } a \in \{B, A\} \\ = & A; \alpha_n(E[v]); a & \text{by ind. hyp.} \\ = & A; \alpha_n(x_r := E[v]) & \text{by [EXP ASSIGN]} \\ = & A; \alpha_n((x_r := E)[v]) & \end{array}$$

$$\begin{array}{rll} & \alpha_x((x_r := E)[e]) & \\ = & \alpha_x(x_r := E[e]) & \\ = & \alpha_x(E[e]) & \text{by [EXP ASSIGN],} \\ = & A; \alpha_x(E[v]) & \text{by ind. hyp.} \\ = & A; \alpha_x(x_r := E[v]) & \text{by [EXP ASSIGN]} \\ = & A; \alpha_x((x_r := E)[v]) & \end{array}$$

Case 3 $\texttt{if}\ E\ e_1\ e_2$

$$\begin{array}{rll} & \alpha_n((\texttt{if}\ E\ e_1\ e_2)[e]) & \\ = & \alpha_n(\texttt{if}\ E[e]\ e_1\ e_2) & \\ = & \alpha_n(E[e]); (a_1 \sqcup a_2) & \text{by [EXP IF],} \\ = & A; \alpha_n(E[v]); (a_1 \sqcup a_2) & \text{by ind. hyp.} \\ = & A; \alpha_n(\texttt{if}\ E[v]\ e_1\ e_2) & \text{by [EXP IF]} \\ = & A; \alpha_n((\texttt{if}\ E\ e_1\ e_2)[v]) & \end{array}$$

$$\begin{array}{rll} & \alpha_x((\texttt{if}\ E\ e_1\ e_2)[e]) & \\ = & \alpha_x(\texttt{if}\ E[e]\ e_1\ e_2) & \\ = & \alpha_x(E[e]) \sqcup (\alpha_n(E[e]); (b_1 \sqcup b_2)) & \text{by [EXP IF],} \\ & & \text{for } b_i \in \mathit{Atomicity} \\ = & (A; \alpha_x(E[v])) \sqcup (A; \alpha_n(E[v]); (b_1 \sqcup b_2)) & \text{by ind. hyp.} \\ = & A; (\alpha_x(E[v]) \sqcup (\alpha_n(E[v]); (b_1 \sqcup b_2)) & \sqcup \text{ distributes over ;} \\ = & A; \alpha_x(\texttt{if}\ E[v]\ e_1\ e_2) & \text{by [EXP IF]} \\ = & A; \alpha_x((\texttt{if}\ E\ e_1\ e_2)[v]) & \end{array}$$

$\square$

RESTATEMENT OF LEMMA 5 (PURE ELIDING). *If the pure annotations in P are correct and $P_0 \overset{n}{\Longmapsto}{}^* P_k$ and $N^n(P_0)$ and $N^n(P_k)$, where $n \geq 1$, then $P_0 \overset{n-1}{\Longrightarrow}{}^* P_k$.*

PROOF. Proof is by contradiction. Suppose that the lemma is false and let $P_0, \ldots, P_k$ be the shortest sequence violating this lemma. Then $P_1, \ldots, P_{k-1}$ must not be in N^n. Otherwise there would be a shorter sequence violating this lemma.

Thus, $N^n(P_0)$, $\neg N^n(P_1)$, $N^n(P_k)$, $\neg N^n(P_{k-1})$, and inspection of the evaluation rules indicates that $\exists i, E \in D^n$ such that

$$\begin{array}{rcl} P_0 & = & H, M, T[i := E[\texttt{pure}\ e]] \\ P_1 & = & H, M, T[i := E[\texttt{in-pure}\ e]] \\ P_{k-1} & = & H', M', T'[i := E[\texttt{in-pure}\ v]] \\ P_k & = & H', M', T'[i := E[v]] \end{array}$$

Note that thread i contains the evaluation context $E[\texttt{in-pure}\ [\ ]]$ for each state $P_1, \ldots, P_{k-1}$. Since [RED FORK] is not applicable inside `in-pure` blocks, Lemma 7 indicates that $T = T'$. By Lemma 8, $H, M, e \rightarrow^* H', M', v$. Since e is pure, $H, M, v \rightarrow H', M', v$. Thus,

$$\begin{array}{rll} P_0 & = & H, M, T[i := E[\texttt{pure}\ e]] \\ & \overset{n-1}{\Longrightarrow}_i & H, M, T[i := E[v]] \quad \text{by rule [ABS PURE SKIP]} \\ & \overset{n-1}{\Longrightarrow}_i & H', M', T[i := E[v]] \text{ by rule [ABS GARBAGE]} \\ & = & P_k \end{array}$$

Therefore $P_0 \overset{n-1}{\Longrightarrow}{}^* P_k$. $\square$

RESTATEMENT OF LEMMA 4 (PURE REDUCTION). *Let P be a well-formed program with correct conflict and call annotations. Suppose $N^n(P)$ and $N^n(Q)$ and $P \overset{n}{\Longrightarrow}{}^* Q$, where $n \geq 1$. Then $P \overset{n}{\Longmapsto}{}^* Q$.*

PROOF. Lemma 16 shows that for we can use Lemma 15 to obtain the desired result by substituting the set R_i^n for $\mathcal{R}_i$, the set L_i^n for $\mathcal{L}_i$, the set U_i^n for $\mathcal{U}_i$, the set N_i^n for $\mathcal{N}_i$, the relation $\overset{n}{\Longrightarrow}_i$ for $\rightharpoonup_i$, the relation $\overset{n}{\Longmapsto}$ for $\vdash\!\rightharpoonup$, the state P for p, and the state Q for q. $\square$

Our proof of the above theorem depends on the following theorem, which is inspired by the reduction theorem of Cohen and Lamport [6].

LEMMA 15 (REDUCTION). *For all i, let $\mathcal{R}_i$, $\mathcal{L}_i$, $\mathcal{U}_i$, and $\mathcal{N}_i$ be sets of states, and $\rightharpoonup_i$ be a transition relation. Suppose for all thread indices i and j such that $j \neq i$,*

A1. $\mathcal{R}_i$, $\mathcal{L}_i$, $\mathcal{U}_i$, and $\mathcal{N}_i$ are pairwise disjoint and their union yields the entire state space,

A2. $(\mathcal{L}_i / \rightharpoonup_i \backslash \mathcal{R}_i)$ is false,

A3. $(\rightharpoonup_i \backslash \mathcal{R}_i)$ right-commutes with $\rightharpoonup_j$,

A4. $(\mathcal{L}_i / \rightharpoonup_i)$ left-commutes with $\rightharpoonup_j$,

A5. if $p \rightharpoonup_i q$, then $\mathcal{R}_j(p) \Leftrightarrow \mathcal{R}_j(q)$, $\mathcal{L}_j(p) \Leftrightarrow \mathcal{L}_j(q)$, $\mathcal{U}_j(p) \Leftrightarrow \mathcal{U}_j(q)$, and $\mathcal{N}_j(p) \Leftrightarrow \mathcal{N}_j(q)$,

A6. if $p \rightharpoonup_i q$ and $\mathcal{U}_i(p)$, then $\mathcal{U}_i(q)$.

Let

$$\begin{array}{rcl} \mathcal{N} & = & \forall i.\, \mathcal{N}_i \\ \rightharpoonup & = & \exists i.\, \rightharpoonup_i \\ \vdash\!\rightharpoonup & = & \exists i.\, (\forall j \neq i.\, \mathcal{N}_j) / \rightharpoonup_i \end{array}$$

Suppose $p \in \mathcal{N}$ and $q \in \mathcal{N}$ and $p \rightharpoonup^* q$. Then $p \vdash\!\rightharpoonup^* q$.

PROOF. Follows by a similar argument to Theorem 1 of [17]. $\square$

LEMMA 16 (PURE REDUCTION HELPER). *Let P be a program with correct conflict and call annotations and let Γ be a valid environment such that $\Gamma \vdash P$. Suppose $n \geq 1$ and $N^n(P)$ and $N^n(Q)$ and $P \overset{n}{\Longrightarrow}{}^* Q$. Then for all thread indices i and j such that $j \neq i$,*

1. *R_i^n, L_i^n, U_i^n, and N_i^n are pairwise disjoint and their union yields the entire state space,*

2. *$(L_i^n / \overset{n}{\Longrightarrow} \backslash R_i^n)$ is false,*

3. *$(\overset{n}{\Longrightarrow}_i \backslash R_i^n)$ right-commutes with $\overset{n}{\Longrightarrow}_j$,*

4. *$(L_i^n / \overset{n}{\Longrightarrow}_i)$ left-commutes with $\overset{n}{\Longrightarrow}_j$,*

5. *if $P \overset{n}{\Longrightarrow}_i Q$, then $R_j^n(P) \Leftrightarrow R_j^n(Q)$, $L_j^n(P) \Leftrightarrow L_j^n(Q)$, and $W_j(P) \Leftrightarrow W_j(Q)$.*

6. *if $P \overset{n}{\Longrightarrow}_i Q$ and $U_i^n(P)$ then $U_i^n(Q)$.*

PROOF.

1. By the definition of R_i^n, L_i^n, U_i^n, and N_i^n.

2. Suppose $L_i^n / \overset{n}{\Longrightarrow} \backslash R_i^n$ is not false. Then there exist P and Q such that $P \rightarrow_i Q$ and $N^{n+1}(P)$ and $L_i^n(P)$ and $R_i^n(Q)$. The following must therefore hold:

$$\begin{array}{rcl} P & = & H, M, T \\ Q & = & H', M', T' \\ T_i & = & E[\texttt{in-pure}\ e] \\ E & \in & D^{n-1} \\ \alpha_n(e) & \sqsubseteq & L \\ T'_i & = & E'[\texttt{in-pure}\ e'] \\ E' & \in & D^{n-1} \\ \alpha_n(e') & \not\sqsubseteq & L \end{array}$$

   Then $E' = E$ by Lemma 13, and $\alpha_n(e') \sqsubseteq \alpha_n(e) \sqsubseteq L$ by Lemma 12, which contradicts the assumption that $\alpha_n(e') \not\sqsubseteq L$.

3. Suppose $P_1 \overset{n}{\Longrightarrow}_i P_2 \overset{n}{\Longrightarrow}_j P_3$ where $i \neq j$ and $P_2 \in R_i^n$. Then $P_1 \rightarrow_i P_2 \rightarrow_j P_3$ and $N^{n+1}(P_1)$ and $N^{n+1}(P_2)$. We proceed by case analysis on the rule for $P_1 \rightarrow_i P_2$, and show that in each case that there exists P_4 such that $P_1 \overset{n}{\Longrightarrow}_j P_4 \overset{n}{\Longrightarrow}_i P_3$.

   (a) [RED READ] We have $P_1 = H, M, T$ such that $T_i \equiv E[x_r]$, and $P_2 = H, M, T[i \mapsto E[H(x)]]$.

   First consider the case where x is unstable. If $P_1 \rightarrow_i P_2$ by rule [RED READ], then $P_1 \rightarrow_i P_2$ by rule [ABS UNSTABLE READ]. Since [ABS UNSTABLE READ] does not access or update the heap or synchronization heap, $\to_i$ right-commutes with $\to_j$ in this case.

Suppose x is a stable variable and $r = \epsilon$. Since the conflict tag is correct and the other thread j is not about to read or write x, so $\to_i$ right-commutes with $\to_j$.

Finally consider the case where x is a stable variable and $r = \bullet$. Then there exists E', E'' such that

$$\begin{array}{rcll}
E'[\texttt{in-pure } E''] & \equiv & E & \text{since } P_2 \in R_i^n \\
E' & \in & D^{n-1} & \text{since } P_2 \in R_i^n \\
\alpha_n(E''[H(x)]) & \in & \{B, L\} & \text{since } P_2 \in R_i^n \\
\alpha_n(E''[H(x)]) & \sqsubseteq & A & [\text{EXP INPURE}] \\
\alpha_n(E''[H(x)]) & \in & \{R, A\} & \text{from prev 2 stmts} \\
\alpha_n(x_\bullet) & = & A & \text{from [EXP READ RACE]} \\
\alpha_n(E''[x_\bullet]) & = & \top & \text{from Lemma 14} \\
 & & & \text{and prev 2 stmts.} \\
\alpha_n(E''[x_\bullet]) & \sqsubseteq & A & \text{from [EXP INPURE]}
\end{array}$$

The last two statements give a contradiction, so this case is impossible.

(b) [RED WRITE] or [ABS UNSTABLE WRITE] First consider the case where x is unstable. If $P_2 \to_j P_3$ does not access x, then it is trivial to show that $\to_i$ right-commutes with $\to_j$. If $P_2 \to_j P_3$ does access x, there are three cases to consider for the rule used to perform this transition:

i. [RED WRITE] or [ABS UNSTABLE WRITE] We have:

$$\begin{array}{rl}
P_1 = & H, M, T \\
\to_i & H[x := v], M, T[i \mapsto E_i[v']] \\
\to_j & H[x := w], M, T[i \mapsto E_i[v']][j \mapsto E_j[w']] \\
= & P_3
\end{array}$$

Using two applications of the rule [ABS UNSTABLE WRITE], we may obtain the following:

$$\begin{array}{rl}
P_1 = & H, M, T \\
\to_j & H[x := v], M, T[j \mapsto E_j[w]] \\
\to_i & H[x := w], M, T[j \mapsto E_j[w']][i \mapsto E_i[v']] \\
= & P_3
\end{array}$$

Thus, the two steps commute.

ii. [ABS GARBAGE] Similar to the previous case.

iii. [RED READ] We have:

$$\begin{array}{rl}
P_1 = & H, M, T \\
\to_i & H[x := v], M, T[i \mapsto E_i[v']] \\
\to_j & H[x := v], M, T[i \mapsto E_i[v']][j \mapsto E_j[v]] \\
= & P_3
\end{array}$$

Using an application of [ABS UNSTABLE READ] followed by [ABS UNSTABLE WRITE], we may obtain the following:

$$\begin{array}{rl}
P_1 = & H, M, T \\
\to_j & H, M, T[j \mapsto E_j[v]] \\
\to_i & H[x := v], M, T[j \mapsto E_j[v]][i \mapsto E_i[v']] \\
= & P_3
\end{array}$$

Thus, the two steps commute.

Suppose x is a stable variable. We have $P_1 = H, M, T$ such that $T_i \equiv E_i[x_r = v]$, and $P_2 = H[x := v], M, T[i \mapsto E_i[v']]$.

Further suppose that $r = \epsilon$. Since the conflict tag is correct and the other thread j is not about to read or write x, so $\to_i$ right-commutes with $\to_j$.

Finally consider the case where x is a stable variable and $r = \bullet$. Then there exists E', E'' such that

$$\begin{array}{rcll}
E'[\texttt{in-pure } E''] & \equiv & E & \text{since } P_2 \in R_i^n \\
E' & \in & D^{n-1} & \text{since } P_2 \in R_i^n \\
\alpha_n(E''[v]) & \in & \{B, L\} & \text{since } P_2 \in R_i^n \\
\alpha_n(E''[v]) & \sqsubseteq & A & [\text{EXP INPURE}] \\
\alpha_n(E''[v]) & \in & \{R, A\} & \text{from prev 2 stmts} \\
\alpha_n(x_\bullet = v) & = & A & \text{from [EXP ASSIGN RACE]} \\
\alpha_n(E''[x_\bullet = v]) & = & \top & \text{from Lemma 14} \\
 & & & \text{and prev 2 stmts.} \\
\alpha_n(E''[x_\bullet = v]) & \sqsubseteq & A & \text{from [EXP INPURE]}
\end{array}$$

The last two statements give a contradiction, so this case is impossible.

(c) [RED PRIM] In this case $P_1 = H, M, T$ where $T_i \equiv E[\texttt{in-pure } E'[p(\overline{v})]]$ and $E \in D^{n-1}$. Again, by similar reasoning, $\alpha_n(p(\overline{v})) \sqsubseteq R$, and hence $\Gamma(p) \sqsubseteq R$. Hence, the application of p right-commutes with any primitive application of thread j (which are the only steps that modify the synchronization heap), so $\to_i$ right-commutes with $\to_j$.

(d) [RED FORK] In this case $P_1 = H, M, T$ such that $T_i \equiv E[\texttt{in-pure } E'[\texttt{fork } e]]$ and $P_2 = H, M, T[i \mapsto E[\texttt{in-pure } E'[0]]$. From the subject reduction theorem, we know that $\alpha(E'[\texttt{fork } e]) \sqsubseteq A$. From Lemma 14, we know that

$$\alpha(E'[\texttt{fork}\, e]) = \alpha(\texttt{fork}\, e); \alpha(E'[0]) = A; \alpha(E'[0]).$$

Therefore, we must have $\alpha(E'[0]) \sqsubseteq L$. Therefore $P_2 \in L_i$ and we have a contradiction. That is, it is impossible to perform a fork operation and remain in the right-mover part of a pure block.

(e) If $P_1 \to_i P_2$ by [RED LET], [RED CALL], or [ABS GARBAGE] then only unstable or unallocated variables in the heap or synchronization heap are accessed. Previously unallocated variables could not be accessed by thread j, and we may use an argument similar to the case for [ABS UNSTABLE WRITE] to show that commutativity is not affected by updates to unstable variables. Thus, $\to_i$ right-commutes with $\to_j$.

(f) If $P_1 \to_i P_2$ by any of [RED WRONG], [RED IF1], [RED IF2], [RED LOOP], [RED BREAK], [RED BLOCK], [RED ATOMIC], [RED IN-ATOMIC], [RED PURE], [RED PURE EX], [RED IN-PURE], [ABS PURE SKIP], or [ABS UNSTABLE READ], then these rules do not access the heap or the synchronization heap, and so $\to_i$ right-commutes with $\to_j$.

4. Suppose $P_1 \stackrel{n}{\Longrightarrow}_j P_2 \stackrel{n}{\Longrightarrow}_i P_3$ where $i \neq j$ and $P_2 \in L_i^n$.

Then $P_1 \rightarrow_j P_2 \rightarrow_i P_3$ and $N^{n+1}(P_1)$ and $N^{n+1}(P_2)$. We proceed by case analysis on the rule for $P_2 \rightarrow_i P_3$.

(a) [RED READ] In this case $P_2 = H, M, T$ such that $T_i \equiv E[x_r]$ and $P_3 = H, M, T[i \mapsto E[H(x)]$.

First consider the case where x is unstable. If $P_2 \rightarrow_i P_3$ by rule [RED READ], then $P_2 \rightarrow_i P_3$ by rule [ABS UNSTABLE READ]. Since [ABS UNSTABLE READ] does not access the heap or synchronization heap, $\rightarrow_i$ right-commutes with $\rightarrow_j$ in this case.

Suppose x is a stable variable and $r = \epsilon$. Since the conflict tag is correct and the other thread j is not about to read or write x, so $\rightarrow_i$ right-commutes with $\rightarrow_j$.

Finally consider the case where x is a stable variable and $r = \bullet$. Then there exists E', E'' such that

$E'[\texttt{in-pure}\ E'']$	$\equiv$	E	since $P_2 \in L_i^n$
E'	$\in$	D^{n-1}	since $P_2 \in L_i^n$
$\alpha_n(E''[x_\bullet])$	$\in$	$\{B, L\}$	since $P_2 \in L_i^n$
$\alpha_n(E''[H(x)])$	$\in$	$\{\bot, B, L\}$	by Lemma 12
$\alpha_n(x_\bullet)$	$=$	A	from [EXP READ RACE]
$\alpha_n(E''[x_\bullet])$	$\in$	$\{\bot, A\}$	from Lemma 14 and prev 2 stmts.

The third and last statements give a contradiction, so this case is impossible.

(b) [RED WRITE] or [ABS UNSTABLE WRITE] First consider the case where x is unstable. If $P_1 \rightarrow_j P_2$ does not access x, then it is trivial to show that $\rightarrow_i$ right-commutes with $\rightarrow_j$. If $P_1 \rightarrow_j P_2$ does access x, there are three cases to consider for the rule used to perform this transition:

i. [RED WRITE] or [ABS UNSTABLE WRITE] We have:

$$\begin{array}{lll} P_1 & = & H, M, T \\ & \rightarrow_j & H[x := v], M, T[j \mapsto E_j[v']] \\ & \rightarrow_i & H[x := w], M, T[j \mapsto E_j[v']][i \mapsto E_i[w']] \\ & = & P_3 \end{array}$$

Using two applications of the rule [ABS UNSTABLE WRITE], we may obtain the following:

$$\begin{array}{lll} P_1 & = & H, M, T \\ & \rightarrow_i & H[x := w], M, T[i \mapsto E_i[w']] \\ & \rightarrow_j & H[x := w], M, T[i \mapsto E_i[w']][j \mapsto E_j[v']] \\ & = & P_3 \end{array}$$

Thus, the two steps commute.

ii. [ABS GARBAGE] Similar to the previous case.

iii. [RED READ] We have:

$$\begin{array}{lll} P_1 & = & H, M, T \\ & \rightarrow_j & H, M, T[j \mapsto E_j[H(x)]] \\ & \rightarrow_i & H[x := v'], M, T[j \mapsto E_j[H(x)]][i \mapsto E_i[v]] \\ & = & P_3 \end{array}$$

Using an application of [ABS UNSTABLE WRITE] followed by [ABS UNSTABLE READ], we may obtain the following:

$$\begin{array}{lll} P_1 & = & H, M, T \\ & \rightarrow_i & H[x := v'], M, T[i \mapsto E_i[v]] \\ & \rightarrow_j & H[x := v'], M, T[i \mapsto E_i[v]][j \mapsto E_j[H(x)]] \\ & = & P_3 \end{array}$$

Thus, the two steps commute.

Now consider the case that x is stable. We have $P_2 = H, M, T$ such that $T_i \equiv E[x_r = v]$, and $P_3 = H[x := v], M, T[i \mapsto E_i[v']]$.

Suppose x is a stable variable and $r = \epsilon$. Since the conflict tag is correct and the other thread j is not about to read or write x, so $\rightarrow_i$ right-commutes with $\rightarrow_j$.

Finally consider the case where x is a stable variable and $r = \bullet$. Then there exists E', E'' such that

$E'[\texttt{in-pure}\ E'']$	$\equiv$	E	since $P_2 \in L_i^n$
E'	$\in$	D^{n-1}	since $P_2 \in L_i^n$
$\alpha_n(E''[x_\bullet = v])$	$\in$	$\{B, L\}$	since $P_2 \in L_i^n$
$\alpha_n(E''[v])$	$\in$	$\{\bot, B, L\}$	by Lemma 12
$\alpha_n(x_\bullet = v)$	$=$	A	from [EXP ASSIGN RACE]
$\alpha_n(E''[x_\bullet = v])$	$\in$	$\{\bot, A\}$	from Lemma 14 and prev 2 stmts.

The third and last statements give a contradiction, so this case is impossible.

(c) [RED PRIM] In this case $P_2 = H, M, T$ where $T_i \equiv E[\texttt{in-pure}\ E'[p(\overline{v})]]$ and $E \in D^{n-1}$. Again, by similar reasoning, $\alpha_n(p(\overline{v})) \sqsubseteq L$, and hence $\Gamma(p) \sqsubseteq L$. Hence, the application of p left-commutes with any primitive application of thread j (which are the only steps that modify the synchronization heap), so $\rightarrow_i$ left-commutes with $\rightarrow_j$.

(d) [RED FORK] In this case $P_2 = H, M, T$ such that $T_i \equiv E[\texttt{in-pure}\ E'[\texttt{fork}\ e]]$ and $P_3 = H, M, T[i \mapsto E[\texttt{in-pure}\ E'[0]]$. From the subject reduction theorem, we know that $\alpha(E'[\texttt{fork}\ e]) \sqsubseteq A$. From Lemma 14, we know that

$$\alpha(E'[\texttt{fork}\,e]) = \alpha(\texttt{fork}\,e); \alpha(E'[0]) = A; \alpha(E'[0]).$$

Therefore $\alpha(E'[\texttt{fork}\ e]) \not\sqsubseteq L$. Therefore $P_2 \notin L_i$ and we have a contradiction. That is, it is impossible to perform a fork operation from the left-mover part of an atomic block.

(e) If $P_2 \rightarrow_i P_3$ by any of [RED LET], [RED CALL], or [ABS GARBAGE], then only unstable variables and locations not previously allocated in the synchronization heap and heap are accessed. Previously unallocated variables could not be accessed by thread j, and we may use an argument similar to the case for [ABS UNSTABLE WRITE] to show that commutativity is not affected by updates to unstable variables. Thus, $\rightarrow_i$ left-commutes with $\rightarrow_j$.

(f) If $P_2 \rightarrow_i P_3$ by any of [RED WRONG], [RED IF1], [RED IF2], [RED LOOP], [RED BREAK], [RED BLOCK], [RED ATOMIC], [RED IN-ATOMIC], [RED PURE], [RED PURE EX], [RED IN-PURE], [ABS PURE SKIP], or [ABS UNSTABLE READ], then these rules do not access the heap or the synchronization heap, and so $\rightarrow_i$ left-commutes with $\rightarrow_j$.

5. Suppose $P \rightarrow_i Q$. If this step is not a fork step, then threads other than i do not change in going from P to Q. Therefore $R_j^n(P)$ iff $R_j^n(Q)$, $L_j^n(P)$ iff $L_j^n(Q)$, $U_j^n(P)$ iff $U_j^n(Q)$, and $W_j^n(P)$ iff $W_j^n(Q)$.

   If this step forks a new thread j, then $P = H, M, T$ and $T_i = E[\texttt{fork}\ e]$ where e does not contain `in-pure`, since it does not occur in an evaluation context position. Let j be the newly forked thread. Then $N_j^n(P)$ (by the definition of N_j^n) and $N_j^n(Q)$. Furthermore, since threads other than i or j do not change in going from P to Q, we have that $R_j^n(P)$ iff $R_j^n(Q)$, $L_j^n(P)$ iff $L_j^n(Q)$, $U_j^n(P)$ iff $U_j^n(Q)$, and $W_j^n(P)$ iff $W_j^n(Q)$.

6. Suppose $P \rightarrow_i Q$ and $P \in U_i^n$. Then $P = H, M, T$ and $T_i = E[\texttt{in-pure}\ e]$ and $E \in D^{n-1}$ and $\alpha_n(e) = \bot$. Hence $Q = H', M', T'$ and $T_i' = E[\texttt{in-pure}\ e']$ and, by Lemma 12, $\alpha_n(e') = \bot$ and so $Q \in U_i^n$.

□

RESTATEMENT OF LEMMA 6 (ATOMIC REDUCTION). *Let P be a program with correct conflict and call annotations and let Γ be a valid effect environment such that $\Gamma \vdash P$. Suppose $N(P)$ and $N(Q)$ and $P \overset{0}{\Longrightarrow}{}^{*} Q$. Then $P \mapsto^* Q$.*

PROOF. By Lemma 17, we can use Lemma 15 to obtain the desired result by substituting the set W_i for $\mathcal{W}_i$, the set R_i for $\mathcal{R}_i$, the set L_i for $\mathcal{L}_i$, the relation $\overset{0}{\Longrightarrow}_i$ for $\rightharpoonup_i$, the relation $\mapsto$ for $\rightharpoonup_c$, the state P for p, and the state Q for q. □

LEMMA 17 (ATOMIC REDUCTION HELPER). *Let P be a program with correct conflict and call annotations and let Γ be a valid type environment such that $\Gamma \vdash P$. Suppose $N(P)$ and $N(Q)$ and $P \overset{0}{\Longrightarrow}{}^{*} Q$. Then the following statements are true for all thread indices i and j such that $i \neq j$,*

1. *R_i, L_i, U_i, and N_i are disjoint and their union yields the entire state space,*
2. *$(L_i / \overset{0}{\Longrightarrow} \backslash R_i)$ is false,*
3. *$(\overset{0}{\Longrightarrow}_i \backslash R_i)$ right-commutes with $\overset{0}{\Longrightarrow}_j$,*
4. *$(L_i / \overset{0}{\Longrightarrow}_i)$ left-commutes with $\overset{0}{\Longrightarrow}_j$,*
5. *if $P \overset{0}{\Longrightarrow}_i Q$, then $R_j(P) \Leftrightarrow R_j(Q)$, $L_j(P) \Leftrightarrow L_j(Q)$, and $W_j(P) \Leftrightarrow W_j(Q)$.*
6. *if $P \overset{0}{\Longrightarrow}_i Q$ and $U_i(P)$ then $U_i(Q)$.*

PROOF.

1. By the definition of R_i, L_i, U_i and N_i.

2. Suppose $P_1 \rightarrow_i P_2$ where $P_1 \in L_i$ and $P_2 \in R_i$. The proof follows by a straightforward case analysis on the transition rule for $P_1 \rightarrow_i P_2$.

3. Suppose $P_1 \rightarrow_i P_2 \rightarrow_j P_3$ where $i \neq j$ and $P_2 \in R_i$ and $N^1(P_1)$ and $N^1(P_2)$.

   We proceed by case analysis on the rule for $P_1 \rightarrow_i P_2$.

   (a) [RED READ] We have $P_1 = H, M, T$ such that $T_i \equiv E[x_r]$, and $P_2 = H, M, T[i \mapsto E[H(x)]]$.

   First consider the case where x is unstable. If $P_1 \rightarrow_i P_2$ by rule [RED READ], then $P_1 \rightarrow_i P_2$ by rule [ABS UNSTABLE READ]. Since [ABS UNSTABLE READ] does not access the heap or synchronization heap, $\rightarrow_i$ right-commutes with $\rightarrow_j$ in this case.

   Now consider the case where x is a stable variable and $r = \epsilon$. Since the conflict tag is correct and the other thread j is not about to read or write x, so $\rightarrow_i$ right-commutes with $\rightarrow_j$.

   Finally consider the case where x is a stable variable and $r = \bullet$. Then there exists E', E'' such that

$E'[\texttt{in-atomic}\ E'']$	$\equiv$	E	since $P_2 \in R_i$
$\alpha(E''[H(x)])$	$\not\sqsubseteq$	L	since $P_2 \in R_i$
$\alpha(E''[H(x)])$	$\sqsubseteq$	A	from [EXP INATOMIC]
$\alpha(E''[H(x)])$	$\in$	$\{R, A\}$	from prev 2 stmts.
$\alpha(x_\bullet)$	$=$	A	from [EXP READ RACE]
$\alpha(E''[x_\bullet])$	$=$	$\top$	from Lemma 14 and prev 2 stmts.
$\alpha(E''[x_\bullet])$	$\sqsubseteq$	A	from [EXP INATOMIC]

   The last two statements give a contradiction, so this case is impossible.

   (b) [RED WRITE] or [ABS UNSTABLE WRITE] Similar to the Pure Reduction Helper Lemma.

   (c) [RED PRIM] In this case $P_1 = H, M, T$ where $T_i \equiv E[\texttt{in-atomic}\ E'[p(\overline{v})]]$. Again, by similar reasoning, $\alpha(p(\overline{v})) \sqsubseteq R$, and hence $\Gamma(p) \sqsubseteq R$. Hence, the application of p right-commutes with any primitive application of thread j (which are the only steps that modify the synchronization heap), so $\rightarrow_i$ right-commutes with $\rightarrow_j$.

   (d) [RED FORK] In this case $P_1 = H, M, T$ such that $T_i \equiv E[\texttt{in-atomic}\ E'[\texttt{fork}\ e]]$ and $P_2 = H, M, T[i \mapsto E[\texttt{in-atomic}\ E'[0]]$. From the subject reduction theorem, we know that $\alpha(E'[\texttt{fork}\ e]) \sqsubseteq A$. From Lemma 14, we know that

   $$\alpha(E'[\texttt{fork}\ e]) = \alpha(\texttt{fork}\ e); \alpha(E'[0]) = A; \alpha(E'[0]).$$

   Therefore, we must have $\alpha(E'[0]) \sqsubseteq L$. Therefore $P_2 \in L_i$ and we have a contradiction. That is, it is impossible to perform a fork operation and remain in the right-mover part of an atomic block.

(e) If $P_1 \rightarrow_i P_2$ by any of [RED LET], [RED CALL], or [ABS GARBAGE], then only unstable variables and locations not previously allocated in the synchronization heap and heap are accessed. As in the Pure Reduction Lemma, $\rightarrow_i$ right-commutes with $\rightarrow_j$ in this case.

(f) If $P_1 \rightarrow_i P_2$ by any of [RED WRONG], [RED IF1], [RED IF2], [RED LOOP], [RED BREAK], [RED BLOCK], [RED ATOMIC], [RED IN-ATOMIC], [RED PURE], [RED PURE EX], [RED IN-PURE], [ABS PURE SKIP], or [ABS UNSTABLE READ], then these rules do not access the heap or the synchronization heap, and so $\rightarrow_i$ right-commutes with $\rightarrow_j$.

4. Suppose $P_1 \rightarrow_j P_2 \rightarrow_i P_3$ where $i \neq j$ and $P_2 \in L_i$ and $N^1(P_1)$ and $N^1(P_2)$. A straightforward case analysis on the rule for $P_2 \rightarrow_i P_3$ will yield the desired result.

(a) [RED READ] In this case $P_2 = H, M, T$ such that $T_i \equiv E[x_r]$ and $P_3 = H, M, T[i \mapsto E[H(x)]$.
First consider the case where x is unstable. If $P_2 \rightarrow_i P_3$ by rule [RED READ], then $P_2 \rightarrow_i P_3$ by rule [ABS UNSTABLE READ]. Since [ABS UNSTABLE READ] does not access the heap or synchronization heap, $\rightarrow_i$ right-commutes with $\rightarrow_j$ in this case.

Now consider the case where x is a stable variable and $r = \epsilon$. Since the conflict tag is correct and the other thread j is not about to read or write x, so $\rightarrow_i$ right-commutes with $\rightarrow_j$.

Finally consider the case where x is a stable variable and $r = \bullet$. Then there exists E', E'' such that

$E'[\texttt{in-atomic}\ E'']$	$\equiv$	E	since $P_2 \in L_i$
$\alpha(E''[x_\bullet])$	$\in$	$\{B, L\}$	since $P_2 \in L_i$
$\alpha(E''[H(x)])$	$\in$	$\{\bot, B, L\}$	by Lemma 12
$\alpha(x_\bullet)$	$=$	A	from [EXP READ RACE]
$\alpha(E''[x_\bullet])$	$=$	$\{\bot, A\}$	from Lemma 14 and prev 2 stmts.

The second and last statements give a contradiction, so this case is impossible.

(b) [RED WRITE] or [ABS UNSTABLE WRITE] Similar to the Pure Reduction Helper Lemma.

(c) [RED PRIM] In this case $P_2 = H, M, T$ where $T_i \equiv E[\texttt{in-atomic}\ E'[p(\overline{v})]]$. Again, by similar reasoning, $\alpha(p(\overline{v})) \sqsubseteq L$, and hence $\Gamma(p) \sqsubseteq L$. Hence, the application of p left-commutes with any primitive application of thread j (which are the only steps that modify the synchronization heap), so $\rightarrow_i$ left-commutes with $\rightarrow_j$.

(d) [RED FORK] In this case $P_2 = H, M, T$ such that $T_i \equiv E[\texttt{in-atomic}\ E'[\texttt{fork}\ e]]$ and $P_3 = H, M, T[i \mapsto E[\texttt{in-atomic}\ E'[0]]$. From the subject reduction theorem, we know that $\alpha(E'[\texttt{fork}\ e]) \sqsubseteq A$. From Lemma 14, we know that

$$\alpha(E'[\texttt{fork}\ e]) = \alpha(\texttt{fork}\ e); \alpha(E'[0]) = A; \alpha(E'[0]).$$

Therefore $\alpha(E'[\texttt{fork}\ e]) \not\sqsubseteq L$. Therefore $P_2 \notin L_i$ and we have a contradiction. That is, it is impossible to perform a fork operation from the left-mover part of an atomic block.

(e) If $P_2 \rightarrow_i P_3$ by any of [RED LET], [RED CALL], or [ABS GARBAGE], then only unstable variables and locations not previously allocated in the synchronization heap and heap are accessed. As in the Pure Reduction Lemma, $\rightarrow_i$ left-commutes with $\rightarrow_j$ in this case.

(f) If $P_2 \rightarrow_i P_3$ by any of [RED WRONG], [RED IF1], [RED IF2], [RED LOOP], [RED BREAK], [RED BLOCK], [RED ATOMIC], [RED IN-ATOMIC], [RED PURE], [RED PURE EX], [RED IN-PURE], [ABS PURE SKIP], or [ABS UNSTABLE READ], then these rules do not access the heap or the synchronization heap, and so $\rightarrow_i$ left-commutes with $\rightarrow_j$.

5. Suppose $P \rightarrow_i Q$. If this step is not a fork step, then threads other than i do not change in going from P to Q. Therefore $R_j(P)$ iff $R_j(Q)$, $L_j(P)$ iff $L_j(Q)$, $U_j^n(P)$ iff $U_j^n(Q)$, and $W_j(P)$ iff $W_j(Q)$.

If this step forks a new thread j, then $P = H, M, T$ and $T_i = E[\texttt{fork}\ e]$ where e does not contain `in-atomic`, since it does not occur in an evaluation context position. Let j be the newly forked thread. Then $N_j(P)$ (by the definition of N_j) and $N_j(Q)$. Furthermore, since threads other than i or j do not change in going from P to Q, we have that $R_j(P)$ iff $R_j(Q)$, $L_j(P)$ iff $L_j(Q)$, $U_j^n(P)$ iff $U_j^n(Q)$, and $W_j(P)$ iff $W_j(Q)$.

6. Suppose $P \rightarrow_i Q$ and $P \in U_i$. Then $P = H, M, T$ and $T_i = E[\texttt{in-atomic}\ e]$ and $\alpha_n(e) = \bot$. Hence $Q = H', M', T'$ and $T'_i = E[\texttt{in-atomic}\ e']$ and, by Lemma 12, $\alpha_n(e') = \bot$ and so $Q \in U_i^n$. □