

Compositional Verification of Relaxed-Memory Program Transformations

Mike Dodds Mark Batty Alexey Gotsman

Abstract

This paper is about verifying program transformations on a relaxed memory model of the kind used in C/C++. For a block of code being transformed, we define a denotation from its behaviour in a set of representative contexts. Our denotation summarises interactions of the code block with the rest of the program both through local and global variables, and through subtle synchronisation effects due to relaxed memory. We can then prove that a transformation does not introduce new program behaviours by comparing the denotations of the code block before and after. Our approach is compositional: by examining only representative contexts, transformations are verified for any context. It is also fully abstract, meaning any valid transformation can be verified. We cover several tricky aspects of C/C++-style memory models, including release-acquire operations, SC fences, and non-atomics. We also define a variant of our denotation that is finite for loop-free blocks of code, at the cost of losing full abstraction. Based on this variant, we have implemented a prototype verification tool and applied it to automatically prove and disprove a range of compiler optimisations.

1. Introduction

This paper is about verifying program transformations in a concurrent relaxed-memory setting. We define a new reasoning principle that characterises whether two blocks of code are related, and we use this principle in a new tool that automatically checks whether program transformations are valid.

Context and objectives. Any program defines a collection of observable behaviours: a sorting algorithm maps unsorted to sorted sequences, and a paint program responds to mouse clicks by updating a rendering. It is often desirable to transform the program without introducing new observable behaviours – for example, a compiler might optimise the program or a programmer might refactor it. Such transformations are called *observational refinements*, and they ensure that properties of the original program will carry over to the transformed version, without needing to be re-checked.

It is also desirable for transformations to be *compositional*, meaning that they can be applied to a block of code irrespective of the surrounding program context. Composi-

tional transformations are particularly useful for automated systems such as compilers, where they are known as *peephole optimisations*.

The semantics of the programming language is highly significant in determining which transformations are valid, because it determines the ways that a block of code being transformed can interact with its context and thereby affect the observable behaviour of the whole program. In a sequential language, a code-block can only interact with the context at its start and end-points; but in a concurrent language, a block can be observed in intermediate states.

Our work applies to a *relaxed memory* concurrent setting, where different threads can observe different, apparently contradictory orders of events. Such behaviour is permitted by programming languages to reflect CPU-level relaxations and to allow compiler optimisations. In this paper we focus on *axiomatic* memory models of the type used in C/C++ and Java. In these models, program executions are represented by structures of events and relations on them, and program semantics is defined by a set of axioms constraining these structures. Reasoning about the correctness of program transformations on such memory models is very challenging, and indeed, compiler optimisations have been repeatedly shown unsound with respect to models they were intended to support [18, 23].

Approach. Suppose we have a proposed transformation from code block P to Q . For any code-block we define a *denotation* which summarises its behaviour in a restricted representative context by a set of *histories*. We then compare the denotations of P and Q to prove that the transformation is valid. This approach is compositional: it requires reasoning only about the code-blocks and representative contexts; the validity of the transformation in an arbitrary program will follow.

More technically, histories track context accesses to variables used in the code-block and temporal relationships between these accesses. They abstract from the syntactic structure of the rest of the program, and numerous other relationships between accesses tracked by the memory model. We prove that this approach is *fully abstract*, meaning that it can verify any valid transformation: considering only representative contexts does not lose generality.

We also define a variant of our denotation that is *finite* for loop-free blocks of code, at the cost of losing full abstraction. We achieve this by further restricting the form of contexts one needs to consider in exchange for tracking more information in histories. For example, we show that, it is unnecessary to consider executions where two distinct operations in the context read the same write.

Using this finite denotation, we have implemented a prototype verification tool, Stellite. Our tool converts an input transformation into a model in the Alloy language [10], and then checks that the transformation is valid using the Alloy* solver [14]. Our tool can prove or disprove a range of introduction, elimination, and exchange compiler optimisations. Many of these were verified by hand in previous work; our tool verifies them automatically.

Contribution. Our contribution is twofold: we define a compositional, fully-abstract reasoning principle for axiomatic relaxed models; and using this we show it is feasible to automatically verify relaxed-memory program transformations. Our model is derived from C/C++ 2011 standard [1], and we handle many of its key features, including release-acquire, non-atomic accesses, and SC fences. However, our aim is not to handle C/C++ per se: rather we target the simplest axiomatic model rich enough to demonstrate our approach.

2. Observation and Transformation

Observational refinement. The notion of *observation* is crucial when determining how different programs are related. For example, observations might be I/O behaviour or writes to special variables. Given program executions X_1 and X_2 , we write $X_1 \preceq_{\text{ex}} X_2$ if the observations in X_1 are replicated in X_2 . Lifting this notion, a program P_1 *observationally refines* another P_2 if every observable behaviour of one could also occur with the other – we write this $P_1 \preceq_{\text{pr}} P_2$. More formally, let $\llbracket - \rrbracket$ be the map from programs to sets of executions. Then we define \preceq_{pr} as:

$$P_1 \preceq_{\text{pr}} P_2 \stackrel{\Delta}{\iff} \forall X_1 \in \llbracket P_1 \rrbracket. \exists X_2 \in \llbracket P_2 \rrbracket. X_1 \preceq_{\text{ex}} X_2 \quad (1)$$

If $P_1 \preceq_{\text{pr}} P_2$, then any property of observable behaviour established of P_2 will also hold of P_1 .

Compositional transformation. Many common program transformations are *compositional*: they modify a sequential fragment of the program without examining the rest of the program. We call the former the *code-block* and the latter its *context*. Code-blocks are analogous to C/C++ blocks, while contexts are analogous to programs with a single block-shaped ‘hole’. Contexts can include sequential code before and after the block, and concurrent code that runs in parallel with it. Code-blocks are sequential, i.e. they do not feature internal concurrency. A context C and code-block B can be composed to give a whole program $C(B)$.

A transformation $B_2 \rightsquigarrow B_1$ replaces some instance of the code-block B_2 with B_1 . To validate such transformation, we must establish whether *every* whole program containing B_1 observationally refines the same program with B_2 substituted. If this holds, we say that B_1 observationally refines B_2 , written $B_1 \preceq_{\text{bl}} B_2$. We define this by lifting the definition of \preceq_{pr} :

$$B_1 \preceq_{\text{bl}} B_2 \stackrel{\Delta}{\iff} \forall C. C(B_1) \preceq_{\text{pr}} C(B_2) \quad (2)$$

If $B_1 \preceq_{\text{bl}} B_2$ holds, then the compiler can replace block B_2 with block B_1 irrespective of the whole program, i.e. $B_2 \rightsquigarrow B_1$ is a valid transformation. Thus, deciding $B_1 \preceq_{\text{bl}} B_2$ is the core problem in validating compositional transformations.

The semantics of the programming language is highly significant in determining which transformations are valid. For example, in a sequential setting, where only initial and final values of variables are observable, the following code-blocks are observationally equivalent, i.e., either of them refines the other:

$$B_1: \text{write}(x,2); \text{write}(x,5) \quad | \quad B_2: \text{write}(x,5)$$

However in a standard concurrent setting where the environment can observe intermediate states, $x = 2$ can be observed in B_2 but not B_1 , meaning the code-blocks are no longer equivalent. A relaxed-memory setting removes the idea of a single state seen by all threads, and so further complicates the notion of observation.

Compositional verification. To establish $B_1 \preceq_{\text{bl}} B_2$, it is undesirable to examine every possible syntactic context: there are an infinite number, with large amounts of irrelevant structure. Instead, the verification process should *also* be compositional. Our approach is to construct a *denotation* for each code-block – a simplified, ideally finite, summary of all possible interactions between the block and its context. We then define a *refinement relation* on denotations and use it to establish observational refinement without examining every context.

We write $B_1 \sqsubseteq B_2$ when the denotation of B_1 refines that of B_2 . Refinement on denotations should be *adequate*, i.e., it should validly approximate observational refinement:

$$B_1 \sqsubseteq B_2 \implies B_1 \preceq_{\text{bl}} B_2$$

Hence, if $B_1 \sqsubseteq B_2$, then $B_2 \rightsquigarrow B_1$ is a valid transformation. It is also desirable for the denotation to be *fully abstract*, i.e.:

$$B_1 \preceq_{\text{bl}} B_2 \implies B_1 \sqsubseteq B_2$$

This means any valid transformation can be verified by comparing denotations. Below we define several versions of \sqsubseteq with different properties.

3. Target Language and Memory Model

Our language’s memory model is derived from the C/C++ 2011 standard (henceforth ‘C11’), as formalized by Batty et

al [1, 3]. However, we simplify in several ways — see end of section for details. In this section we describe a model without non-atomic operations: we restore them in §7.

Relaxed memory. In a sequentially consistent (SC) concurrent system, there is a total temporal order on reads and writes, and reads take the value of the most recent write; in particular, they cannot read values that have been overwritten, or that are written in the future. A *relaxed* (or *weak*) memory system weakens this total order, which allows behaviour forbidden under SC.

A standard relaxed behaviour is *store buffering*:

```

write(x,0); write(y,0);
write(x,1);   || write(y,1);           (SB)
v1 := read(y); || v2 := read(x);

```

Perhaps counter-intuitively, in most relaxed models $v1 = v2 = 0$ is a possible post-state. This cannot occur on an SC system: if $v1 = 0$ then `write(y,1)` must be ordered after the read of y , which would order `write(x,1)` before the read of x , forcing it to assign $v2 = 1$.

Another important example is *message passing*:

```

write(f,0); write(x,0);
write(x,1);   || b := read(f);         (MP)
write(f,1);   || if (b==1)
               ||   r := read(x);

```

In many relaxed models, $b = 1 \wedge r = 0$ is a possible post-state. This is undesirable if, for example, x is a complex data-structure and f is a flag indicating it has been safely created.

Language syntax. Programs in the language we consider manipulate *thread-local variables* $l, l_1, l_2 \dots \in \text{LVar}$ and *global variables* $x, y, \dots \in \text{GVar}$, coming from disjoint sets LVar and GVar . Each variable stores a value from a finite set Val , and is initialised to $0 \in \text{Val}$. We assume that each thread uses the same set of local variable names LVar (though the identity of local variables with the same name used in different threads is distinct).

The syntax of the programming language is as follows:

```

C ::= l := E | write(x,l) | l := read(x) |
     l := RMW(x, l1, l2) | l := LL(x) | l' := SC(x, l) |
     fence | C1 || C2 | C1; C2 |
     if (l) {C1} else {C2} | while (l) {C} | {-}
E ::= l | l1 = l2 | l1 ≠ l2 | ...

```

Many of the constructs are standard. $\text{RMW}(x, l_1, l_2)$ is a *read-modify-write*. It reads the value of x and compares it to the value in l_1 . If they are equal, then it assigns x to the value in l_2 and returns 1; otherwise it returns 0. $\text{LL}(x)$ and $\text{SC}(x, l)$ are load-link and store-conditional, which are used by many platforms to implement RMW . A load-link behaves as a standard read, but if it is followed by a store-conditional to the same location, the store fails if there are intervening writes to the same location. The `fence` command is an SC

fence: interleaving such fences between all statements in a program would guarantee sequential consistent behaviour.

The construct $\{-\}$ represents a block-shaped hole in the program. The set Prog of *whole programs* consists of programs without holes, while the set Contx of *contexts* consists of programs with a single hole. To simplify our presentation, we assume that holes in a context only appear in loop-free positions. The set Block of *code-blocks* are whole programs without parallel composition. We often write $P \in \text{Prog}$ for a whole program, $B \in \text{Block}$ for a code-block, and $C \in \text{Contx}$ for a context. Given a context C and a code-block B , the composition $C(B)$ is C with its single block-shaped hole syntactically replaced by B .

Memory model structure. The semantics of a whole program P is given by a set $\llbracket P \rrbracket$ of *executions*, which consist of *actions*, representing memory events on global variables, and several relations on these. Actions are tuples:

$$\text{Actions} \triangleq \text{ActID} \times \text{Kind} \times \text{Option}(\text{GVar}) \times \text{Val}^*$$

In an action $(a, k, z, b) \in \text{Action}$: $a \in \text{ActID}$ is the unique action identifier; $k \in \text{Kind}$ is the kind of action – we use `read`, `write`, `ll`, `sc`, and their failed variants `llf` and `scf` in the semantics, and will introduce further kinds as needed; $z \in \text{Option}(\text{GVar})$ is an option type consisting of either a single global variable $\text{Just}(x)$ or None ; and $b \in \text{Val}^*$ is the vector of values. The kind of the action dictates the number of variables and values: each read or write accesses one global variable and has one value.

Given an action v , we use $\text{kind}(v)$, $\text{gvar}(v)$ and $\text{var}(v)$ as selectors for the different fields. We often write actions so as to elide action identifiers and the option type. For example, `read(x, 3)` stands for $\exists i. (i, \text{read}, \text{Just}(x), [3])$. We also sometimes elide values. Given a set of actions \mathcal{A} , we write e.g. `reads(\mathcal{A})` and `writes(\mathcal{A})` to identify actions of each kind in \mathcal{A} . We range over actions by u, v , read actions by r and write actions by w .

The semantics of a program $P \in \text{Prog}$ is defined in two stages. First, a *thread-local semantics* of P produces a set $\langle P \rangle$ of *pre-executions* $(\mathcal{A}, \text{sb}) \in \text{PreExec}$. A pre-execution contains a finite¹ set of memory actions $\mathcal{A} \in \text{Action}$ that could be produced by the program. It has a transitive and irreflexive *sequence-before* relation $\text{sb} \subseteq \mathcal{A} \times \mathcal{A}$, which defines the sequential order imposed by the program syntax.

Thread-local semantics takes into account control flow in P 's threads and operations on local variables. However, it does not constrain the behaviour of global variables: the values threads read from them are chosen arbitrarily. This is addressed by extending pre-executions with extra relations, and filtering these executions using *validity axioms*.

Thread-local semantics. We show selected clauses in Figure 1. The full semantics is given in §A.

¹ We do not consider infinite computations in this paper, since their semantics in axiomatic memory models is not settled in the literature [5].

$$\begin{aligned}
\langle \text{write}(x, l), \sigma \rangle &\triangleq \{(\{\text{write}(x, a)\}, \emptyset, \sigma) \mid \sigma(l) = a\} \\
\langle l := \text{read}(x), \sigma \rangle &\triangleq \{(\{\text{read}(x, a)\}, \emptyset, \sigma[l \mapsto a]) \mid a \in \text{Val}\} \\
\langle C_1; C_2, \sigma \rangle &\triangleq \{(\mathcal{A}_1 \cup \mathcal{A}_2, \text{sb}_1 \cup \text{sb}_2 \cup (\mathcal{A}_1 \times \mathcal{A}_2), \sigma_2) \mid (\mathcal{A}_1, \text{sb}_1, \sigma_1) \in \langle C_1, \sigma \rangle \wedge (\mathcal{A}_2, \text{sb}_2, \sigma_2) \in \langle C_2, \sigma_1 \rangle\}
\end{aligned}$$

Figure 1. Read, write, and sequential composition in the thread-local semantics. The full semantics is given in §A. We write $\mathcal{A}_1 \cup \mathcal{A}_2$ for a union that is defined only when actions in \mathcal{A}_1 and \mathcal{A}_2 use disjoint sets of identifiers.

To track the values of local variables in the thread-local semantics, we use maps $\sigma \in \text{VMap} \triangleq \text{LVar} \rightarrow \text{Val}$. The thread-local semantics is defined using a function

$$\langle -, - \rangle : \text{Prog} \times \text{VMap} \rightarrow \mathcal{P}(\text{PreExec} \times \text{VMap})$$

Given a program P and an initial local variable map σ , $\langle P, \sigma \rangle$ yields the set of pre-executions of P paired with final variable maps and final open load-linked regions. Note that in Figure 1, the clause for `read`, a is unrestricted, and hence the read can take any value in `Val`.

We take a simplified approach to local variables at thread creation: the initial variable map σ is copied to both threads in $C_1 \parallel C_2$, and the original map is restored when they complete. The RMW command is encoded either as a plain read, or as a successful `ll/sc` pair. The `fence` command is encoded by a successful `ll/sc` pair to a distinguished variable $\text{fen} \in \text{GVar}$ that is not otherwise read or written.²

Let σ_0 map every local variable to 0, then the thread-local semantics of a program $P \in \text{Prog}$ is defined as follows:

$$\langle P \rangle \triangleq \{(\mathcal{A}, \text{sb}) \mid (\mathcal{A}, \text{sb}, \sigma') \in \langle P, \sigma_0 \rangle\}$$

Execution structure. The semantics of a program P is a set $\llbracket P \rrbracket$ of *executions* $X = (\mathcal{A}, \text{sb}, \text{at}, \text{mo}, \text{rf}, \text{hb}) \in \text{Exec}$, where (\mathcal{A}, sb) is a pre-execution and the relations $\text{at}, \text{mo}, \text{rf}, \text{hb} \subseteq \mathcal{A} \times \mathcal{A}$. Given an execution X we sometimes write $\mathcal{A}(X), \text{sb}(X), \dots$ as selectors for the appropriate set or relation. The relations have the following purposes.

- *Reads-from* (`rf`) is an injective map from write and `sc` actions to read, `ll` and `llf` actions. A write and reading action are related $w \xrightarrow{\text{rf}} r$ if r takes its value from w .
- *Modification order* (`mo`) is an irreflexive, total order on write and `sc` actions to each distinct variable. This is a per-variable order in which *all* threads observe writes; two threads cannot observe writes to a variable in different orders.
- *Atomicity* ($\text{at} \subseteq \text{sb}$) is a bijection which associates each successful load-link action to a successful store-conditional on the same location.

- *Happens-before* (`hb`) is the closest the memory model has to global time – however, unlike the SC notion of time, it is partial. Happens-before is defined as $(\text{sb} \cup \text{rf})^+$: therefore statements ordered in the program syntax are ordered in time, as are reads with the writes they observe.

Validity axioms. The semantics $\llbracket P \rrbracket$ of a program P is the set of executions $X \in \text{Exec}$ that are compatible with the thread-local semantics $\langle P \rangle$ and that satisfy certain *validity axioms*, denoted $\text{valid}(X)$:

$$\llbracket P \rrbracket \triangleq \{X \mid (\mathcal{A}(X), \text{sb}(X)) \in \langle P \rangle \wedge \text{valid}(X)\} \quad (3)$$

The axioms on an execution $(\mathcal{A}, \text{sb}, \text{at}, \text{rf}, \text{mo}, \text{hb})$ are:

- **HBDEF:** $\text{hb} = (\text{sb} \cup \text{rf})^+$ and `hb` is acyclic.

This axiom defines `hb` and enforces the intuitive property that there are no cycles in the temporal order. It also prevents an action reading from its `hb`-future: as `rf` is included in `hb`, this would result in a cycle.

- **HBvsMO:** $\neg \exists w_1, w_2. w_1 \xrightarrow{\text{hb}} w_2 \xrightarrow{\text{mo}} w_1$

This axiom requires that the the order in which writes to a location become visible to threads cannot contradict the temporal order. However, note that writes may be ordered in `mo` but not `hb`.

- **COHERENCE:** $\neg \exists w_1, w_2, r. w_1 \xrightarrow{\text{mo}} w_2 \xrightarrow{\text{hb}} r \xrightarrow{\text{rf}} w_1$

This axiom generalises the SC prohibition on reading overwritten values. If two writes are ordered in `mo`, then intuitively the second overwrites the first. A read / RMW that follows some write in `hb` or `mo` cannot read from writes earlier in `mo` – these earlier writes have been overwritten. However, unlike in SC, `hb` is partial, so there may be multiple writes that an action can legally read.

- **RFVAL:**

$$\begin{aligned}
\forall r. (\neg \exists w'. w' \xrightarrow{\text{rf}} r) \implies \\
(\text{rval}(r) = 0 \wedge (\neg \exists w. w \xrightarrow{\text{hb}} r \wedge \text{var}(w) = \text{var}(r)))
\end{aligned}$$

Most reads must take their value from a write, represented by an `rf` edge. However, the **RFVAL** axiom allows the `rf` edge to be omitted if the read takes the initial value

²This provides a stronger semantics than C/C++11; here we follow Lahav et al. [12], who argue this strengthening is sound under existing compilation strategies to common multiprocessors.

0 and there is no hb-earlier write to the same location. Intuitively, an hb-earlier write would supersede the initial value in a similar way to COHERENCE.

- **ATOM:** $\neg \exists w_1, l, w_2, s. w_1 \xrightarrow[\text{rf}]{\text{mo}} l \xrightarrow[\text{at}]{\text{mo}} w_2 \xrightarrow{\text{mo}} s$

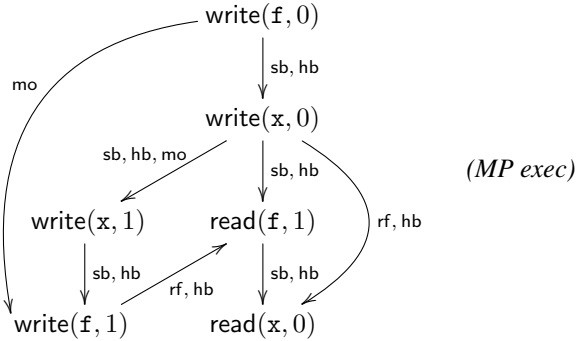
For load-link store conditional-pairs, this axiom (adapted from Lahav et al. [12]) ensures that there is no mo-intervening write that would invalidate the store.

- **NOFAIL:**


$$\neg \exists l, f, s \in . l \xrightarrow[\text{at}]{\text{sb}} f \xrightarrow{\text{sb}} s \wedge \text{gvar}(l) = \text{gvar}(f) \wedge f \in (\text{ll}_f(\mathcal{A}) \cup \text{sc}_f(\mathcal{A}))$$

This final axiom ensures that ll/sc pairs cannot have intervening failed accesses to the same location.

Example executions. Our model forbids the problematic MP relaxed behaviour discussed above. We can try to construct an execution exhibiting the relaxed behaviour $\mathbf{b} = 1 \wedge \mathbf{r} = 0$ as follows. To avoid clutter, we omit hb edges obtained by transitivity.



Although this execution is allowed by the thread-local semantics of the MP program, it is ruled out by the COHERENCE axiom. As hb is transitively closed, there is a derived hb edge $\text{write}(x, 1) \xrightarrow{\text{hb}} \text{read}(x, 0)$, which forms a COHERENCE violation.



Thus this is not an execution of the MP program. Indeed, any execution ending in $\text{read}(x, 0)$ is forbidden for the same reason, meaning that the MP relaxed behaviour cannot occur.

Relaxed observations. Finally, we define a notion of observational refinement suitable for our relaxed model. We assume a subset of *observable* global variables, $\text{OVar} \subseteq \text{GVar}$, which can only be accessed by the context and not by the code-block. We consider the actions and the hb relation on these variables to be the observations. We write $X|_{\text{OVar}}$ for the projection of X 's action set and relations to OVar , and

use this to define \preceq_{ex} for our model:

$$X \preceq_{\text{ex}} Y \iff \mathcal{A}(X|_{\text{OVar}}) = \mathcal{A}(Y|_{\text{OVar}}) \wedge \text{hb}(Y|_{\text{OVar}}) \subseteq \text{hb}(X|_{\text{OVar}})$$

This is lifted to programs and blocks as in §2, def. (1) and (2).

Note that in the more abstract execution, actions on observable variables must be the same, but hb can be weaker. This is because we interpret hb as a constraint on time order: two actions that are unordered in hb could have occurred in either order, or in parallel. Thus, weakening hb allows more observable behaviours (see §2).

Differences from C11. Our model is derived from Batty et al's C11 formalization [3], with a number of simplifications.

We omit SC accesses to reduce complexity in the history. We also assume a static set of shared variables, and omit pointers. In both cases, we believe these limitations could be lifted at the cost of more complexity and more challenging automation: see §10 for a discussion.

We omit RLX accesses to avoid well-known problems with thin-air values [5]. We know of at least four current proposals for fixing these problems, and it would be premature to select one.

We add LL-SC atomic instructions in addition to C11's RMW. This increases the context's observational power and gives us a full abstraction result. LL/SC is commonly available as a hardware instruction on platforms supporting C11.

4. Denotations of Code-Blocks

An obvious way to generate the denotation of a code-block would be to quantify over all its contexts. However, we show one needs only to consider a set of limited representative contexts. (In §5 we show that even these contexts can be cut down further, giving a finite denotation for code-blocks whose set of thread-local executions is finite.) We construct the denotation for a code-block in two steps:

1. Generate the *block-local* executions of the code-block. These are executions under a set of special cut-down contexts that are sufficient to represent interactions of all possible contexts with the block.
2. From these, extract an interaction summary called a *history*. The denotation is the set of all such histories.

Block-local executions. The block-local executions of a block $B \in \text{Block}$ omit context structure such as syntax and actions on variables not accessed in the block. Instead the context is represented by special actions *call* and *ret*, a set \mathcal{A}_B , and relations R_B and S_B , each covering an aspect of the interaction of the block and an arbitrary unrestricted context.

- **Local variables.** A context can include code that precedes and follows the block on the same thread, with interaction through local variables. We capture this with special action $\text{call}(\sigma)$ at the start of the block, and $\text{ret}(\sigma')$ at the

end, where $\sigma, \sigma' : \text{LVar} \rightarrow \text{Val}$ record the values of local variables at these points. Assume that variables in LVar are ordered: l_1, l_2, \dots, l_n . Then call has the following encoding as an action, with fresh identifier i :

$$\text{call}(\sigma) \triangleq (i, \text{call}, \text{None}, [\sigma(l_1), \dots, \sigma(l_n)])$$

We encode ret in the same way.

- *Global variable actions.* The context can also interact with the block through concurrent reads and writes to global variables. These interactions are represented by set \mathcal{A}_B of actions added to the ones generated by the thread-local semantics of the block. This set only contains actions on the variables VS_B that B can access (VS_B can be constructed syntactically).

- *Context happens-before.* The context can generate hb edges between its actions – to get adequacy (§2), we track these with a relation R_B over actions in \mathcal{A}_B , call and ret:

$$R_B \subseteq (\mathcal{A}_B \times \mathcal{A}_B) \cup (\mathcal{A}_B \times \{\text{call}\}) \cup (\{\text{ret}\} \times \mathcal{A}_B) \quad (4)$$

The context can generate hb edges between actions directly if they are on the same thread, or indirectly through inter-thread reads. Likewise call / ret may be related to context actions on the same or different threads.

- *Context atomicity.* The context can generate at edges between its actions that we capture in the relation $S_B \subseteq (\mathcal{A}_B \times \mathcal{A}_B)$. We consider only cases where LL/SC pairs do not cross block boundaries, so we need not consider boundary-crossing at edges.

When constructing block-local executions, we represent all possible interactions by quantifying over all possible choices of $\sigma, \sigma', \mathcal{A}_B, R_B$ and S_B .

Together, call, ret, \mathcal{A}_B, R_B , and S_B represent a limited context, stripped of syntax, relations at, rf, sb, and mo, and actions on global variables other than VS_B . The set $\llbracket B, \mathcal{A}_B, R_B, S_B \rrbracket$ contains all executions of B under this special limited context. Formally, an execution $X = (\mathcal{A}, \text{sb}, \text{mo}, \text{rf}, \text{at}, \text{hb})$ is in this set if:

1. $\mathcal{A}_B \subseteq \mathcal{A}$ and there exist variable maps σ, σ' such that $\{\text{call}(\sigma), \text{ret}(\sigma')\} \subseteq \mathcal{A}$. That is, the call, return, and extra context actions are included in the execution.
2. There exists a set \mathcal{A}_l and relation sb_l such that
 - $(\mathcal{A}_l, \text{sb}_l, \sigma') \in \langle B, \sigma \rangle$
 - $\mathcal{A}_l = \mathcal{A} \setminus (\mathcal{A}_B \cup \{\text{call}, \text{ret}\})$
 - $\text{sb} = \text{sb}_l \cup \{(\text{call}, u), (u, \text{ret}) \mid u \in \mathcal{A}_l\}$

That is, actions from the code-block satisfy the thread-local semantics, beginning with map σ , and deriving map σ' . All actions arising from the block are between call and ret in sb.

3. X satisfies the validity axioms, but with modified axioms HBDEF' , ATOM' , and NOFAIL' . We define HBDEF' as:

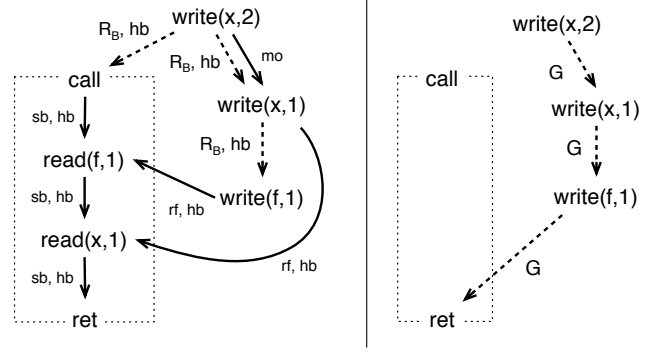


Figure 2. Left: block-local execution. Right: associated history.

HBDEF' : $\text{hb} = (\text{sb} \cup \text{rf} \cup R_B)^+$ and hb is acyclic.

That is, the context relation R_B is added to hb. ATOM' and NOFAIL' are defined analogously with S_B added to at.

We say that \mathcal{A}_B, R_B and S_B are *consistent with B* if they act over variables in the set VS_B . In the rest of the paper we only consider choices that are consistent with the code-block. The *block-local executions* of B are all executions $X \in \llbracket B, \mathcal{A}_B, R_B, S_B \rrbracket$ for consistent choices of \mathcal{A}_B, R_B, S_B .

Example block-local execution. The left side of Figure 2 shows a block-local execution for the code-block

$$l1 := \text{read}(f); l2 := \text{read}(x) \quad (5)$$

Here the set VS_B of accessed global variables is $\{f, x\}$, the context action set \mathcal{A}_B consists of the three writes, and the context relation R_B is denoted by dotted edges.

In this execution, both \mathcal{A}_B and R_B affect the behaviour of the code-block. The following path is generated by R_B and the read of $f = 1$:

$$\text{wr}(x, 2) \xrightarrow{\text{mo}} \text{wr}(x, 1) \xrightarrow{R_B} \text{wr}(f, 1) \xrightarrow{\text{rf}} \text{rd}(f, 1) \xrightarrow{\text{sb}} \text{rd}(x, 1)$$

Because hb includes sb, rf, and R_B , there is a transitive edge $\text{wr}(x, 1) \xrightarrow{\text{hb}} \text{rd}(x, 1)$. The edge $\text{wr}(x, 2) \xrightarrow{\text{mo}} \text{wr}(x, 1)$ is forced because the HBvsMO axiom prohibits mo from contradicting hb. Consequently, the COHERENCE axiom forces the program to read $x = 1$. This is a product of the interaction between code, reads of context actions \mathcal{A}_B , and hb edges in R_B .

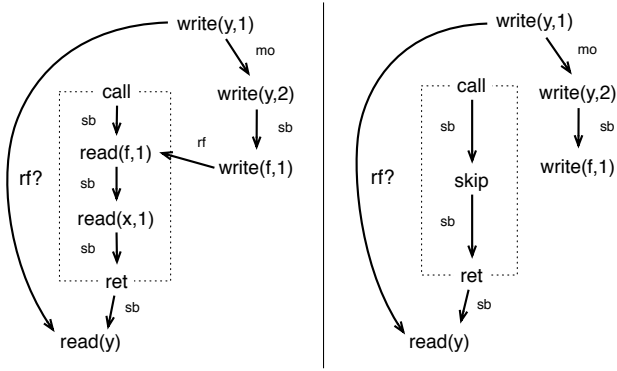
Histories. From any block-local execution X , its *history* summarises the interactions between the code-block and the context. Informally, the history records hb over context actions, call, and ret. More formally the history, written $\text{hist}(X)$, is a pair (A, G) consisting of an action set A and *guarantee relation* $G \subseteq A \times A$. We write $\text{contx}(X)$ to denote actions in $\mathcal{A}(X)$ outside the code-block, and define the history as follows:

- The action set A is the projection of X 's action set to call, ret, and $\text{ctx}(X)$.
- The guarantee relation G is the projection of $\text{hb}(X)$ to

$$\begin{aligned} & (\text{ctx}(X) \times \text{ctx}(X)) \cup \\ & (\text{ctx}(X) \times \{\text{ret}\}) \cup (\{\text{call}\} \times \text{ctx}(X)) \end{aligned} \quad (6)$$

The guarantee summarises the code-block's effect on its context: it suffices to only track hb and ignore other relations. Note the guarantee definition is similar to the context relation R_B (4). The difference is that call and ret are switched: this is because the guarantee represents hb edges generated by the code-block, while R_B represents the edges generated by the context. The right of Figure 2 shows the history corresponding to the block-local execution on the left.

To see why the history needs to record the guarantee at all, compare the code-block $\text{read}(f)$; $\text{read}(x)$ with a block consisting of just `skip`. Note that neither modifies any variables. Below we give executions of both blocks in a context that also reads and writes to a variable y .



On the left, the context read of y cannot take value 1 because there is a path $\text{wr}(y, 1) \xrightarrow{\text{mo}} \text{wr}(y, 2) \xrightarrow{\text{hb}} \text{rd}(f) \xrightarrow{\text{rf}} \text{wr}(f, 1) \xrightarrow{\text{sb}} \text{rd}(y)$ created by the read of f ³. This means that reading $y = 1$ contradicts the COHERENCE axiom. Conversely, on the right there is no such path and the context read of y can read 1. Thus these two code-blocks have different effects on the context through happens-before, even though they are identical in their writes to variables. The guarantee tracks these kinds of hb effects.

Comparing denotations. The denotation of a code-block B is the set of histories of block-local executions of B under each possible context, i.e.

$$\{\text{hist}(X) \mid \exists \mathcal{A}_B, R_B, S_B. X \in \llbracket B, \mathcal{A}_B, R_B, S_B \rrbracket\}$$

To allow us to compare the denotations of two code-blocks, we first define a *refinement relation* on histories:

$$(A_1, G_1) \sqsubseteq_{\text{h}} (A_2, G_2) \iff A_1 = A_2 \wedge G_2 \subseteq G_1$$

³The read of f could also take the initial value, but to illustrate the effects of hb we assume this does not happen.

Here the history (A_2, G_2) places fewer restrictions on the context than (A_1, G_1) . This matches our definition of observational refinement in §3 – a weaker guarantee corresponds to more observable behaviours.

We write $B_1 \sqsubseteq_{\text{q}} B_2$ to state that the denotation of B_1 *refines* that of B_2 . To distinguish from later definitions, the subscript ‘q’ stands for the fact we *quantify* over both \mathcal{A} and R . We define \sqsubseteq_{q} by lifting \sqsubseteq_{h} :

$$\begin{aligned} B_1 \sqsubseteq_{\text{q}} B_2 & \iff \forall \mathcal{A}, R, S. \forall X_1 \in \llbracket B_1, \mathcal{A}, R, S \rrbracket. \\ & \exists X_2 \in \llbracket B_2, \mathcal{A}, R, S \rrbracket. \text{hist}(X_1) \sqsubseteq_{\text{h}} \text{hist}(X_2) \end{aligned} \quad (7)$$

In other words, two code-blocks are related $B_1 \sqsubseteq_{\text{q}} B_2$ if for every block-local execution of B_1 , there is a corresponding execution of B_2 with a related history. Note that the corresponding history must be constructed under the same \mathcal{A}, R, S to ensure that the behaviour can be replicated in the same context.

THEOREM 1 (\sqsubseteq_{q} adequacy). $B_1 \sqsubseteq_{\text{q}} B_2 \implies B_1 \preceq_{\text{bl}} B_2$

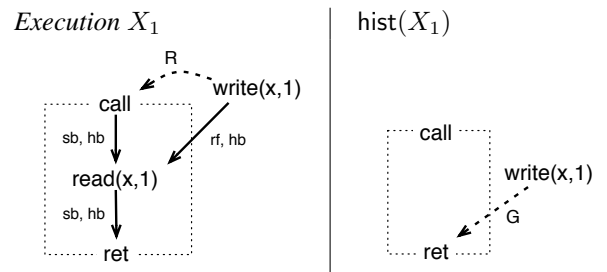
THEOREM 2 (Full abs. of \sqsubseteq_{q}). $B_1 \preceq_{\text{bl}} B_2 \implies B_1 \sqsubseteq_{\text{q}} B_2$.

We prove Theorem 1 in §B, and discuss Theorem 2 in §8. As a corollary these two theorems, a program transformation $B_2 \rightsquigarrow B_1$ is valid if and only if $B_1 \sqsubseteq_{\text{q}} B_2$ holds.

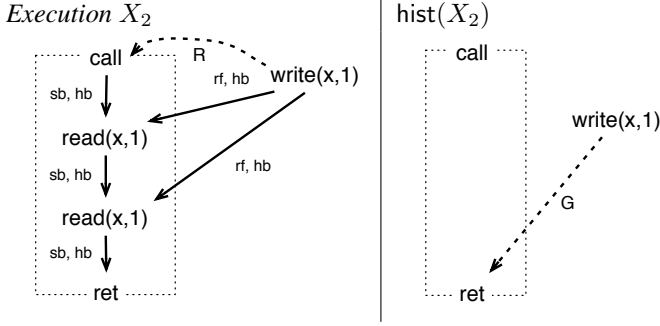
Validating a transformation. We now apply our approach to a simple program transformation. To verify this transformation, we require that $B_1 \sqsubseteq_{\text{q}} B_2$.

$$B_2: l := \text{read}(x); l := \text{read}(x) \rightsquigarrow B_1: l := \text{read}(x)$$

We illustrate the necessary reasoning for a single block-local execution $X_1 \in \llbracket B_1, \mathcal{A}, R, S \rrbracket$, with a context action set \mathcal{A} consisting of a single write $x = 1$, a context relation R relating the write to call, and an empty S relation. This choice of R forces the read to read from the context write:



We can exhibit an execution $X_2 \in \llbracket B_2, \mathcal{A}, R, S \rrbracket$ with a matching history by making both reads read from the same context write:



5. A Finite Denotation

The approach described above simplifies $\mathcal{A}/R/S$ contexts, but there are still infinitely many for any code-block, so we cannot simply enumerate in order to verify a transformation. To solve this, we modify our denotation. This gives us finiteness in some cases, meaning we can automatically check transformations (see our tool in §6). However, this comes at a cost: the new approach is no longer fully abstract.

We modify our denotation as follows:

- We eliminate redundant block-local executions from the denotation by only considering those that satisfy a predicate cut . Intuitively, the denotation must consider each *pattern* of context behaviour, but it need not consider every execution. For example, we need not have two context reads reading the same code-block write: one is sufficient.
- We remove the quantification over context relation R from definition (7) by fixing it as \emptyset . In exchange, we extend the history with an extra component called a *deny*.

Before defining these steps in detail, we give the structure of our modified refinement \sqsubseteq_c . In the definition, $\text{hist}_E(X)$ stands for the *extended history* of an execution X , and \sqsubseteq_E for refinement on extended histories.

$$B_1 \sqsubseteq_c B_2 \iff \forall \mathcal{A}, S. \forall X_1 \in \llbracket B_1, \mathcal{A}, \emptyset, S \rrbracket. \text{cut}(X_1) \implies \exists X_2 \in \llbracket B_2, \mathcal{A}, \emptyset, S \rrbracket. \text{hist}_E(X_1) \sqsubseteq_E \text{hist}_E(X_2)$$

THEOREM 3 (\sqsubseteq_c adequacy). $B_1 \sqsubseteq_c B_2 \implies B_1 \preceq_{bl} B_2$.

The proof of adequacy is given in §D.

By *finiteness*, we mean that a code-block has a finite number of block-local executions satisfying cut . Because block-local executions are derived from pre-executions in the thread-local semantics, finiteness only holds if the set of the latter is finite, for example for loop-free blocks. Note that we also assume a finite domain of values in Val .

THEOREM 4 (Finiteness). *If for any σ the set $\langle B, \sigma \rangle$ is finite, then so is the set $\{X \mid \exists \mathcal{A}. X \in \llbracket B, \mathcal{A}, \emptyset, S \rrbracket \wedge \text{cut}(X)\}$.*

Cutting predicate. We first identify the actions in an execution that are *visible*, meaning they directly affect the behaviour of the block. We write $\text{code}(X)$ for the set of actions

in X generated by the code-block. Visible actions belong to $\text{code}(X)$, are read from $\text{code}(X)$, or are read by $\text{code}(X)$:

$$\text{vis}(X) \triangleq \text{code}(X) \cup \{u \mid \exists v \in \text{code}(X). u \xrightarrow{rf} v \vee v \xrightarrow{rf} u\}$$

The predicate cut is the conjunction of cutR for reads, and cutW for writes. We consider a successful LL-SC pair as a single operation: if cut allows us to keep one half of the LL-SC, we can implicitly keep the other.


$$\text{cut}(X) \triangleq \text{cutR}(X) \wedge \text{cutW}(X)$$

$$\begin{aligned} \text{cutR}(X) &\triangleq \text{reads}(X) \subseteq \text{vis}(X) \wedge \forall r_1, r_2 \in \text{contx}(X). \\ &r_1 \neq r_2 \implies \neg \exists w. w \xrightarrow{rf} r_1 \wedge w \xrightarrow{rf} r_2 \\ \text{cutW}(X) &\triangleq \forall w_1, w_2 \in (\text{contx}(X) \setminus \text{vis}(X)). w_1 \xrightarrow{mo} w_2 \\ &\implies \exists w_3 \in \text{vis}(X). w_1 \xrightarrow{mo} w_3 \xrightarrow{mo} w_2 \end{aligned}$$

The predicate cutR requires that all reads are visible and that pairs of reads must read from distinct writes. In particular, this rules out multiple context reads all reading from the same write. Unlike reads, cutW permits writes that are not visible. However any two non-visible writes to a location must be separated in mo by a visible write. This still achieves finiteness (Theorem 4) because for a given pre-execution B , any two non-visible writes must be distinguished by a visible write, limiting their number.

The non-visible writes in cutW are required for adequacy (§D): they represent contexts where the block interleaves between a context write and read in mo .

Extended history (hist_E). The definition of \sqsubseteq_c removes the context relation R , which records the hb edges enforced by the context, and replaces it with a history component which records the hb edges that *cannot* be enforced due to the execution structure. For example, consider the following block-local execution⁴:



It represents a set of larger execution contexts, but it cannot be embedded into a context that generates the dashed edge D as a hb. We represent such ‘forbidden’ edges D into a separate history component called a *deny*.

The *extended history* of an execution X , written $\text{hist}_E(X)$ is a triple (A, G, D) , consisting of the familiar notions of action set A and guarantee, together with deny as defined

⁴We use this execution to illustrate the extended history, but in fact cut would not allow the context read.

below: $G, D \subseteq A \times A$.

$$\{(u, v) \mid \text{HBvsMO-d}(u, v) \vee \text{Cohere-d}(u, v) \vee \text{RFval-d}(u, v)\} \\ \cap \left(\begin{array}{c} (\text{contx}(X) \times \text{contx}(X)) \cup \\ ((\text{contx}(X) \times \{\text{call}\}) \cup (\{\text{ret}\} \times \text{contx}(X))) \end{array} \right)$$

Each of the predicates HBvsMO-d, Cohere-d, and RFval-d generates the deny for one validity axiom. In the diagrammatic definitions below, dashed edges represent the deny edge, and hb^* is the reflexive-transitive closure of hb:

$$\text{HBvsMO-d}(u, v): \exists w_1, w_2. w_1 \xrightarrow{\text{hb}^*} u \xrightarrow{\text{D}} v \xrightarrow{\text{hb}^*} w_2 \\ \xleftarrow{\text{mo}} w_1$$

$$\text{Coherence-d}(u, v): \exists w_1, w_2, r. w_1 \xrightarrow{\text{mo}} w_2 \\ \downarrow \text{hb}^* \\ u \\ \downarrow \text{D} \\ v \\ \downarrow \text{hb}^* \\ r \\ \swarrow \text{rf} \\ w_1$$

$$\text{RFval-d}(u, v): \exists w, r. \text{var}(w) = \text{var}(r) \wedge \neg \exists w'. w' \xrightarrow{\text{rf}} r \\ \wedge w \xrightarrow{\text{hb}^*} u \xrightarrow{\text{D}} v \xrightarrow{\text{hb}^*} r$$

One can think of a deny edge as an ‘almost’ violation of an axiom. For example, if $\text{HBvsMO-d}(u, v)$ holds, then the context cannot generate an extra hb-edge $u \xrightarrow{\text{hb}} v$ – to do so would violate HBvsMO.

Because deny edges represent constraints on the context, weakening the deny places fewer constraints, allowing more behaviours, so we compare them with relational inclusion:

$$(A_2, G_2, D_2) \sqsubseteq_E (A_1, G_1, D_1) \iff \\ A_1 = A_2 \wedge G_2 \subseteq G_1 \wedge D_2 \subseteq D_1$$

Counter-example to fully abstraction. Finiteness has a cost: the modified denotation is not fully abstract. To see this, consider blocks, $B_1: \text{skip}$ and $B_2: \text{read}(x)$. It is easy to see that $B_1 \sqsubseteq_q B_2$ holds: the new read can read from either a hb-earlier write, or the initialisation if none exists. Neither case introduces an extra guarantee edge.

However, $B_1 \sqsubseteq_c B_2$ does not hold. If the context contains a write W , then the read can either read from it or the initialisation. The former generates a hb-edge in the history, while the latter generates a deny from RFval-d – thus history inclusion does not hold. The problem is that the \sqsubseteq_c does not distinguish the position of W in hb, which removes precision from the denotation.

6. Verification Tool

In this section we describe Stellite⁵, a tool that checks our \sqsubseteq_c relation using the Alloy system [10].

Input transformations are written in a simple C-like language. Stellite supports transformations with atomic reads

⁵Stellite is a cobalt-chromium alloy designed for wear-resistance.

and writes, and SC fences. It does not yet support NA accesses, ll/sc or branching control-flow, but these would not present fundamental difficulties. The Alloy encoding would be similar, albeit with a downclosure on histories, and the increased search space arising from this downclosure.

Stellite converts an input transformation $B_2 \rightsquigarrow B_1$ into an Alloy* model encoding the check $B_1 \sqsubseteq_c B_2$. We use the higher-order Alloy* solver because standard Alloy cannot handle the existential quantification on histories in the definition of \sqsubseteq_c .

As previously noted [24], there is a natural fit between Alloy models and C/C++-style axiomatic memory models. The actions of an execution, as well as locations, values, and all other basic notions are modelled as Alloy objects. Values are modelled abstractly, tracking only equality or inequality, which avoids having to enumerate them explicitly. The relations hb, rf, sb etc. are modelled by Alloy relations. Finally, validity axioms are modelled by Alloy predicates.

If a counter-example is discovered, the execution and history of B_1 can be viewed using the Alloy model visualiser, which has a similar appearance to the diagrams in this paper. Understanding why no corresponding history of B_2 exists is left to the user, but is usually obvious. As \sqsubseteq_c is not fully abstract, this counter-example could of course be spurious.

The Alloy* search is parameterised by the maximum size of the model it will examine. For this reason, Stellite is a bounded checking tool: it only examines histories up to a specified size, measured in total number of actions. However, our cutting predicate bounds the size of histories that must be considered, meaning that a sufficiently large bound guarantees that tool results are sound. Given a query $B_1 \sqsubseteq_c B_2$, the bound needed varies depending on the number of internal actions on distinct locations in B_1 . In our experiments we ran the tool with size bound of 10, sufficient to give validity all the optimisations we consider. Note that most transformations do not require such a large bound, and execution times improve considerably if it is reduced.

Experimental results. Figure 3 gives results for different transformations, many derived from [23] – we test all the examples from [23] that fit into our input language. Transformations of the sort that we check have led to bugs in GCC [15] and LLVM [8].

Two of the transformations cause the tool to timeout. This does not establish validity, but as with other bounded model-checking work, our experience is that counter-examples are found at shallow positions in the search space.

Note that some transformations are invalid because of their effect on local variables, e.g. $\text{skip} \rightsquigarrow l := \text{read}(x)$. The closely related transformation $\text{skip} \rightsquigarrow \text{read}(x)$ throws away the result of the read, and is consequently valid.

7. Transformations with Non-Atomicity

We now extend our memory model and denotation to non-atomic (i.e. unsynchronised) accesses. In C11, these are

Here $X|_A$ is the projection of the execution X to actions in A . We lift the downclosure to sets of executions in the standard way. Now we define $B_1 \sqsubseteq_q^{NA} B_2$ as follows:

$$\begin{aligned}
B_1 \sqsubseteq_q^{NA} B_2 &\iff \forall R, S. \\
&\forall X_1 \in \llbracket B_1, \mathcal{A}, R, S \rrbracket_v^{NA}. \\
&\exists X_2' \in (\llbracket B_2, \mathcal{A}, R, S \rrbracket_v^{NA})^\downarrow. \\
&\exists X_1' \in (X_1)^\downarrow. \text{hist}(X_1') \sqsubseteq_h \text{hist}(X_2') \wedge \\
&\quad (\text{safe}(X_2') \implies \text{safe}(X_1) \wedge (X_1' = X_1) \\
&\quad \wedge X_2' \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_v^{NA})
\end{aligned}$$

There are essentially two cases in this definition: either the X_2' we choose is safe or unsafe.

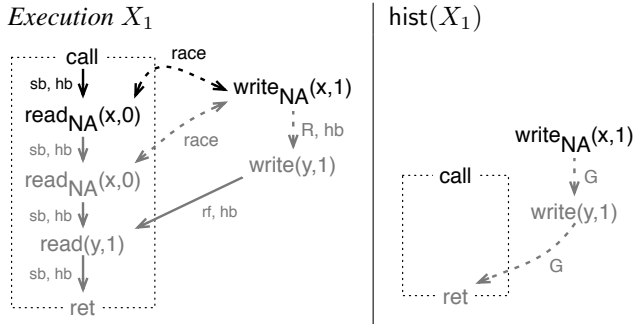
- X_2' **is safe**. In this case the situation resembles \sqsubseteq_q , i.e. X_2' is a complete, not prefixed, execution, and it has a related history to X_1 . If the code-block B_2 is guaranteed to be safe, for example because it has no non-atomic accesses, the above definition is equivalent to \sqsubseteq_q .
- X_2' **is unsafe**. In this case, any program $C(B_2)$ which has X_2' as a sub-execution has semantics \top . Therefore, we need only match histories *up to the point that X_2' becomes unsafe*. This ensures the unsafety will actually occur in the (unknown) whole program $C(B_2)$. In \sqsubseteq_q^{NA} , if X_2' unsafe, then it can be a prefix, not a complete execution. We then must witness a prefix X_1' of X_1 with a corresponding history. After this point, X_1 and X_2 can behave entirely differently.

This prefixing in the definition of \sqsubseteq_q^{NA} is required for full abstraction—it would be adequate to always require complete executions with related histories.

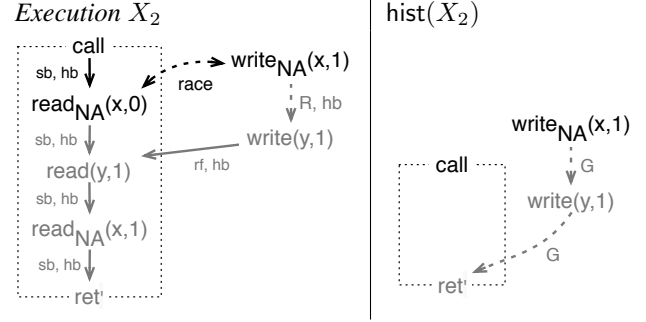
Validating a transformation. Consider the following *anti-roach-motel* transformation:

$$\begin{array}{ll}
B_2: 11 := \text{read}_{NA}(x); & \rightsquigarrow B_1: 11 := \text{read}_{NA}(x); \\
12 := \text{read}(y); & 13 := \text{read}_{NA}(x); \\
13 := \text{read}_{NA}(x); & 12 := \text{read}(y);
\end{array}$$

To verify the transformation, we must establish that $B_1 \sqsubseteq_q^{NA} B_2$. We illustrate the necessary reasoning for a single block-local execution $X_1 \in \llbracket B_1, \mathcal{A}, R, S \rrbracket$, with a context action set \mathcal{A} consisting of a non-atomic write of $x = 1$ and an atomic write of $y = 1$, and a context relation R relating the write of x to the write of y :



Note the data races between reads and a write over x . We can exhibit an unsafe execution $X_2 \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_v$:



The histories of the *complete* executions X_1 and X_2 differ in their return action. In X_2 The read of y takes the value of the context write, so COHERNA forces the second read of x to read from the context write of x . This changes the values of local variables recorded in ret' .

However, because X_2 is unsafe, we can select a prefix X_2' which includes the race (we denote in grey the parts that we do not include). Similarly, we can select a prefix X_1' of X_1 . We have that $\text{hist}(X_1') \sqsubseteq_h \text{hist}(X_2')$, even though the histories $\text{hist}(X_1)$ and $\text{hist}(X_2)$ do not correspond. Reasoning in this way shows that the transformation is valid.

Finite denotation with NA. We have also defined a finite variant of \sqsubseteq_q^{NA} , using the cutting strategy described in §5. We defer the details to §C.


8. Full Abstraction

We give a proof of full abstraction (Theorems 2 and 6) in §E. Our proof depends on constructing a special syntactic context that is sensitive to one particular history. Given an execution X produced from a block B , we write C_X for the resulting special context. The construction of C_X guarantees (1) that X is the block portion of an execution of $C_X(B)$; and (2) for any block B' , if $C_X(B')$ has a different block history from X , this is visible in different observable behaviour. Therefore for any blocks that are distinguished by different histories, our construction can produce a program with different observable behaviour, establishing full abstraction.


Context construction. The definition of C_X is given in §E – here we sketch its behaviour. C_X executes the context operations from X in parallel with the block. It wraps these operations in auxiliary wrapper code to enforce R and check the history. If wrapper code fails, it writes to an error variable, which thereby alters the observable behaviour.

The context must generate edges in R . This is enforced by wrappers $\text{Rrel}_{u,v}$ and $\text{Rac}_{u,v}$, one for each edge $(u, v) \in R$. Each wrapper uses a watchdog variable to create a hb-edge: if $\text{Rac}_{u,v}$ does not read the value written by $\text{Rrel}_{u,v}$, then the error variable is written. A successful read gives the

following shape:



The context must also prohibit history edges beyond those in the original guarantee G . This is checked by wrappers $\text{Nrel}_{u,v}$ and $\text{Nacq}_{u,v}$, one for each (u, v) not in G . Once again happens-before edges are detected using watchdog variables. If $\text{Nacq}_{v,u}$ *does* read the value written by $\text{Nrel}_{v,u}$, then there is an errant history edge, and the error location is written. An erroneous execution has the following shape (omitting the write to the error location):



Our context construction requires LL/SC, not just C11's standard RMW. This is because it is necessary to inject wrapper code between LL and SC in order to detect and enforce hb-edges. This is not possible with RMW alone.

9. Related Work

Our approach builds on Batty et al. [4], who generalise linearizability [9] to the C11 memory model. They represent interactions between a library and its clients by sets of histories consisting of a guarantee and a deny; we do the same for code and context. However, we cannot assume *information hiding*: Batty et al. assume that the variables used by the library cannot be directly accessed by clients. Also, we establish both adequacy and full abstraction, propose a finite denotation, and build an automated verification tool.

Our approach is broadly similar to Brookes's seminal concurrency semantics [6]. In both cases, a code block is represented by a denotation capturing possible interactions with an abstracted context. In [6], denotations are sets of traces, consisting of sequences of global program states; context actions are represented by changes in these states. To handle the more complex axiomatic memory model, our denotation consists of sets of context actions and relations on them, with context actions explicitly represented as such.

In order to achieve full abstraction, Brookes assumes a powerful `await()` instruction which blocks until the global state satisfies a boolean expression. Our full abstraction result does not require this: our strongest instruction is LL-SC, which is commonly available on hardware platforms.

Brookes-like approaches have been applied to several relaxed models: operational hardware models [7]; TSO [11]; and SC-DRF [17]. Also, [7, 17] define tools for verifying program transformations. All three approaches are based on traces, and are therefore not directly portable to C11-style axiomatic models. All three also target substantially stronger (i.e. more restrictive) relaxed models than C11.

Methods for verifying code transformations, either manually or using proof assistants, have been proposed for several relaxed models: TSO [19, 20, 22], Java [18] and C/C++ [23]. These methods are non-compositional in the sense that verifying a transformation requires considering the trace set of the entire program — there is no abstraction of the context. We abstract both the sequential and concurrent context and thereby support automated verification. The above methods also model transformations as rewrites on program executions, whereas we treat them directly as modifications of program syntax; the latter corresponds more closely to actual compilers. There has also been various work on automatically verifying compiler optimisations under sequential consistency [13, 16, 25].

Several tools exist for testing (not verifying) program transformations on axiomatic models: for example Morisset et al. [12] for GCC, and Chakraborty et al. [8] for LLVM. Similarly, Alloy is used in the MemSAT tool for simulation of the Java memory model [21]. Our Alloy encoding of the memory model is also similar to the input files for the Herd/Cat memory model simulator [2].

10. Conclusions and Extensions

We have proposed the first compositional method for verifying program transformation on a C/C++-style memory model. Our method only requires reasoning about the code-block being transformed in minimal contexts, which simplifies proofs and allows push-button verification of transformations. We have defined two variants on our approach with different motivations: one that achieves maximal precision by guaranteeing full abstraction and one that supports automatic verification by guaranteeing finiteness.

Extensions. We omit SC and RLX accesses, but both are handled in [4]. We believe very similar techniques could be applied. For SC this would require an extra relation in the history. For RLX, this requires a non-compositional safety check. Note that [4] gives RLX an unrealistically strong semantics as a stopgap solution the thin-air problem.

Our language is based on shared variables, not addressable memory, so for example we cannot write $y := *x$; $z := *y$. To support addressable memory, the block-local execution construction would need to quantify over actions on all memory locations, not just a static variable set VS . The rest of our theory would remain the same however, because C11-style models grant no special status to pointer values. Cutting down to a finite denotation would then require some extra abstraction over memory.

References

- [1] *Programming Languages — C++*. 2011. ISO/IEC JTC1 SC22 WG21.
- [2] J. Alglave, L. Maranget, and M. Tautschnig. Herding cats: modelling, simulation, testing, and data-mining for weak memory. In *PLDI*, page 7, 2014.
- [3] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL*, pages 55–66, 2011.
- [4] M. Batty, M. Dodds, and A. Gotsman. Library abstraction for C/C++ concurrency. In *POPL*, pages 235–248, 2013.
- [5] M. Batty, K. Memarian, K. Nienhuis, J. Pichon-Pharabod, and P. Sewell. The problem of programming language concurrency semantics. In *ESOP*, pages 283–307, 2015.
- [6] S. D. Brookes. Full abstraction for a shared-variable parallel language. *Inf. Comput.*, 127(2):145–163, 1996.
- [7] S. Burckhardt, M. Musuvathi, and V. Singh. Verifying local transformations on relaxed memory models. In *CC*, 2010.
- [8] S. Chakraborty and V. Vafeiadis. Validating optimizations of concurrent C/C++ programs. In *CGO*, pages 216–226, 2016.
- [9] M. Herlihy and J. M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.
- [10] D. Jackson. *Software Abstractions – Logic, Language, and Analysis*. MIT Press, revised edition, 2012.
- [11] R. Jagadeesan, G. Petri, and J. Riely. Brookes is relaxed, almost! In *FOSSACS*, pages 180–194, 2012.
- [12] O. Lahav, N. Giannarakis, and V. Vafeiadis. Taming release-acquire consistency. In *POPL*, pages 649–662, 2016.
- [13] N. P. Lopes, D. Menendez, S. Nagarakatte, and J. Regehr. Provably correct peephole optimizations with Alive. In *PLDI*, pages 22–32, 2015.
- [14] A. Milicevic, J. P. Near, E. Kang, and D. Jackson. Alloy*: A general-purpose higher-order relational constraint solver. In *ICSE*, pages 609–619, 2015.
- [15] R. Morisset, P. Pawan, and F. Zappa Nardelli. Compiler testing via a theory of sound optimisations in the C11/C++11 memory model. In *PLDI*, pages 187–196, 2013.
- [16] E. Mullen, D. Zuniga, Z. Tatlock, and D. Grossman. Verified peephole optimizations for CompCert. In *PLDI*, pages 448–461, 2016.
- [17] D. Poetzl and D. Kroening. Formalizing and checking thread refinement for data-race-free execution models. In *TACAS*, 2016.
- [18] J. Sevcík and D. Aspinall. On validity of program transformations in the Java memory model. In *ECOOP*, pages 27–51, 2008.
- [19] J. Sevcík, V. Vafeiadis, F. Zappa Nardelli, S. Jagannathan, and P. Sewell. Relaxed-memory concurrency and verified compilation. In *POPL*, pages 43–54, 2011.
- [20] J. Sevcík, V. Vafeiadis, F. Zappa Nardelli, S. Jagannathan, and P. Sewell. CompCertTSO: A verified compiler for relaxed-memory concurrency. *J. ACM*, 60(3):22:1–22:50, 2013.
- [21] E. Torlak, M. Vaziri, and J. Dolby. MemSAT: checking axiomatic specifications of memory models. In *PLDI*, pages 341–350, 2010.
- [22] V. Vafeiadis and F. Zappa Nardelli. Verifying fence elimination optimisations. In *SAS*, pages 146–162, 2011.
- [23] V. Vafeiadis, T. Balabonski, S. Chakraborty, R. Morisset, and F. Zappa Nardelli. Common compiler optimisations are invalid in the C11 memory model and what we can do about it. In *POPL*, pages 209–220, 2015.
- [24] J. Wickerson, M. Batty, T. Sorensen, and G. A. Constantinides. Automatically comparing memory consistency models. In *POPL*, 2017.
- [25] J. Zhao, S. Nagarakatte, M. M. K. Martin, and S. Zdancewic. Formal verification of SSA-based optimizations for LLVM. In *PLDI*, pages 175–186, 2013.

A. Collected Definitions

Execution observational refinement

$$X \preceq_{\text{ex}} Y \iff \mathcal{A}(X|_{\text{OVar}}) = \mathcal{A}(Y|_{\text{OVar}}) \wedge \text{hb}(Y|_{\text{OVar}}) \subseteq \text{hb}(X|_{\text{OVar}})$$

Program observational refinement

$$P_1 \preceq_{\text{pr}} P_2 \iff \forall X_1 \in \llbracket P_1 \rrbracket. \exists X_2 \in \llbracket P_2 \rrbracket. X_1 \preceq_{\text{ex}} X_2$$

Program observational refinement with NA

$$P_1 \preceq_{\text{pr}}^{\text{NA}} P_2 \iff (\text{safe}(P_2) \implies \text{safe}(P_1) \wedge P_1 \preceq_{\text{pr}} P_2)$$

Block observational refinement

$$B_1 \preceq_{\text{bl}} B_2 \iff \forall C. C(B_1) \preceq_{\text{pr}} C(B_2)$$

History abstraction

$$(A_1, G_1) \sqsubseteq_{\text{h}} (A_2, G_2) \iff A_1 = A_2 \wedge G_2 \subseteq G_1$$

Quantified abstraction

$$B_1 \sqsubseteq_{\text{q}} B_2 \iff \forall \mathcal{A}, R, S. \forall X_1 \in \llbracket B_1, \mathcal{A}, R, S \rrbracket. \exists X_2 \in \llbracket B_2, \mathcal{A}, R, S \rrbracket. \text{hist}(X_1) \sqsubseteq_{\text{h}} \text{hist}(X_2)$$

Extended history abstraction

$$(A_2, G_2, D_2) \sqsubseteq_{\text{E}} (A_1, G_1, D_1) \iff A_1 = A_2 \wedge G_2 \subseteq G_1 \wedge D_2 \subseteq D_1$$

Cut abstraction

$$B_1 \sqsubseteq_{\text{c}} B_2 \iff \forall \mathcal{A}, S. \forall X_1 \in \llbracket B_1, \mathcal{A}, \emptyset, S \rrbracket. \text{cut}(X_1) \implies \exists X_2 \in \llbracket B_2, \mathcal{A}, \emptyset, S \rrbracket. \text{hist}_{\text{E}}(X_1) \sqsubseteq_{\text{E}} \text{hist}_{\text{E}}(X_2)$$

Cut predicates

$$\text{vis}(X) \iff \text{code}(X) \cup \{u \mid \exists v \in \text{code}(X). u \xrightarrow{\text{rf}} v \vee v \xrightarrow{\text{rf}} u\}$$

$$\text{cut}(X) \iff \text{cutR}(X) \wedge \text{cutW}(X)$$

$$\text{cutR}(X) \iff \text{reads}(X) \subseteq \text{vis}(X) \wedge \forall r_1, r_2 \in \text{contx}(X). r_1 \neq r_2 \implies \neg \exists w. w \xrightarrow{\text{rf}} r_1 \wedge w \xrightarrow{\text{rf}} r_2$$

$$\text{cutW}(X) \iff \forall w_1, w_2 \in (\text{contx}(X) \setminus \text{vis}(X)). w_1 \xrightarrow{\text{mo}} w_2 \implies \exists w_3 \in \text{vis}(X). w_1 \xrightarrow{\text{mo}} w_3 \xrightarrow{\text{mo}} w_2$$

Execution downclosure

$$X^\downarrow \triangleq \{X' \mid \exists A. X' = X|_A \wedge \forall (a, a') \in (\text{hb}(X) \cup \text{rf}(X))^+. a' \in A \implies a \in A\}$$

Quantified abstraction with NA

$$\begin{aligned} B_1 \sqsubseteq_{\text{q}}^{\text{NA}} B_2 \iff & \forall R, S. \forall X_1 \in \llbracket B_1, \mathcal{A}, R, S \rrbracket_v^{\text{NA}}. \exists X'_2 \in (\llbracket B_2, \mathcal{A}, R, S \rrbracket_v^{\text{NA}})^\downarrow. \\ & \exists X'_1 \in (X_1)^\downarrow. \text{hist}(X'_1) \sqsubseteq_{\text{h}} \text{hist}(X'_2) \wedge \\ & (\text{safe}(X'_2) \implies \text{safe}(X_1) \wedge (X'_1 = X_1) \wedge X'_2 \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_v^{\text{NA}}) \end{aligned}$$

Thread-local semantics. We write $\mathcal{A}_1 \cup \mathcal{A}_2$ for a union that is defined only when actions in \mathcal{A}_1 and \mathcal{A}_2 use disjoint sets of identifiers.

$$\begin{aligned}
\langle \text{write}(x, l), \sigma \rangle &\triangleq \{(\{\text{write}(x, a)\}, \emptyset, \sigma) \mid \sigma(l) = a\} \\
\langle l := \text{read}(x), \sigma \rangle &\triangleq \{(\{\text{read}(x, a)\}, \emptyset, \sigma[l \mapsto a]) \mid a \in \text{Val}\} \\
\langle l := \text{RMW}(x, l_1, l_2), \sigma \rangle &\triangleq \{(\{r, w\}, (r, w), \sigma[l \mapsto 1]) \mid \sigma(l_1) = a \wedge \sigma(l_2) = b \wedge r = \text{ll}(x, a) \wedge w = \text{sc}(x, b)\} \cup \\
&\quad \{(\{\text{read}(x, a)\}, \emptyset, \sigma[l \mapsto 0]) \mid \sigma(l_1) \neq a\} \\
\langle l := \text{LL}(x), \sigma \rangle &\triangleq \{(\{\text{ll}(x, a)\}, \emptyset, \sigma[l \mapsto a]) \mid a \in \text{Val}\} \cup \{(\{\text{ll}_f(x, a)\}, \emptyset, \sigma[l \mapsto a]) \mid a \in \text{Val}\} \\
\langle l' := \text{SC}(x, l), \sigma \rangle &\triangleq \{(\{\text{sc}(x, a)\}, \emptyset, \sigma[l' \mapsto 1]) \mid \sigma(l) = a\} \cup \{(\{\text{sc}_f(x)\}, \emptyset, \sigma[l' \mapsto 0])\} \\
\langle \text{fence}, \sigma \rangle &\triangleq \{(\{r, w\}, (r, w), \sigma) \mid r = \text{ll}(\text{fen}, 0) \wedge w = \text{sc}(\text{fen}, 0)\} \\
\langle C_1 \parallel C_2, \sigma \rangle &\triangleq \{(\mathcal{A}_1 \cup \mathcal{A}_2, \text{sb}_1 \cup \text{sb}_2, \sigma) \mid (\mathcal{A}_1, \text{sb}_1, \sigma_1) \in \langle C_1, \sigma \rangle \wedge (\mathcal{A}_2, \text{sb}_2, \sigma_2) \in \langle C_2, \sigma \rangle\} \\
\langle C_1; C_2, \sigma \rangle &\triangleq \{(\mathcal{A}_1 \cup \mathcal{A}_2, \text{sb}_1 \cup \text{sb}_2 \cup (\mathcal{A}_1 \times \mathcal{A}_2), \sigma_2) \mid \\
&\quad (\mathcal{A}_1, \text{sb}_1, \sigma_1) \in \langle C_1, \sigma \rangle \wedge (\mathcal{A}_2, \text{sb}_2, \sigma_2) \in \langle C_2, \sigma_1 \rangle\} \\
\langle \text{if}(l) \{C_1\} \text{else} \{C_2\}, \sigma \rangle &\triangleq \begin{cases} \langle C_2, \sigma \rangle, & \text{if } \sigma(l) = 0 \\ \langle C_1, \sigma \rangle, & \text{otherwise} \end{cases}
\end{aligned}$$

B. Proof of Theorem 1

We now prove adequacy of $\sqsubseteq_q^{\text{NA}}$. As $\sqsubseteq_q^{\text{NA}} \implies \sqsubseteq_q$, this suffices to prove adequacy of \sqsubseteq_q . Our proof need several auxiliary notions:

- $\text{codeE}(X)$ is the projection of an execution X to actions in $(\text{codeE}(X) \cup \text{interf}(X) \cup \{\text{call}, \text{ret}\})$.
- The *interface actions* are actions on variables in VS_B that occur in the context. These are context actions that can affect the behaviour of the code-block. We write $\text{interf}(X)$ for this set.
- $\text{contxE}(X)$ is the projection of an execution X to the context. This is a more complex projection than $\text{codeE}(X)$ because it removes mo and rf over actions in $\text{interf}(X)$. Let $\mathcal{I} = \text{contxE}(X) \cup \{\text{call}, \text{ret}\}$ and $\mathcal{C} = \text{contxE}(X) \setminus \text{interf}(X)$. Then

$$\text{contxE}(X) = (A(X)|_{\mathcal{I}}, \text{hb}(X)|_{\mathcal{I}}, \text{sb}(X)|_{\mathcal{I}}, \text{mo}(X)|_{\mathcal{C}}, \text{rf}(X)|_{\mathcal{C}})$$

- $\text{hbC}(X)$ is the context-side projection of hb to interface actions. In other words, the projection of $\text{hb}(X)$ to pairs in:

$$(\text{interf}(X) \times \text{interf}(X)) \cup (\text{interf}(X) \times \{\text{call}\}) \cup (\{\text{ret}\} \times \text{interf}(X))$$

- $\text{atC}(X)$ is the context-side projection of at to context actions: i.e. the projection of $\text{at}(X)$ to pairs in $(\text{interf}(X) \times \text{interf}(X))$.
- $\llbracket C, R, S \rrbracket_v$ is the *context-local* execution of a single-hole context C – this is an analogous notion to the block-local execution, except that rf and mo are not generated for the interface. Here R is a relation representing dependencies in hb arising from the code and S represents code at edges. An execution X is in this set iff:

- R is a code-side relation on interface actions $\text{interf}(X)$:

$$R \subseteq (\text{interf}(X) \times \text{interf}(X)) \cup (\text{interf}(X) \times \{\text{ret}\}) \cup (\{\text{call}\} \times \text{interf}(X))$$

- S is a code-side relation on interface actions $\text{interf}(X)$:

$$S \subseteq (\text{interf}(X) \times \text{interf}(X))$$

- The execution satisfies the thread-local semantics:

$$(A(X), \text{sb}(X)) \in \langle C \rangle$$

We assume that a singleton hole has the following thread-local semantics:

$$\langle \{-\}, \sigma \rangle \triangleq \{(\{c, r\}, \{c \rightarrow r\}, \sigma') \mid c = \text{call}(\sigma) \wedge r = \text{ret}(\sigma')\}$$

- X satisfies HBDEF' , ATOM' , NOFAIL' , ACYCLICITY , RFWF , HBVSMO , COHERENCE , RFHBNA , COHERNA .

- The projection $X|_{\text{contxE}(X)}$ satisfies RFVAL, MOWF. mo and rf are disjoint from actions in $\text{interf}(X)$.

We sometimes write $\llbracket C \rrbracket_v$ to stand for $\llbracket C, \emptyset, \emptyset \rrbracket$, i.e. the set of context-local executions with empty code-side relations.

LEMMA 7 (Decomposition). *Assume $X \in \llbracket C(B) \rrbracket_v$, and no there are no at edges in C spanning B , nor any between the actions of B and C . Then $\text{codeE}(X) \in \llbracket B, \text{interf}(X), \text{hbC}(X), \text{atC}(X) \rrbracket_v$ and $\text{contxE}(X) \in \llbracket C, \text{hbL}(X), \text{atL}(X) \rrbracket_v$.*

Proof (code). We have several proof obligations.

- $\text{hbC}(X)$ and $\text{atC}(X)$ are context-side relations on interface actions (trivial by definition).
- $(\text{codeE}(\text{codeE}(X)), \text{sb}(\text{codeE}(X))) \in \langle B \rangle$, i.e. the execution satisfies the thread-local semantics.
- The actions in $\text{codeE}(\text{codeE}(X))$ are in between a call / ret pair in sb. We assume we can introduce call / ret freely to satisfy this requirement.
- $\text{codeE}(X)$ satisfies the validity axioms for a block-local execution – note that this replaces HBDEF with HBDEF', ATOM with ATOM' and NOFAIL with NOFAIL'.

For the first obligation, we argue inductively over the structure of C . First assume that $C = \{-\}$, i.e. C consists only of a hole. In this case the result holds immediately from the thread-local semantics. For the inductive case, assume C is a composite one-hole context, e.g. $C_1; C_2(-) / C_1(-); C_2 / C_1 \parallel C_2(-) /$ etc.

For the fourth obligation, we prove $\text{codeE}(X)$ satisfies the validity axioms by arguing in turn about each. Assume the following shorthand:

$$\text{codeE}(X) = (A(l), \text{hb}(l), \text{at}(l)\text{sb}(l), \text{mo}(l), \text{rf}(l))$$

HBDEF': Let $R = \text{hbC}(X)$. Now prove in both directions:

$$(a, b) \in \text{hb}(l) \implies (a, b) \in (\text{sb}(l) \cup \text{rf}_{\text{AT}}(l) \cup R)^+ \quad (8)$$

$$(a, b) \in (\text{sb}(l) \cup \text{rf}_{\text{AT}}(l) \cup R)^+ \implies (a, b) \in \text{hb}(l) \quad (9)$$

For the first case, any (a, b) in $\text{hb}(l)$ must have code or interface actions at both ends, and must have originated from a path $(a, b) \in (\text{sb}(X) \cup \text{rf}_{\text{AT}}(X))^+$. By construction, there are no rf-edges between $\text{codeE}(X)$ and $\text{contxE}(X)$. Therefore, portions of the path which stray into the context must enter and leave through call, ret, or actions in $\text{interf}(X)$. These portions of the path will be summarised by $\text{hbC}(X)$. As a result, for any such path, there must be an equivalent path $(a, b) \in (\text{sb}(l) \cup \text{rf}_{\text{AT}}(l) \cup \text{hbC}(X))^+$.

For the second case, we make a similar argument. For any pair $(c, d) \in \text{hbC}(X)$, there must be a path $(c, d) \in (\text{sb}(X) \cup \text{rf}_{\text{AT}}(X))$. As a consequence, for any (a, b) in $(\text{sb}(l) \cup \text{rf}(l) \cup \text{hbC}(X))^+$, there must be a path $(a, b) \in (\text{sb}(X) \cup \text{rf}_{\text{NA}}(X))$. Thus $(a, b) \in \text{hb}(X)$. As $\text{hb}(l)$ is a projection of $\text{hb}(X)$, this completes the proof.

ATOM', NOFAIL', ACYCLICITYRFWF, MOWF, COHERENCE, RFHBNA, COHERNA: all hold immediately by the fact that $\text{codeE}(X)$ is a projection of X .

RFVAL: holds because $\text{code}(X)$ contains exactly the actions in X that are on locations $a \in \text{gv}_B$. Therefore, the projection cannot remove the origin write for a read. □

Proof (context). Similar argument to the code. □

LEMMA 8 (Completion lemma). *Let X be an execution. If $\text{valid}(X)$ and $(A(X), \text{sb}(X)) \in \langle Q \rangle^\downarrow$, then $X \in \llbracket Q \rrbracket_v^\downarrow$.*

Proof. We require the existence of a $Y \in \llbracket Q \rrbracket_v$ such that $X \in Y^\downarrow$. To prove this, we iteratively extend X by adding sb-final actions, and show that the new execution can in each case be made valid. As all executions are finite, this proves the result.

Assume the current execution is X_i . We choose an $A(X_{i+1})$ and $\text{sb}(X_{i+1})$ such that the new execution is extended by a single sb-final action, and that $(A(X_{i+1}), \text{sb}(X_{i+1})) \in \langle Q \rangle^\downarrow$. We now need to show that we can construct a valid X_{i+1} .

Case-split on the type of the new action. Non-atomics read from their immediate hb predecessor, or the init value if none exists. Atomic reads read from the end of mo, and writes can be added to the end of mo. Read-modify-writes read from the end of mo. All of these cases preserve the validity axioms.

Note that if the new action is a read, we may need to fix its value depending on an earlier write. This depends on the property of *receptiveness* – given a prefix $(A, \text{sb}) \in \langle Q' \rangle$ and a read r that is sb-maximal, any value can be given to the read. This property follows from the thread-local semantics: the only tricky cases are conditionals and LL-SC, where receptiveness is guaranteed by the fact that any possible value is represented in the set of possible reads. □

LEMMA 9 (Safety completion). *Let X, Y be valid executions. $\neg\text{safe}(X)$ and $X \in Y^\downarrow$ implies $\neg\text{safe}(Y)$.*

Proof. Prove the contrapositive: $\text{safe}(Y) \implies \text{safe}(X)$. This holds immediately from the fact that in a safe execution, potentially racy actions must be related in hb. \square

LEMMA 10 (Composition). *Let X and Y be executions such that $X \in \llbracket B, \mathcal{A}, \text{hbC}(Y), \text{atC}(Y) \rrbracket_v^\downarrow$ and $Y \in \llbracket C, \mathcal{A}, R', S' \rrbracket_v^\downarrow$ with no ll/sc pairs crossing the block boundary in each case, with $\text{hist}(Y) \sqsubseteq_{\text{h}} \text{hist}(X)$ and with $\text{atL}(X) = S'$. Then there exists an execution Z such that $Z \in \llbracket C(B) \rrbracket_v^\downarrow$. Furthermore:*

- If $\neg\text{safe}(X)$ or $\neg\text{safe}(Y)$, then $\neg\text{safe}(Z)$.
- If $\text{safe}(X)$, $\text{safe}(Y)$, and $\text{safe}(Z)$, and $X \in \llbracket B, \mathcal{A}, \text{hbC}(Y), \text{atC}(Y) \rrbracket_v$ and $Y \in \llbracket C, \mathcal{A}, R', S' \rrbracket_v$, then $Z \in \llbracket C(B) \rrbracket_v$ and $\text{ctxxE}(Y) \preceq_{\text{ex}} \text{ctxxE}(Z)$.

Proof. We begin by defining Z . Taking each term of the execution in turn:

- The action set $A(Z)$ is the union of the two action sets $A(X)$ and $A(Y)$, merging call, return and interface actions.
- $\text{sb}(Z) = (\text{sb}(X) \cup \text{sb}(Y))^+$.
- $\text{mo}_Z = (\text{mo}(X) \cup \text{mo}(Y))$ – as the two mo relations are disjoint, no transitive closure is needed.
- $\text{rf}_Z = (\text{rf}(X) \cup \text{rf}(Y))$ – likewise.
- $\text{hb}_Z = (\text{sb}(Z) \cup \text{rf}_{\text{AT}}(Z))^+$, ie, according to HBDEF.
- $\text{at}_Z = \text{at}(X) \cup \text{at}(Y)$.

We first need to show that $Z \in \llbracket C(L) \rrbracket_v^\downarrow$. To do this we use the completion lemma: thus our proof obligations are $(A(Z), \text{sb}(Z)) \in \langle C(B_2) \rangle^\downarrow$ and $\text{valid}(Z)$.

We observe that that $(A(Z), \text{sb}(Z)) \in \langle C(B_2) \rangle^\downarrow$ is obvious from the thread-local semantics.

Next prove that $\text{valid}(Z)$. HBDEF holds by construction. RFWF, RFVAL, MOWF, RBDEF are true trivially as for each variable, validity is checked solely in either the code or context. This leaves ACYCLICITY, HBVSMO, COHERENCE, COHERNA, ATOM and NOFAIL. (RFHBNA needs to be treated specially – see below).

- For ACYCLICITY, a violation would correspond to a path in $(\text{sb}(Z) \cup \text{rf}_{\text{AT}}(Z) \cup \text{rf}_{\text{NA}})^+$. As this path cannot appear in either X or Y , it must cross between the two: each point where it does so must be an interface action or call / return. As a result, a corresponding violation can be constructed in X .
Call-to-return paths are in $(\text{sb}(X) \cup \text{rf}_{\text{AT}}(X))^+ \cup \text{rf}_{\text{NA}}(X))^+$. Conversely, return-to-call paths are in $(\text{sb}(Y) \cup \text{rf}_{\text{AT}}(Y) \cup \text{rf}_{\text{NA}}(Y))^+$. As Y satisfies RFHBNA, $\text{rf}_{\text{NA}}(Y) \in \text{hb}(Y)$. Thus the return-to-call portions of the path are in $\text{hbC}(Y)$. This contradicts the assumption that X satisfies ACYCLICITY.
- For HBVSMO, a violation consists of a write pair w_1, w_2 such that $(w_1, w_2) \in \text{hb}(Z)$ and $(w_2, w_1) \in \text{mo}(Z)$. As mo is partitioned between code and context, either both writes are in X or both in Y . By assumption, the violation is not solely in X or Y , so the path from w_1 to w_2 in $(\text{sb} \cup \text{rf}_{\text{AT}})^+$ must contain a sequence of interface actions or call / return.
 1. If the writes are in X , then mo is replicated immediately. The block-local portions of the path are in $(\text{sb}(X) \cup \text{rf}_{\text{AT}}(X))^+$, while the context-local portions are in $\text{hbC}(X)$. Thus we can replicate the violation.
 2. If the writes are in Y , we can use a similar argument. However, we also appeal to the fact that $\text{hist}(Y) \sqsubseteq_{\text{h}} \text{hist}(X)$, which means that $\text{hbL}(X) \subseteq \text{hbL}(Y)$. This means that any code-side hb edge in X can be replicated in Y to recreate the violation.
- For COHERENCE and COHERNA, we note that rf and mo are partitioned between X and Y . Therefore we can apply the same argument as for the previous axioms to show the hb edges for a violation must exist in either X or Y .
- Similarly, for ATOM and NOFAIL, we note that at is partitioned between X and Y so any violation must exist in either X or Y .

Finally, we consider RFHBNA. As $\text{hist}Y \sqsubseteq_{\text{h}} \text{hist}X$, composing the two may weaken hb and generate violations on the context side. To solve this, we convert the RFHBNA violation to a safety violation. Take a $Z' \in Z^\downarrow$ such that there is a single $(\text{hb} \cup \text{rf})$ -final RFHBNA violation. We redirect the origin of this read to its immediate hb-predecessor, or the initialisation value if this does not exist. This gives an execution Z'' which satisfies RFHBNA, but violated DRF. All the other validity axioms are preserved under prefixing, so by the completion lemma, $Z'' \in \llbracket C(B) \rrbracket_v^\downarrow$. We use Z'' as our constructed execution.

We now need to show that $\neg\text{safe}(X)$ or $\neg\text{safe}(Y)$ implies $\neg\text{safe}(Z)$. If we had to fix an RFHBNA violation, the new execution Z'' is unsafe by construction. Otherwise, composition can only weaken hb, meaning any violation is trivially replicated.

Conversely, we need to show that if $\text{safe}(X)$, $\text{safe}(Y)$, and $\text{safe}(Z)$, and $X \in \llbracket B, \mathcal{A}, \text{hbC}(Y), \text{atC}(Y) \rrbracket_v$ and $Y \in \llbracket C, \mathcal{A}, R', S' \rrbracket_v$, then $Z \in \llbracket C(B) \rrbracket_v$ and $\text{contxE}(Y) \preceq_{\text{ex}} \text{contxE}(Z)$. As Z is safe, we know we did not have to fix a RFHBNA violation. For the rest of the proof, the same arguments as above give us a valid execution $Z \in \llbracket C(B) \rrbracket_v$.

It remains to show that $\text{contxE}(Y) \preceq_{\text{ex}} \text{contxE}(Z)$. Inclusion of context actions follows from the construction of Z . Inclusion on hb follows from the fact that $\text{hist}(Y) \sqsubseteq_{\text{h}} \text{hist}(X)$. Thus the composition can only weaken hb over context actions. \square

THEOREM 11 (Adequacy). $B_1 \sqsubseteq_{\text{q}}^{\text{NA}} B_2 \implies B_1 \preceq_{\text{bl}} B_2$ for blocks that include only matched ll/sc pairs.

Proof. Our objective from the definition of \preceq_{bl} is the following property:

$$\forall C, V. \neg \text{safe}(C(B_2)) \vee (\text{safe}(C(B_1)) \wedge \forall X \in \llbracket C(B_1) \rrbracket_v. \exists Y \in \llbracket C(B_2) \rrbracket_v. X|_V \preceq_{\text{ex}} Y|_V)$$

Begin the proof by picking an arbitrary C, V . The proof then proceeds by the normal steps: decomposition, abstraction, then composition.

- Case-split on whether $C(B_2)$ is safe or unsafe. If unsafe, we are done immediately. Therefore we can assume $\text{safe}(C(B_2))$.
- Pick an arbitrary execution $X \in \llbracket C(B_1) \rrbracket_v$.
- Apply the decomposition lemma to show that $\text{contxE}(X) \in \llbracket C, \text{hbL}(X), \text{atL}(X) \rrbracket_v$ and $\text{codeE}(X) \in \llbracket B_1, \text{hbC}(X), \text{atC}(X) \rrbracket_v$.
- Expand the definition of $\sqsubseteq_{\text{q}}^{\text{NA}}$, and pick $R = \text{hbC}(X)$ and $S = \text{atC}(X)$. This gives us executions $Y \in \llbracket B_2, \mathcal{A}, \text{hbC}(X), \text{atC}(X) \rrbracket_v^\downarrow$ and $X' \in \text{codeE}(X)^\downarrow$ such that:

$$\begin{aligned} \text{hist}(X') \sqsubseteq_{\text{h}} \text{hist}(Y) \wedge \\ \text{safe}(Y) \implies (\text{safe}(X') \wedge (X' = \text{codeE}(X)) \wedge Y \in \llbracket B_2, \mathcal{A}, \text{hbC}(X) \rrbracket_v) \end{aligned}$$

- Case-split on whether $\text{safe}(Y) \wedge \text{safe}(\text{contxE}(X))$ holds. If not, then apply the composition lemma to build an execution $Z \in \llbracket C(B_2) \rrbracket_v^\downarrow$ such that $\neg \text{safe}(Z)$. By lemma 9, there must exist a $Z' \in \llbracket C(B_2) \rrbracket_v$ such that $\neg \text{safe}(Z')$, which contradicts our assumption that $C(B_2)$ is safe.

Conversely, suppose $\text{safe}(Y) \wedge \text{safe}(\text{contxE}(X))$ holds. In this case, we apply the context lemma to build a $Z \in \llbracket C(B_2) \rrbracket_v$ such that $\text{contxE}(X) \preceq_{\text{ex}} \text{contxE}(Z)$. All actions on observable variables in V must be in the context, which means that $X|_V \preceq_{\text{ex}} Z|_V$ must also hold.

It remains to prove that $\text{safe}(X)$ holds. First we observe that $\text{safe}(\text{codeE}(X))$ holds by the abstraction theorem. As $\text{safe}(\text{contxE}(X))$ also holds, the result follows immediately. \square

C. Non-atomics and Denies

We now define $\sqsubseteq_{\text{c}}^{\text{NA}}$, a refinement between denotations which includes both cutting and non-atomics.


To do this we first need extra deny shapes. In the following, the variables obey the following constraint:

$$u, a, c \in \text{ret} \cup \text{interf}(X) \quad v, b, d \in \text{call} \cup \text{interf}(X)$$

All the actions a, b, c, d, u, v are pairwise distinct. Note that some of the hb-edges are transitively closed, meaning that syntactically distinct actions might be the same – e.g. w_1 and u in HBvsMO-d.

CoNA-d(u, v): $\exists w_1, w_2, r \in \text{NA}.$

The diagram illustrates two execution paths from w_1 to r . In the left path, w_1 is connected to w_2 by a solid black arrow labeled 'hb'. From w_2 , a solid black arrow labeled 'hb*' points down to u . From u , a dotted black arrow labeled 'D' points down to v . From v , a solid black arrow labeled 'hb*' points down to r . A red fox arrow labeled 'rf' points from w_1 down to r . In the right path, w_1 is connected to u by a solid black arrow labeled 'hb*'. From u , a dotted black arrow labeled 'D' points to v . From v , a solid black arrow labeled 'hb*' points to w_2 . From w_2 , a solid black arrow labeled 'hb*' points down to r . A red fox arrow labeled 'rf' points from w_1 down to r . The two paths are separated by a large 'V' symbol.



As before, we need a few notions to define the deny theorem.

- $\text{denyL}(X)$ contains all the binary denies:

$$\text{denyL}(X) \triangleq \text{HBvsMO-d} \cup \text{CoWR-d} \cup \text{Init-d} \cup \text{CoNA-d}$$

- $\text{denyNA}(X)$ contains the quaternary denies: $\text{denyNA}(X) \triangleq \text{CoNA-d2}$
- $\text{guarNA}(X)$ is the projection of $(\text{rf}_{\text{NA}} \cup \text{hb})^+$ to pairs in

$$(\text{interf}(X) \times \text{interf}(X)) \cup (\text{interf}(X) \times \{\text{ret}\}) \cup (\{\text{call}\} \times \text{interf}(X))$$

- Let \mathcal{I} be the set of actions $\text{interf}(X) \cup \{\text{call}, \text{ret}\}$. The *augmented history* of X , written $\text{hist}_E(X)$, is defined as

$$\text{hist}_E(X) \triangleq (A(X)|_{\mathcal{I}}, \text{hbL}(X), \text{denyL}(X), \text{guarNA}(X), \text{denyNA}(X))$$

- Two augmented histories, $H = (A, G, D, M, N)$, $H' = (A', G', D', M', N')$ are related $H \sqsubseteq_h H'$ iff

$$A = A' \wedge G' \subseteq G \wedge D' \subseteq D \wedge M' \subseteq M \wedge N' \subseteq N$$

- $\text{FinalNA}(X, a)$ holds if the action a is (1) an NA action, and (2) is the hb-final action in the code block in X .
- $\text{hbA}(X, a)$, for an execution X and action a is the projection of $\text{hb}(X)$ to pairs in $(\{a\} \times \text{interf}(X)) \cup (\text{interf}(X) \times \{a\})$
- The comparison $a \leq_{\text{na}}^{X, Y} b$ ensures that b participates in a race if a does. Formally, the comparison holds if: (1) a and b are actions on the same location; (2) b is a write if a is a write; and (3) $\text{hbA}(Y, b) \subseteq \text{hbA}(X, a)$. The final condition is needed to ensure that history edges cannot spuriously prevent a race in Y .
- The set of *almost-valid* executions $\llbracket B, \mathcal{A}, R, S \rrbracket_{vr}$ is defined identically to the standard semantics, except that it permits RFHBNA not to hold. We write $\llbracket B, \mathcal{A}, S \rrbracket_{vr}$ stands for $\llbracket B, \mathcal{A}, \emptyset, S \rrbracket_{vr}$

We then define *deny abstraction* as follows:

$$B_1 \sqsubseteq_d^{\text{NA}} B_2 \triangleq \forall S. \forall X \in \llbracket B_1, \mathcal{A}, S \rrbracket_{vr}^\downarrow. \exists Y \in \llbracket B_2, \mathcal{A}, S \rrbracket_{vr}^\downarrow. \text{hist}_E(X) \sqsubseteq_E \text{hist}_E(Y) \wedge \\ (\forall a. \text{FinalNA}(X, a) \implies \exists b \in A(Y). a \leq_{\text{na}}^{X, Y} b) \wedge \\ (X \in \llbracket B_1, \mathcal{A}, S \rrbracket_{vr} \implies Y \in \llbracket B_2, \mathcal{A}, S \rrbracket_{vr})$$

In addition to the cutting predicated defined in the body of the paper, we need the following to cover NA cuts.

$$\text{NAcutR}(X) \triangleq \forall r_1, r_2 \in (\text{interf}(X) \cap \text{Read} \cap \text{NA}). \\ (\text{val}(r_1) = \text{val}(r_2) = \text{init} \vee \exists w. w \xrightarrow{\text{rf}} r_1 \wedge w \xrightarrow{\text{rf}} r_2) \\ \implies (r_1 = r_2)$$

$$\text{NAcutW}(X) \triangleq \forall w_1, w_2 \in (\text{interf}(X) \cap \text{Write} \cap \text{NA}). \\ (\text{loc}(w_1) = \text{loc}(w_2)) \implies \\ (w_1 = w_2) \vee (\exists r \in \text{code}(X). w_1 \xrightarrow{\text{rf}} r \vee w_2 \xrightarrow{\text{rf}} r)$$

The context cutting predicate is defined as the conjunction of these predicates:

$$\text{cut}^{\text{NA}}(X) \triangleq \text{cutR}(X) \wedge \text{cutW}(X) \wedge \text{NAcutR}(X) \wedge \text{NAcutW}(X)$$

We then define *cut abstraction* as follows:

$$B_1 \sqsubseteq_c^{\text{NA}} B_2 \triangleq \forall X \in \llbracket B_1 \rrbracket_{vr}^\downarrow. \text{cut}^{\text{NA}}(X) \implies \\ \exists Y \in \llbracket B_2 \rrbracket_{vr}^\downarrow. \text{hist}_E(X) \sqsubseteq_E \text{hist}_E(Y) \wedge \\ (\forall a. \text{FinalNA}(X, a) \implies \exists b \in A(Y). a \leq_{\text{na}}^{X,Y} b) \wedge \\ (X \in \llbracket B_1 \rrbracket_{vr} \implies Y \in \llbracket B_2 \rrbracket_{vr})$$

THEOREM 12. $B_1 \sqsubseteq_d^{\text{NA}} B_2 \implies B_1 \sqsubseteq_q^{\text{NA}} B_2$

Proof. Pick a context-side \mathcal{A} , R and an execution $X \in \llbracket B_1, \mathcal{A}, R, S \rrbracket_v$. Case-split on $\text{safe}(X)$ – suppose first that it does not hold.

- Pick a prefix $X' \in X$ and action $a \in A(X')$ such that (1) X' contains precisely one safety violation, which includes a ; and (2) $\text{FinalNA}(X', a)$ holds.
- Generate a new execution X'' by building hb as $(\text{sw} \cup \text{sb})^+$ (i.e. kick out R). As all axioms but RFHBNA are preserved under reduction of hb , $X' \in \llbracket B_1, \mathcal{A}, S \rrbracket_{vr}^\downarrow$.
- Apply the assumption to give an execution $Y' \in \llbracket B_2, \mathcal{A}, S \rrbracket_{vr}^\downarrow$, such that $\text{hist}_E(X'') \sqsubseteq_E \text{hist}_E(Y')$. By the theorem, there must exist an action b to the same location such that $a \leq_{\text{na}} b$.
- Build Y from Y' by defining $\text{hb}(Y)$ as $\text{sb}(Y') \cup \text{rf}_{\text{NA}}(Y') \cup R$, and keeping other relations the same. We now need to establish that (1) $\text{hist}(X') \sqsubseteq_h \text{hist}(Y)$; (2) $\neg \text{safe}(Y)$; and (3) $\text{valid}(Y)$.
- $\text{hist}(X') \sqsubseteq_h \text{hist}(Y)$ holds from the fact that $\text{hist}_{\text{NA}}(X'') \sqsubseteq_E \text{hist}_E(Y')$, and both X' and Y are derived by adding the same relation R .
- To show $\neg \text{safe}(Y)$ we observe that action a in X' participates in a race. As actions in a code-block are sb -sequenced, the other action c forming the race must be in $\text{interf}(X')$. If (b, c) does not form a race in Y , then (b, c) or (c, b) must be in $\text{hb}(Y)$. Any such path must be in $R \cup \text{hbL}(Y) \cup \text{hbA}(Y', b)$. The corresponding path must exist in $R \cup \text{hbL}(X) \cup \text{hbA}(X'', b)$, which rules out the race in X and contradicts the assumption.
- Finally, we need to prove that $\text{valid}(Y)$. HBDEF' holds by construction. RFWF, RFVAL, MOWF, ATOM, NOFAIL are invariant under adding hb -edges, and so follow immediately from $\text{valid}(Y')$. This leaves ACYCLICITY, COHERENCE, HBVS MO, COHERNA, and RFHBNA.

All but RFHBNA are covered by a deny (RFHBNA requires special treatment). A new violation of an axiom caused by edges from R would induce a corresponding deny shape in $\text{hist}_E(Y')$. As $\text{hist}_E(X') \sqsubseteq_E \text{hist}_E(Y)$ this deny shape must also be in X' . However, this means that the corresponding violation can be replicated in X' , which contradicts the assumption that $\text{valid}(X')$ holds.

- Thus, we have an almost-valid execution $Y \in \llbracket B_2, \mathcal{A}, R \rrbracket_{vr}^\downarrow$ such that $\text{hist}(X') \sqsubseteq_h \text{hist}(Y)$; (2) $\neg \text{safe}(Y)$. To complete the proof, we need to fix violations of RFHBNA. We use the same approach as in the proof of Theorem 10: (1) build a shorter prefix in Y^\downarrow which contains precisely one violation of RFHBNA; (2) redirect the read to a valid origin, using the receptiveness of the thread-local semantics. This redirection does alter the history, because non-atomic reads do not appear in the quantified history. This gives an execution Y'' such that (1) $Y'' \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_{vr}^\downarrow$; (2) $\neg \text{safe}(Y'')$; and (3) $\text{hist}(Y'') \in \text{hist}(Y^\downarrow)$.

We finally need to show that there exists an $X''' \in X^\downarrow$ such that $\text{hist}(X''') \sqsubseteq_h \text{hist}(Y'')$. This necessarily exists by application of the history prefixing lemma. Note that X''' may not necessarily be unsafe, but Y'' is guaranteed to be unsafe by construction.

Now suppose that $\text{safe}(X)$ holds. We use essentially the same proof structure as above: the constructed Y may be safe or unsafe, depending whether we need to fix violations of RFHBNA. \square

D. Proof of Theorems 3 and 5

We now prove that $\sqsubseteq_c^{\text{NA}}$ is adequate. Note that because $\sqsubseteq_c^{\text{NA}} \implies \sqsubseteq_c$, we implicitly prove \sqsubseteq_c adequate. We define several versions of the abstractions with different levels of context cutting:

$$B_1 \sqsubseteq_c^i B_2 \triangleq \forall X \in \llbracket B_1, \mathcal{A} \rrbracket_{vr}^\downarrow. \text{cut}^i(X) \implies \\ \exists Y \in \llbracket B_2, \mathcal{A} \rrbracket_{vr}^\downarrow. \text{hist}_E(X) \sqsubseteq_E \text{hist}_E(Y) \wedge \\ (\forall a. \text{FinalNA}(X, a) \implies \exists b \in A(Y). a \leq_{\text{na}}^{X,Y} b) \wedge \\ (X \in \llbracket B_1, \mathcal{A} \rrbracket_{vr} \implies Y \in \llbracket B_2, \mathcal{A} \rrbracket_{vr})$$

We define several versions of the cutting predicate, incrementally cutting more of the context.

$$\begin{aligned}\text{cut}^1(X) &\triangleq \text{cutR}(X) \\ \text{cut}^2(X) &\triangleq \text{cutR}(X) \wedge \text{cutW}(X) \\ \text{cut}^3(X) &\triangleq \text{cutR}(X) \wedge \text{cutW}(X) \wedge \text{NAcutR}(X)\end{aligned}$$

LEMMA 13 (Atomic read cutting). $B_1 \sqsubseteq_c^1 B_2 \implies B_1 \sqsubseteq_d^{\text{NA}} B_2$

Proof. • Pick an execution $X \in \llbracket B_1, \mathcal{A}, S \rrbracket_d^\downarrow$. We now want to build a corresponding execution such that cutR holds.

- Identify a subset $\mathcal{A}' \subseteq A(X)$ such that $\text{cutR}(X|_{\mathcal{A}'})$ holds, and no larger subset exists. We call this maximal projected execution X' . We use \mathcal{A}_R to refer to the removed actions $\mathcal{A} \setminus \mathcal{A}'$.
- It's straightforward to see that $\mathcal{A}_R \subseteq \text{Read} \cap \text{interf}(X)$. Context actions aren't required by the thread-local semantics, and removing context reads preserves validity, so $X' \in \llbracket B_1, \mathcal{A}, S \rrbracket_d^\downarrow$.

It's also straightforward from the definition of cutR to see that any read r in \mathcal{A}_R is removed for one of two reasons:

- *context-read.* The associated write for r is in the context.
- *duplicate-read.* The associated write is read by another context read r' which is not removed. We call this r' the *representative* for r .
- We have an execution $X' \in \llbracket L_1, \mathcal{A}, S \rrbracket_d^\downarrow$ such that $\text{cutR}(X')$ holds. Now apply the assumption to produce an execution $\exists Y' \in \llbracket L_2, \mathcal{A}, S \rrbracket_d^\downarrow$ such that $\text{hist}_E(X') \sqsubseteq_E \text{hist}_E(Y')$.

Build a new execution Y by re-injecting the actions from \mathcal{A}_R . As all of these actions are context reads, the only relation that must change is rf.

- If the action is a context-read, direct rf to the context write it pointed to in X . This must still exist by history inclusion.
- If the action is a duplicate-read, direct rf to the write read by its representative. The origin for the representative write must exist by validity of Y' .

It now remains to show that that $Y \in \llbracket L_2, \mathcal{A}, S \rrbracket_d^\downarrow$ and $\text{hist}_E(X) \sqsubseteq_E \text{hist}_E(Y)$.

- To show that $Y \in \llbracket B_2, \mathcal{A}, S \rrbracket_d^\downarrow$, we only need to show that Y is valid. Adding new atomic context reads to a valid execution is guaranteed to preserve validity, as long as they are equipped with valid origin writes in rf.
- To show that $\text{hist}_E(X) \sqsubseteq_E \text{hist}_E(Y)$, we have two obligations: $\text{hbL}(Y) \subseteq \text{hbL}(X)$, and $\text{denyL}(Y) \subseteq \text{denyL}(X)$. The former is a trivial consequence of the way we construct Y .

For the latter, we reason by contradiction for each of the deny shapes:

- HBvsMO-d and Acyc-d: As context reads are terminal in hb, the only case we need to consider is the one where $u \in \mathcal{A}_R$ and the remainder of the shape is not removed. Otherwise the deny is entirely replicated in Y' , and thus in X . If u is a duplicate-read, the deny is replicated in Y using its representative. If u is a context-read, a deny edge exists $w_1 \xrightarrow{d} v$. In either case, it is easy to see that the deny $u \xrightarrow{d} v$ must be replicable in X , contradicting the assumption.
- Cohere-d and Init-d: Similarly, the cases we need to consider are (1) $u \in \mathcal{A}_R$, (2) $v = r$ and $r \in \mathcal{A}_R$, and (3) both. In the first case, the same argument applies as with HBvsMO. In the second, we can replace r with its representative. In both cases, it's straightforward to replicate the deny $u \xrightarrow{d} v$ is replicated in X . The third case just combines the arguments from the other two.
- CoNA-d and CoNA-d2: Ruled out as actions in \mathcal{A}_R must be hb-terminal. This precludes any such action participating in one of these non-atomic shapes.
- Finally, we need to show that any final NA action in X is replicated in Y , and that Y is complete if X is complete. Both properties are inherited trivially from Y' . □

LEMMA 14 (Atomic write cutting). $B_1 \sqsubseteq_c^2 B_2 \implies B_1 \sqsubseteq_c^1 B_2$

Proof. • Pick an $X \in \llbracket B_1, \mathcal{A}, S \rrbracket_d^\downarrow$ such that $\text{cutR}(X)$ holds.

- Now we build an X' such that $X' \in \llbracket B_1, \mathcal{A}, S \rrbracket_d^\downarrow$ and $\text{cutR}(X) \wedge \text{cutW}'(X)$ holds. First identify the set of non-visible write actions for each location z :

$$\mathcal{A}^z = \{a \in \mathcal{A} \mid \text{loc}(a) = z \wedge a \in (\text{Write} \cap \text{Atomic}) \wedge \neg \text{visible}(a)\}$$

Partition this set into maximal disjoint nonempty subsets $\mathcal{B}_1^z, \mathcal{B}_2^z \dots$ such that:

$$\mathcal{B}_i^z \subseteq \mathcal{A}^z \wedge (\forall a_1, a_2 \in \mathcal{B}_n^z. \neg \exists w \notin \mathcal{B}_n^z. a_1 \xrightarrow{\text{mo}} w \xrightarrow{\text{mo}} a_2)$$

In other words, each set \mathcal{B} is a maximal set of non-visible writes so that there is no intervening write in mo. Thus, either a set \mathcal{B} is mo-minimal / maximal, or it has a visible action which is its immediate mo-predecessor / successor. We call these actions $w_p^{\mathcal{B}}$ and $w_s^{\mathcal{B}}$ respectively. (The cases where \mathcal{B} is minimal / maximal are ignored as they are simpler versions of the case where $w_p^{\mathcal{B}}$ and $w_s^{\mathcal{B}}$ exist.

Note that due to COHERENCE, if there is a LL-SC in \mathcal{B} , it must either read from a write in \mathcal{B} , or from $w_p^{\mathcal{B}}$. Similarly, if $w_s^{\mathcal{B}}$ is a LL-SC, it must read from a write in \mathcal{B} .

To build X' , replace each \mathcal{B} with a single LL-SC pair $w_n^{\mathcal{B}}$ (as above, call this a *representative*). Take as the value that is read the value of $w_p^{\mathcal{B}}$, and take as the written value the mo-final value written in \mathcal{B} . We modify the rest of the execution as follows:

- As each set \mathcal{B} is mo-contiguous in X , we don't need to modify mo other than to insert the new LL-SC pair.
- As the execution satisfies cutR, we have already kicked out all the context reads. We direct rf so that $w_p^{\mathcal{B}} \xrightarrow{\text{rf}} w_n^{\mathcal{B}}$. If $w_s^{\mathcal{B}}$ is a LL-SC, we direct rf so that $w_n^{\mathcal{B}} \xrightarrow{\text{rf}} w_s^{\mathcal{B}}$.
- Introducing $w_n^{\mathcal{B}}$ may generate new hb edges, so we regenerate hb according to HBDEF'.
- We now need to show that X' is valid. This is simple for most of the axioms because the writes that are removed can only be read by their immediate mo-successor. However, if $w_s^{\mathcal{B}}$ is a LL-SC, then we might generate an hb-edge $w_p^{\mathcal{B}} \xrightarrow{\text{hb}} w_s^{\mathcal{B}}$ which did not previously exist. We therefore need to show that ACYCLICITY, HBVSMO, COHERENCE, COHERNA still hold in X' .
 - HBVSMO, ACYCLICITY: The two writes w_1, w_2 responsible must be on a different location from $w_p^{\mathcal{B}}$ and $w_s^{\mathcal{B}}$: otherwise the violation would be an HBVSMO violation in X . Any hb-path between two actions on different locations must pass through the code. If the $w_p^{\mathcal{B}}$ and $w_s^{\mathcal{B}}$ are not themselves in the code, we can replicate the violation immediately using the hb-adjacent internal actions a_p/a_s .
 - COHERENCE, COHERNA: Again, the responsible writes / reads must be to a different location from $w_p^{\mathcal{B}}$ and $w_s^{\mathcal{B}}$. Otherwise we can generate a violation using the LL-SC $w_n^{\mathcal{B}}$, and the fact that mo follows hb. Apply the same reasoning as the previous point to replicate the violation in X .

We also need to show that $\text{cutR}(X') \wedge \text{cutW}(X')$ holds. It's obvious that $\text{cutR}(X')$ still holds – we have introduced no extra reads. $\text{cutW}(X')$ holds because each new write $w_n^{\mathcal{B}}$ is separated in mo by a visible action.

- Apply the assumption to give an execution $Y' \in \llbracket L_2, \mathcal{A}, S \rrbracket$ such that $X' \sqsubseteq_E Y'$. Now build the execution Y . To do this, replace each representative LL-SC $w_n^{\mathcal{B}}$ in Y' with the corresponding actions in \mathcal{B} . In other words, for any pair of actions in a single set \mathcal{B} , take mo the same as in $\text{mo}(X)$. For an action in \mathcal{B} and some other action, relate it in mo as in $\text{mo}(Y')$ for the set representative. We need to show that (1) Y is valid, (2) $\text{hist}_E(X) \sqsubseteq_E \text{hist}_E(Y)$.
- *Validity*. Modifying Y' to Y alters rf, mo, and hb. RFWF, RFVAL, MOWF, ATOM, NOFAIL are obvious by construction. ACYCLICITY holds because hb edges are only removed between existing writes, and introduced between actions represented in \mathcal{B} , which are by definition unrelated to context actions aside from at $w_p^{\mathcal{B}}$ and $w_s^{\mathcal{B}}$. Therefore any cycle would exist inside \mathcal{B} , and thus in X . HBVSMO holds because actions in \mathcal{B} are introduced at a single point in mo represented by $w_n^{\mathcal{B}}$. Any hb-edges inside \mathcal{B} must be consistent with mo, or a similar violation could be replicated in X . COHERENCE, COHERNA holds because any violation for non- \mathcal{B} reads/write could be replicated in Y' using the representative LL-SC $w_n^{\mathcal{B}}$. A violation inside \mathcal{B} could immediately be replicated in X . RFHBNA holds because any hb-path in Y' that is broken in Y must pass through the code. Therefore, the path must be replicable through sb, which contradicts the violation.
- $\text{hbL}(Y) \subseteq \text{hbL}(X)$. Actions in \mathcal{B} in X are only related to each other and $w_p^{\mathcal{B}} / w_s^{\mathcal{B}}$ in hb. For paths in hb outside some \mathcal{B} , it must be that $\text{hbL}(X) = \text{hbL}(X') \subseteq \text{hbL}(Y') = \text{hbL}(Y)$. As paths inside \mathcal{B} are identical between X and Y , any hb-path can be replicated.
- $\text{guarNA}(Y) \subseteq \text{guarNA}(X)$. Trivial by the previous argument, and the fact \mathcal{B} -sets only cover atomic actions.
- $\text{denyL}(Y) \subseteq \text{denyL}(X)$. Prove by contradiction: assume a deny shape in Y that is not in X .
 - HBvsMO-d / Acyc-d: Suppose a deny shape involving writes w_1/w_2 .
 - w_1/w_2 not in any \mathcal{B} . As actions in \mathcal{B} are not read/written in the code, any hb path which includes actions in \mathcal{B} and which passes through the code, must enter and exit \mathcal{B} through other context actions, a, b . There is a deny $a \xrightarrow{d} b$ in Y' by construction, and thus in X . hb-paths inside \mathcal{B} are identical in X and Y . Combining this gives us a deny in X .
 - w_1/w_2 entirely inside \mathcal{B} : reproducible trivially as mo/hb are identical between Y and X .

- w_1 outside \mathcal{B} , w_2 inside. In this case, there must be a deny in Y' and X' with the representative: $w_p^{\mathcal{B}} \xrightarrow{d} b$ (using the same argument as above). Substituting \mathcal{B} for the representative in X builds the violation.
- w_2 outside \mathcal{B} , w_1 inside. Symmetrical to previous case.
- Cohere-d / Init-d: the deny shape involves writes $w_1/w_2/r$.
 - w_1, w_2, r all in \mathcal{B} : replicated trivially.
 - w_1, w_2, r all outside \mathcal{B} : replicated trivially.
 - w_1, w_2 in \mathcal{B} , r outside: r can only be $w_s^{\mathcal{B}}$, shape ruled out by construction.
 - w_1 in \mathcal{B} , w_2, r outside: deny replicated in Y' / X' using representative. Rebuild the violation when reintroducing \mathcal{B} in X .
 - All three outside \mathcal{B} : trivial.
 - w_2, r in \mathcal{B} , w_1 outside: w_1 can only be $w_p^{\mathcal{B}}$, shape ruled out by construction.
 - w_2 in \mathcal{B} , w_1, r outside: deny replicated in Y' using representative, rebuild in X when adding \mathcal{B} .
 - r in \mathcal{B} , w_1, w_2 outside: w_1 can only be $w_p^{\mathcal{B}}$, shape ruled out by construction.
- CoNA-d2: similar argument to Cohere-d, except that some cases are ruled out by the fact that elements in \mathcal{B} are necessarily atomic.
- $\text{denyNA}(X) \subseteq \text{denyNA}(Y)$. Similar argument to CoNA-d2 above.
- The Final NA and completeness properties are satisfied for the same reason as in the previous proof. □

LEMMA 15 (NA read cutting). $B_1 \sqsubseteq_c^3 B_2 \implies B_1 \sqsubseteq_c^2 B_2$

Proof. • Pick an $X \in \llbracket B_1, \mathcal{A}, S \rrbracket^\downarrow$ such that $\text{cut}^2(X)$ holds. Build X' using the same approach as in atomic read cutting: X' is a maximal sub-execution such that $\text{NAcutR}(X)$ holds.

From the structure of NAcutR , the actions \mathcal{A}_R removed from X must all be non-atomic reads. Just as before, removed reads have a *representative* that remains in X' and that reads from the same write. Unlike in the atomic case, reads from context writes also have representatives. This is necessary to detect new writes that might violated CoNA-d2 (which in turn is needed because NA writes are not ordered in mo).

X' is valid because the axioms are invariant under read removal.

- We then apply the assumption to build an execution $Y' \in \llbracket B_2, \mathcal{A}, S \rrbracket_{vrr}^\downarrow$. Finally we build Y by restoring the cut actions, with $\text{rf}(Y')$ built in the same way as for the atomic cutting case. Almost-validity is preserved trivially because the inserted reads are not part of hb. Deny inclusion is ensured by the fact that the inserted reads are placed at the same position as their representatives: any violation would immediately be replicated by the representative. The FinalNA and completion property are unaffected from Y' . □

LEMMA 16 (NA write cutting). $B_1 \sqsubseteq_c B_2 \implies B_1 \sqsubseteq_c^3 B_2$

Proof. • Pick an $X \in \llbracket B_1, \mathcal{A}, S \rrbracket^\downarrow$ such that $\text{cut}^3(X)$ holds.

- As NAcutW doesn't discriminate on the basis of mo, we can replace the set of all context writes to a location with a single representative write. We build X' as a maximal sub-execution such that NAcutW holds, and 'orphan' context reads are removed. As NAcutR holds, each NA write has at most one context read. Note that as the execution is maximal, if at least one write to a location had an associated read, then the representative will have an associated read.
- Validity for X' is trivial as the removed writes cannot participate in hb, or be read in the code. We then build Y' by applying the theorem to give an almost-valid execution of B_2 . Finally, we build Y by re-inserting the removed reads and writes. The only relation that needs to be updated is rf , which associates removed reads with their origin write. Preservation of almost-validity follows from the fact that the inserted writes are disjoint from all other actions in the execution relations. Deny inclusion holds because any deny shape in Y that involves a removed write / read can be easily replicated using the representative. □

THEOREM 17 (Cut adequacy). $B_1 \sqsubseteq_c B_2 \implies B_1 \sqsubseteq_d^{\text{NA}} B_2$

Proof. Prove this as a sequence of implications:

$$B_1 \sqsubseteq_c B_2 \implies B_1 \sqsubseteq_c^3 B_2 \implies B_1 \sqsubseteq_c^2 B_2 \implies B_1 \sqsubseteq_c^1 B_2 \implies B_1 \sqsubseteq_d^{\text{NA}} B_2$$

Each implication step is proved in a lemma above. □

E. Proof of Theorem 6

E.1 Proof structure.

Assume $B_1 \preceq_{bl} B_2$; we have to prove $B_1 \sqsubseteq_q B_2$. At the high level, we prove this with the following steps:

1. Following (7), consider arbitrary \mathcal{A}, R, S and $X_1 \in \llbracket B_1, \mathcal{A}, R, S \rrbracket$.
2. We use X_1 to construct a special context C . The context performs the actions specified by \mathcal{A} and monitors executions to ensure that they do not significantly diverge from X_1 , e.g., by checking that the values returned by context reads match those in \mathcal{A} . If C detects a mismatch with X_1 , it writes to a special observable error variable $e \in \text{OVar}$. The context C is constructed in such a way that for any code-block B' and any execution $Y \in \llbracket C(B') \rrbracket$ in which e is not written, the following hold:
 - (a) the actions of \mathcal{A} appear in Y , and the actions by B' in Y transform local variables in a way consistent with the call and ret actions in X_1 ;
 - (b) $\text{hb}(Y)$ includes the edges in R ;
 - (c) $\text{hb}(Y)$ does not include any of the edges in the complement of the guarantee in $\text{hist}(X_1)$.
3. We show that there is an execution $Z_1 \in \llbracket C(B_1) \rrbracket$ where the actions generated by B_1 match those in X_1 , and where e is not written; the latter implies that the above properties (a), (b) and (c) hold of Z_1 .
4. Since $B_1 \preceq_{bl} B_2$, applying (2) to the context C , we get an execution $Z_2 \in \llbracket C(B_2) \rrbracket$ where e is never written.
5. By the construction of C , we have (a) and (b). Using this fact, we construct an execution $X_2 \in \llbracket B_2, \mathcal{A}, R, S \rrbracket$ where the actions generated by B_2 match those in Z_2 and the call and ret actions match those in X_1 . Let $\text{hist}(X_1) = (A_1, G_1)$ and $X_2 = (A_2, G_2)$. Using (a), we show $A_1 = A_2$ and using (c) we show $G_2 \subseteq G_1$. This establishes $\text{hist}(X_1) \sqsubseteq_h \text{hist}(X_2)$, and by (7), the required $B_1 \sqsubseteq_q B_2$.

E.2 Full abstraction with non-atomics

We now prove Theorem 6: full abstraction of $\sqsubseteq_q^{\text{NA}}$ for programs and contexts that do not use read-modify-write accesses, $B_1 \preceq_{bl}^{\text{NA}} B_2 \implies B_1 \sqsubseteq_q^{\text{NA}} B_2$.

Proof. • Start by choosing an arbitrary R, S and $X \in \llbracket B_1, \mathcal{A}, R, S \rrbracket_v$. It remains to show that:

$$\begin{aligned} & \exists Y \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_v^\downarrow. \\ & \exists X' \in X^\downarrow. \text{hist}(X') \sqsubseteq_h \text{hist}(Y) \wedge \\ & \text{safe}(Y) \implies (\text{safe}(X) \wedge (X' = X) \wedge Y \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_v) \end{aligned} \quad (10)$$

- Apply the construction lemma to B_1, R, S and X to find a context C_X , and an execution Z :

$$\begin{aligned} & Z \in \llbracket C_X(B_1) \rrbracket_v \wedge \\ & \text{code}(Z) = X \wedge \\ & \text{hbC}(Z) = R \wedge \text{atC}(Z) = S \wedge \\ & \forall B'. \forall Z' \in \llbracket C_X(B') \rrbracket_v. \\ & ((A(\text{contx}(Z)) = A(\text{contx}(Z'))) \implies (\text{hist}(Z) \sqsubseteq_h \text{hist}(Z')) \wedge \text{at}(\text{contx}(Z)) = \text{at}(\text{contx}(Z'))) \wedge \\ & (\neg \text{safe}(Z') \implies \exists X' \in X^\downarrow. \exists W \in \llbracket B', \mathcal{A}, R, S \rrbracket_v^\downarrow. \neg \text{safe}(W) \wedge \text{hist}(X') \sqsubseteq_h \text{hist}(W)) \end{aligned} \quad (11)$$

- Specialise observation with the context C_X and the set of all variables used in C_X, V_{C_X} , to get:

$$\begin{aligned} & (\neg \text{safe}(C_X(B_2)) \vee \\ & (\text{safe}(C_X(B_1)) \wedge \\ & \forall X \in \llbracket C_X(B_1) \rrbracket_v. \exists Y \in \llbracket C_X(B_2) \rrbracket_v. \\ & \mathcal{A}(X|_{V_{C_X}}) = \mathcal{A}(Y|_{V_{C_X}})) \wedge \\ & \text{hb}(X|_{V_{C_X}}) \subseteq \text{hb}(Y|_{V_{C_X}})) \end{aligned} \quad (12)$$

and then case split on $\text{safe}(C_X(B_2))$.

- **Case 1:** $\text{safe}(C_X(B_2))$.

- By 12, there is an execution, Z' of $C_X(B_2)$ with:

$$\text{hb}(Z|_{V_{C_X}}) = \text{hb}(Z'|_{V_{C_X}}) \wedge \mathcal{A}(Z|_{V_{C_X}}) = \mathcal{A}(Z'|_{V_{C_X}})$$

- By construction of C_X , the variables V_{C_X} cover all context variables, so we have:

$$\text{hb}(\text{contx}(Z)) = \text{hb}(\text{contx}(Z')) \wedge A(\text{contx}(Z)) = A(\text{contx}(Z'))$$

- Appealing to 11, we have:

$$\text{hist}(Z) \sqsubseteq_h \text{hist}(Z')$$

- Now apply the decomposition lemma to Z' to get the execution Y :

$$Y \in \llbracket B_2, \mathcal{A}, \text{hbC}(Z'), \text{atC}(Z') \rrbracket_v \wedge \text{code}(Z') = Y$$

- Now simplify using the definition of hbC and atC .

$$Y \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_v \wedge \text{hbC}(Z') = R = \text{hbC}(Z) \wedge \text{atC}(Z') = S = \text{atC}(Z)$$

- Choose Y as the witness for our goal 10, and X as the witness for the inner quantification. It is left to show that:

$$\text{hist}(X) \sqsubseteq_h \text{hist}(Y)$$

- Unfolding the definition of \sqsubseteq_h , we have:

$$A(Z) = A(Z') \wedge \text{hbL}(Z') \subseteq \text{hbL}(Z)$$

and it is left to show that,

$$A(X) = A(Y) \wedge \text{hbL}(Y) \subseteq \text{hbL}(X)$$

- Note that X and Y are the code-block projections of Z and Z' respectively, and we are done.

- **Case 2:** $\neg \text{safe}(C_X(B_2))$.

- Identify an unsafe valid execution of $C_X(B_2)$, Z' , and specify the final conjunct of 11 with B_2 and Z' to get:

$$X' \in X^\downarrow \wedge W \in \llbracket B_2, \mathcal{A}, R, S \rrbracket_v^\downarrow \wedge \neg \text{safe}(W) \wedge \text{hist}(X')^\downarrow \sqsubseteq_h \text{hist}(W)$$

This case is completed by noting that W satisfies 10. □

F. Proof of Theorems 2 and 6


The proof of full abstraction relies on a special context construction that is sensitive enough to pick apart any observable difference in one execution and another. The construction is built from one execution, replicating context hb and at , or writing a context error variable. The context ensures that when composed with any block, all executions with matching context hb and at produce histories that are abstractions of the original execution. For unsafe executions, this history abstraction need only hold over a hb -prefix that includes a safety violation.

The context construction. Figure 4 describes the construction of context C from an execution $X \in \llbracket B, \mathcal{A}, R, S \rrbracket$ (for brevity, we use syntactic sugar that elides manipulations of local variables). The context is a parallel composition of threads: one for the parameter code-block $\{-\}$, one for each ll/sc pair in at , ensuring replication of S , and one each for all other actions in \mathcal{A} — these are collectively ranged over by m in Fig. 4.


We introduce functions call and ret on the indices m , mapping $\{-\}$ to the call and ret actions in X , respectively, and acting as the identity otherwise. The construction then consists of several wrapper functions:

1. Innermost is $\text{check}(m)$: for a read or a write action $u \in \mathcal{A}$, $\text{check}(u)$ executes the corresponding operation and, in the case of a read or a sc , compares the value recorded with the one specified by u . The command $\text{check}(\{-\})$ initialises local variables to the values specified by the call action in X , runs the parameter code-block $\{-\}$ and then compares the local variables with the values specified by the ret action in X . If there is a mismatch in the above cases, check writes to the error variable e and halts the thread.

2. The wrappers $Rrel_m$ and $Racq_m$ replicate R . Each $Rrel_m$ is built up of a sequence of invocations of $Rrel_{u,v}$, one for each edge $(u, v) \in R$ outgoing from $u = ret(m)$; $Racq_m$ is constructed symmetrically. These wrappers use watchdog variables $h_{u,v}$ to create happens-before by writing to and reading as in the MP test of §3. If $Racq_{u,v}$ does *not* read the value written by $Rrel_{u,v}$, then it writes to the error variable. Otherwise, $Rrel_{u,v}$ is sequenced after u and $Racq_{u,v}$ before v , reproducing R with an execution of the following shape:



3. The wrappers $Nrel_m$ and $Nacq_m$ ensure that the history of an error-free execution of $C(B')$ abstracts the history of the original execution, i.e. there are no new guarantee edges beyond those of the original G . Each $Nrel_m$ is built up of a sequence of invocations of $Nrel_{u,v}$, one for each edge (u, v) that might be in the history, but is not covered by R or G . $Nacq_m$ is constructed symmetrically. Errant happens-before edges are detected using watchdog variables $g_{u,v}$, again relying on the mechanics of MP (§3): $Nrel_{u,v}$ and $Nacq_{u,v}$ respectively write to and read from $g_{u,v}$. If $Nacq_{u,v}$ *does* read the value written by $Nrel_{u,v}$, then there is a new guarantee edge, e is written. $Rrel_{u,v}$ is sequenced *before* u and $Racq_{u,v}$ *after* v , so an erroneous execution has the following shape:



Here, the RFVAL axiom (§3) forces $Nacq_{u,v}$ to read from $Nrel_{u,v}$ and write to e . Consequently, a non-erroneous execution does not contain errant happens-before edges. As required by the theorem.

To illustrate the construction, we use the execution X in Figure 2, for B defined by (5). The set \mathcal{A} consists of the three writes outside the dashed rectangle, and R is given by the dashed edges. The edge $(write(x, 2), call) \in R$ is reproduced by the invocation of $Racq_{write(x,2),call}$ on the first thread and $Rrel_{write(x,2),call}$ on the second. The edges from G that we need to consider are $(call, write(f, 1))$ and $(call, write(x, 1))$, and the edge $(call, write(f, 1)) \in \bar{G}$ is covered by the invocation of $Nrel_{call,write(f,1)}$ on the first thread and $Nacq_{call,write(f,1)}$ on the fourth.

F.1 Construction Lemma with non-atomics, with read-modify-writes

LEMMA 18 (Construction Lemma).

$$\begin{aligned}
& \forall B \mathcal{A}, R, S . \forall X \in \llbracket B, \mathcal{A}, R, S \rrbracket_v . \\
& \quad \exists C_X . \exists Z \in \llbracket C_X(B) \rrbracket_v . \\
& \quad (\text{code}(Z) = X) \wedge \\
& \quad (\text{hbC}(Z) = R) \wedge (\text{atC}(Z) = S) \wedge \\
& \quad \forall B' . \forall Z' \in \llbracket C_X(B') \rrbracket_v . \\
& \quad ((A(\text{contx}(Z)) = A(\text{contx}(Z'))) \implies \text{hist}(Z) \sqsubseteq_{\text{h}} \text{hist}(Z') \wedge \text{at}(\text{contx}(Z)) = \text{at}(\text{contx}(Z'))) \wedge \\
& \quad (\neg \text{safe}(Z') \implies \exists X' \in X^\downarrow . \exists W \in \llbracket B', \mathcal{A}, R, S \rrbracket_v^\downarrow . \neg \text{safe}(W) \wedge \text{hist}(X') \sqsubseteq_{\text{h}} \text{hist}(W))
\end{aligned}$$

Proof. • Start by choosing an arbitrary B, \mathcal{A}, R, S and $X \in \llbracket B, \mathcal{A}, R, S \rrbracket$.

- Construct the client C_X as specified in Figure 4 with one minor change: have check halt the thread if the error variable is written.

It remains to show that there exists $Z \in \llbracket C_X(B) \rrbracket_v$ such that:

$$\begin{aligned}
C(-) &= (;(m_1, m_2) \in \text{at } a(m_1); a(m_2)) \parallel (\parallel_{m \setminus \text{at}} a(m)) \\
a(m) &= \text{Racq}_m(\text{Nrel}_m; \text{check}(m)(\text{Nacq}_m(\text{Rrel}_m))) \\
\text{Rrel}_m &= \text{Rrel}_{\text{ret}(m), v_1}; \dots; \text{Rrel}_{\text{ret}(m), v_n}, \\
&\quad \text{where } \{v_1, \dots, v_n\} = \{v \mid (\text{ret}(m), v) \in R\} \\
\text{Rrel}_{u,v} &= \text{write}(h_{u,v}, 1) \\
\text{Racq}_m(-) &= \text{Racq}_{u_1, \text{call}(m)}(\dots \text{Racq}_{u_n, \text{call}(m)}(\{-\}) \dots), \\
&\quad \text{where } \{u_1, \dots, u_n\} = \{u \mid (u, \text{call}(m)) \in R\} \\
\text{Racq}_{u,v}(-) &= \text{if } (\text{read}(h_{u,v})) \{-\} \text{ else write}(e, 1) \\
\text{Nrel}_m &= \text{Nrel}_{\text{call}(m), v_1}; \dots; \text{Nrel}_{\text{call}(m), v_n}, \\
&\quad \text{where } \{v_1, \dots, v_n\} = \{v \mid (\text{call}(m), v) \in H\} \\
\text{Nrel}_{u,v} &= \text{write}(g_{u,v}) \\
\text{Nacq}_m(-) &= \text{Nacq}_{u_1, \text{ret}(m)}(\dots \text{Nacq}_{u_n, \text{ret}(m)}(\{-\}) \dots), \\
&\quad \text{where } \{u_1, \dots, u_n\} = \{u \mid (u, \text{ret}(m)) \in H\} \\
\text{Nacq}_{u,v}(-) &= \text{if } (\neg \text{read}(g_{u,v})) \{-\} \text{ else write}(e, 1)
\end{aligned}$$

Figure 4. Context construction.

$$\begin{aligned}
\text{check}(\{-\}) &= l1 := 0; l2 := 0; \\
&\quad \{-\}; \\
&\quad \text{if } (l1 \neq 1) \{\text{write}(e, 1)\}; \\
&\quad \text{if } (l2 \neq 1) \{\text{write}(e, 1)\} \\
C(-) &= \text{Racq}_{\text{write}(x,2), \text{call}}(\text{Nrel}_{\text{call}, \text{write}(x,1)}; \text{Nrel}_{\text{call}, \text{write}(f,1)}; \\
&\quad \text{check}(\{-\})) \\
&\quad \parallel \text{write}(x, 2); \text{Rrel}_{\text{write}(x,2), \text{ret}}; \text{Rrel}_{\text{write}(x,2), \text{write}(x,2)}; \\
&\quad \parallel \text{Racq}_{\text{write}(x,2), \text{write}(x,1)}(\text{write}(x, 1); \\
&\quad \quad \text{Nacq}_{\text{ret}, \text{write}(x,1)}(\text{Rrel}_{\text{write}(x,1), \text{write}(f,1)})) \\
&\quad \parallel \text{Racq}_{\text{write}(x,1), \text{write}(f,1)}(\text{write}(f, 1); \\
&\quad \quad \text{Nacq}_{\text{ret}, \text{write}(f,1)}(\text{skip}))
\end{aligned}$$

Figure 5. Example of the check function and the context construction for the execution in Figure 2.

1. $(\text{code}(Z) = X) \wedge (\text{hbC}(Z) = R) \wedge (\text{atC}(Z) = S)$
 2. $\forall B'. \forall Z' \in \llbracket C_X(B') \rrbracket_v.$
 $((A(\text{contx}(Z)) = A(\text{contx}(Z'))) \implies \text{hist}(Z) \sqsubseteq_h \text{hist}(Z') \wedge \text{at}(\text{contx}(Z)) = \text{at}(\text{contx}(Z'))) \wedge$
 $(\neg \text{safe}(Z') \implies \exists X' \in X^\downarrow. \exists W \in \llbracket B', \mathcal{A}, R, S \rrbracket_v^\downarrow. \neg \text{safe}(W) \wedge \text{hist}(X') \sqsubseteq_h \text{hist}(W))$
- We first establish 1.
 - Appealing to the thread local semantics and the structure of $C_X(B)$, choose Z_p , a pre-execution of $C_X(B)$ that does not write the error variable, and whose code projection matches X .
 - Construct at as follows: for code actions choose these edges to match X , and for the context part, note that there is no choice: each \parallel/sc pair is in its own thread, so there is only one way to link them.
 - Construct mo as follows: for code actions choose these edges to match X , and for the context part, note that there is no choice: at each location there is only one write after the initialisation.
 - Construct rf as follows: for code actions choose these edges to match X , and for that context actions set rf to be coincident with an R edge in the case of Racq or from the initialisation write in the case of Nacq . Note that the context projection of happens-before matches R by construction.
 - Let Z be the combination of Z_p , mo and rf . Show that Z is valid.

- The only axioms that could fail are *Acyclicity* over some *Racq* or *Coherence* over some *Nacq*.
- In the first case, any cycle would be made up of code *hb* and *R* edges, and would also be a cycle in *X*, a contradiction.
- A *Coherence* violation over some *Nacq* implies the existence of a *hb* edge from the associated *Nrel* to the *Nacq*. This violates the rules used to construct C_X , and is a contradiction.

• Now establish 2.

- Start by choosing arbitrary B' and $Z' \in \llbracket C_X(B') \rrbracket_v$.
- First, show that $(A(\text{ctx}(Z)) = A(\text{ctx}(Z'))) \implies \text{hist}(Z) \sqsubseteq_h \text{hist}(Z') \wedge \text{at}(\text{ctx}(Z)) = \text{at}(\text{ctx}(Z'))$
 - Z does not write e , so neither does Z' (they have an equal context projection).
 - By construction of C_X , the histories of Z abstract the histories of Z' and the *at* relations match.
- Now suppose Z' is unsafe. It remains to show:

$$\exists X' \in X^\downarrow. \exists W \in \llbracket B', \mathcal{A}, R, S \rrbracket_v^\downarrow. \neg \text{safe}(W) \wedge \text{hist}(X') \sqsubseteq_h \text{hist}(W)$$

- C_X uses only atomic and local variables that cannot exhibit safety violations: each violation must be amongst the actions of \mathcal{A} and the actions generated by B' .
- Identify a safety violation in Z' and consider the prefix Z'_p containing only $\text{hb} \cup \text{rf}$ predecessors of the actions of the violation.
- There are no writes to e in Z'_p : after any such write, the thread is stopped, so it cannot appear in the prefix Z'_p .
- Below, we establish that for every thread of Z'_p except those that contain the safety violation from which it is constructed, the error variable is not written in the corresponding thread of Z' .
 - Consider the $\text{hb} \cup \text{rf}$ edges that draw actions into the prefix Z'_p from some thread t , there are two cases: the edge arises from a *Rrel/Racq* pair, or it is created by a read, from write w , in the code block or a context action.
 - In the first case e is not written in *check* or *Nacq* on t , because that would halt the thread before the call to *Rrel*.
 - In the second case, no call to *Nacq* writes e , and calls of *commit* only write to e in the case of a failing store conditional, contradicting the existence of write w in Z' , so w is only performed in threads that never write to e .
- From Z'_p , we construct W by applying the decomposition lemma Z'_p to get an execution W , completing this to an execution in $\llbracket B', \mathcal{A}, R, S \rrbracket_v$, and observing that W is in $\llbracket B', \mathcal{A}, R, S \rrbracket_v^\downarrow$. W is unsafe by construction.
- Take A to be the set of all context actions in W together with the call and *ret* actions present in W . Let X' be the projection of X to the $\text{hb} \cup \text{rf}$ predecessors of A , so that $X' \in X^\downarrow$.
- It remains to show that there is no edge in $\text{hb}(W)$ between the actions of A that is not present in the guarantee of X'_p, G' . By construction of W , any extraneous $\text{hb}(W)$ edge must end at one of the threads hosting the violation. There is only one code block, so at least one of the threads has to be executing a context action. There is no single action that both creates an incoming hb edge and causes a safety violation, so any additional $\text{hb}(W)$ edge must end at the code block. In that case, W does not include the *ret* action following the racy action in hb , and neither does X' , and there can be no new edge in G' .

□