

Formalizing the Concurrency Semantics of an LLVM Fragment

Soham Chakraborty Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS), Germany

Abstract

The LLVM compiler follows closely the concurrency model of C/C++ 2011, but with a crucial difference. While in C/C++ a data race between a non-atomic read and a write is declared to be undefined behavior, in LLVM such a race has defined behavior: the read returns the special ‘undef’ value. This subtle difference in the semantics of racy programs has profound consequences on the set of allowed program transformations, but it has been not formally studied before.

This work closes this gap by providing a formal memory model for a substantial fragment of LLVM and showing that it is correct as a concurrency model for a compiler intermediate language: (1) it is stronger than the C/C++ model, (2) weaker than the known hardware models, and (3) supports the expected program transformations. In order to support LLVM’s semantics for racy accesses, our formal model does not work on the level of single executions as the hardware and the C/C++ models do, but rather uses more elaborate structures called event structures.

1. Introduction

With the advent of the multi-core era, programming languages such as C/C++ and Java have introduced first-class platform-independent support for concurrent programming. For example, the 2011 ISO C standard (termed C11) introduced a library of atomic operations together with a concurrency model, a complex set of rules detailing which outcomes a concurrent program may produce. This concurrency model is a set of promises that a compiler has to fulfill and that programmers can rely upon.

LLVM is a *state-of-the-art* optimizing compiler that fully supports the C11 concurrency primitives and has some rudimentary support for Java concurrency. As such, it has its own concurrency model, which determines how the various concurrency primitives should be compiled and optimized. In the LLVM documentation, the LLVM concurrency model is described as a slight variant of the C11 model, though the exact correspondence is not clearly specified. In fact, the entire specification of the LLVM model is extremely informal. It consists of some informal prose that often refers to the C11 model, a couple of examples, and more importantly a collection of transformations which should be allowed or disallowed. This list of transformations is arguably the most valu-

able part of the specification, because it serves as a guideline to the developers implementing the compiler optimizations.

To date, no formal definition of the LLVM concurrency model exists. (Prior LLVM formalizations, such as VeLLVM [28], have restricted attention to sequential programs.) This lack of a formal model has had a negative effect on LLVM concurrency compilation. First, the LLVM compiler is often overly cautious in optimizing code involving shared memory accesses missing out some optimization opportunities (e.g., eliminating redundant atomic accesses). Second, the difference between the C11 and LLVM semantics remains unappreciated, which has led to subtle compiler bugs [7].

In this paper, we formalize the concurrency semantics for a substantial subset of LLVM, and show that the transformations intended to be correct according to the LLVM documentation are indeed so in our formal model. We also show that our model is stronger than the C11 model, meaning that the standard compilation from C/C++ to the LLVM intermediate representation is correct, and weaker than the Total Store Order (TSO) and Power memory models, meaning that compilation to these hardware models is also correct.

The key challenge that our formalization addresses is the different semantics for racy programs between LLVM and C11. In C11, racy programs are completely undefined. In LLVM, however, read-write races are always well-defined: the read may simply return an *arbitrary* value. This seemingly innocent difference has an important impact on the set of allowed program transformations: it enables hoisting loads outside of conditionals and loops but disallows common subexpression elimination (CSE) over acquire-atomic accesses and fences [7].

To model LLVM’s semantics for read-write races, we depart from the standard “per candidate execution” axiomatic style of defining memory models (e.g., [1, 2, 16, 22, 24, 25]). As noted by Batty et al. [3], this standard style of defining memory models cannot adequately distinguish between the next two programs in the case they return $a = b = 1$.

$$\begin{array}{l} a = X; \quad \left\| \quad b = Y; \right. \\ \text{if}(a) \quad \left\| \quad \text{if}(b) \quad (\text{CYC}) \quad a = X; \quad \left\| \quad b = Y; \right. \\ Y = 1; \quad \left\| \quad X = 1; \quad \quad \quad Y = 1; \quad \left\| \quad X = 1; \quad (\text{LB}) \end{array}$$

For both programs, let the initial state be $X = Y = 0$ and consider whether the outcome $a = b = 1$ should be allowed. In the case of CYC, the outcome is clearly undesirable because it violates the *data-race-freedom* (DRF) property of




Figure 1. Event structure of the CYC program.





Figure 2. Event structure of LB and LB+false-dep.

a memory model: that is, under *sequential consistent* (SC) semantics [18], the program is race-free; and thus should not have weak behaviors. In the case of LB, however, this outcome must be allowed because it can be observed on ARM [9]. The problem is that the C11-style candidate executions of these two programs yielding $a = b = 1$ are identical (treating the accesses as C11 “relaxed”), and hence any model based on C11-style executions cannot distinguish them. Hardware models [1] typically resolve this problem is by recording (syntactic) control, data, and address dependencies. This, however, does not work for us, because compiler optimizations often remove such dependencies. For example, in the program below may be optimized into LB and so the outcome $a = b = 1$ should be allowed.

$$\begin{array}{l} a = X; \\ \text{if}(a) Y = 1; \\ \text{else } Y = 1; \end{array} \parallel \begin{array}{l} b = Y; \\ \text{if}(b) X = 1; \end{array} \quad (\text{LB+false-dep})$$

Instead, we base our model on an adaptation of *prime event structures* [27], which we build incrementally via an operational semantics that checks that each appended read event is justified by some existing write event. The beauty of event structures is that they appropriately distinguish between CYC and LB, because they capture multiple executions in a single structure. In our formal model, CYC produces only the event structure shown in Figure 1 where the po-edges represent the program order (the initialization occurs before the two threads) and the rf-edges denote the read-from relation matching each (non-racy) load to the store, from where the load got its value. The event structure justifies only the outcome $a = b = 0$, whereas LB and LB+false-dep also produce the structure in Figure 2, which contains conflicts (\sim) among the read events generated from a single load operation along with the po and rf edges. Among others, this event structure justifies the outcome $a = b = 1$, when each \mathbf{u} evaluates to 1 (see §2 for an explanation of \mathbf{u}).

In our formal development, we omit LLVM’s *monotonic accesses*, which correspond the C11’s relaxed accesses, because of the known problems with “out of thin air” reads [3, 5, 12, 23, 26] and the long outstanding problem of providing an adequate solution. We believe that the treatment



§4: DRF theorems

Figure 3. LLVM compilation and paper overview

of monotonic accesses is essentially orthogonal to the treatment of racy non-atomic accesses, which constitutes the focus of this work. Nevertheless, event structures might still be a good mechanism to also model monotonic accesses (e.g., following Jeffrey and Riely [12]). An alternative very promising approach to model monotonic accesses was recently proposed in [13]; incorporating it with our semantics of the remaining accesses is left as future work.

Outline The structure of the remainder of the paper follows largely that of the LLVM compiler (see Figure 3). After some background material in §2, we present our formal LLVM concurrency model in §3, and derive some DRF theorems in §4. Next, in §5, we show that compilation from C11 to our model is correct. In §6, we show that the intended re-ordering and elimination transformations are allowed under our model. In §7, we show that compilation to the TSO and Power models is correct. We conclude with a discussion of related and future work. The proofs are available online [8].

2. Background

In this section, we introduce some basic notation and a simple programming language. We also discuss the semantics of LLVM’s undefined value and of racy programs.

Notation Given a set E and binary relations $R, S \subseteq E \times E$, we write $R; S$ for the relational composition of R and S ; formally, $(R; S)(x, y) \triangleq \exists z. R(x, z) \wedge S(z, y)$. We write $R^?$, R^+ , R^* for the reflexive, the transitive and the reflexive-transitive closures of R respectively, and R^{-1} for the inverse of R . The $[A]$ notation denotes an identity relation on set A , i.e. $[A](x, y) \triangleq x = y \wedge x \in A$. Finally, we write $\mathbf{one}(A)$ for the relation saying that at least one of its components belongs to the set A ; i.e. $\mathbf{one}(A)(x, y) \triangleq (x \in A \vee y \in A)$.

Language Our formal model is essentially orthogonal to the syntax of the programming language, but for concreteness we present a simple concurrent imperative language in Figure 4. As already explained in the introduction, in this work we focus on a subset of operations and memory order, excluding monotonic, unordered accesses and fences.

LLVM values, v , can be either constants or the special *undefined value* \mathbf{u} , which is roughly a placeholder for any possible value. Given two values v and v' , we write $v \preceq v'$ if $v = v'$ or $v' = \mathbf{u}$. Expressions, E , can be built from local variables and values using arithmetic operations.

Commands, C , are sequences of instructions including assignments, shared memory operations, as well as uncon-

$$\begin{aligned}
v &::= c \mid \mathbf{u} && \text{(Value)} \\
E &::= t \mid v \mid X \mid E + E \mid E * E \mid E \leq E \mid \dots && \text{(Expr)} \\
C &::= \text{skip} \mid C; C \mid t = E \mid t = \text{load}_{(\text{NA}|\text{ACQ}|\text{SC})} X \\
&\quad \mid \text{store}_{(\text{NA}|\text{REL}|\text{SC})}(X, E) \mid \text{CAS}_{(\text{ACQ}|\text{REL}|\text{SC})}(X, E, E) \\
&\quad \mid \text{label}: C \mid \text{br label} \mid \text{br } t \text{ label label} && \text{(Cmd)} \\
P &::= X = v; \dots X = v; \{C \parallel \dots \parallel C\} && \text{(Program)}
\end{aligned}$$

Figure 4. Syntax of a minimal programming language. Here, X denotes a shared and t a thread-local location.

ditional and conditional branches. The shared memory operations (load, store, compare-and-swap) are annotated with a *memory order*, o . For loads, this can be non-atomic (NA), acquire (ACQ), or sequentially consistent (SC). For stores, it can be non-atomic, release (REL), or sequentially consistent. For CAS, it can be either acquire-release (ACQ_REL) or sequentially consistent. In increasing strength, these orders are: $\text{NA} \sqsubset \{\text{ACQ}, \text{REL}\} \sqsubset \text{ACQ_REL} \sqsubset \text{SC}$.

Finally, programs, P , consist of a sequence of initialization writes followed by a fixed parallel composition of thread commands. In our examples, we freely use C-like syntax.

2.1 The Semantics of the Undefined Value

In LLVM, the special undefined value \mathbf{u} is introduced as the result of erroneous computations, such as reading from an uninitialized memory location as a replacement of an arbitrary constant value. This special value propagates through every assignment and arithmetic operation. So, for example, $\mathbf{u} + 1 \rightsquigarrow \mathbf{u}$ and even $\mathbf{u} * 0 \rightsquigarrow \mathbf{u}$.¹

The intended semantics is that the compiler may replace \mathbf{u} with any concrete value it finds most convenient, and that moreover different uses of the same \mathbf{u} may be even replaced by different values by the compiler. This weak semantics leads to some rather unexpected behaviors. For example,

```
int t; if(t ≤ 1 && t > 1) printf("Hi");
```

may print “Hi” even though the if-condition seems unsatisfiable. The reason is that t is uninitialized and hence returns \mathbf{u} in each use, which can be used to satisfy the condition.

Strange though it may seem, LLVM’s treatment of uninitialized reads is allowed by the C standard, which says that performing any computation with a value returned by an uninitialized read results in *undefined* behavior.

2.2 The Semantics of Data Races

A program execution is racy if it has two concurrent memory accesses to the same location, such that at least one of them is a write and at least one of them is non-atomic. There are two types of races: *read-write* races, and *write-write* races.

A read-write race occurs between a load and a store or update operation. The intended semantics for LLVM is that

¹This is needed to justify the distributivity of $+$ over $*$. Consider the transformations: $\mathbf{u} * 0 \rightsquigarrow \mathbf{u} * (1 - 1) \rightsquigarrow \mathbf{u} * 1 + \mathbf{u} * (-1) \rightsquigarrow \mathbf{u} + \mathbf{u} \rightsquigarrow \mathbf{u}$.

in such cases the racy load returns \mathbf{u} . Although stronger than the C11 model, where races result in undefined behavior, the LLVM semantics may still lead to unintuitive behavior. Consider the following program where initially $Y = 0$.

```
YNA = 1; || t = YNA; if(t ≤ 1 && t > 1) printf("Hi");
```

As with the program in §2.1, the current program may also print “Hi” just because the non-atomic load of Y is racy and thus returns \mathbf{u} . The reason that the treatment of read-write races in the LLVM semantics differs from that in C11 is because LLVM readily performs the following transformation

$$\text{if}(cond) t = X_{\text{NA}}; \rightsquigarrow t' = X_{\text{NA}}; t = cond ? t' : t;$$

that converts a conditional branch into a conditional move instruction. This transformation may, however, introduce a read-write race if there were some other parallel thread writing to X only when the condition $cond$ is false. The transformation is correct because the target execution uses the racy read value only when the source execution is also racy.

A write-write race occurs whenever both of the accesses racing with one another are stores or updates. In this case, the intended semantics according to the LLVM documentation [19, Section “Optimization outside atomic”] is the same as in C11: even a single consistent execution with a write-write race results in the program having *undefined behavior*. This semantics allows the read-after-write elimination over an acquire access as shown in the following example:

$$\begin{aligned}
&X_{\text{NA}} = 4; \quad \parallel \quad \left\| \begin{array}{l} X_{\text{NA}} = 8; \\ Y_{\text{REL}} = 1; \\ t = X_{\text{NA}}; \end{array} \right. \rightsquigarrow \text{if}(Y_{\text{ACQ}}) \left\| \begin{array}{l} X_{\text{NA}} = 4; \\ t = 4; \end{array} \right. \parallel \left\| \begin{array}{l} X_{\text{NA}} = 8; \\ Y_{\text{REL}} = 1; \end{array} \right.
\end{aligned}$$

Because of the write-write race on X , the source program has undefined behavior, and hence the transformation is trivially sound. If, however, write-write races were not considered to be undefined behavior, but rather that one of the accesses occurred before the other, then the transformation would be unsound, because in the source program, t would have to contain the final value of X (which may well be 8).

This optimization was performed by LLVM version 3.6 but was later dropped in version 3.7 while fixing another concurrency bug (Bug #22514 [6]). This demonstrates that it is important for LLVM to have a clear concurrency semantics because it affects the validity of basic optimizations.

3. The Formal LLVM Concurrency Model

In this section, we present our formalization of LLVM’s concurrency model.

Events The unit of execution in our model is called an *event* and represents either a shared memory access or the creation of a thread. An event, $e = \langle id, C, lab \rangle$, is a tuple consisting of a (unique) identifier for the event, a command representing the thread’s continuation after the event, and a

label representing the type of the event. We use the notation $e.id$, $e.code$, and $e.lab$ to return the components of an event e . Event labels, lab , are given by the following grammar:

$$lab ::= Ld_o(\ell, v) \mid St_o(\ell, v) \mid U_o(\ell, v, v') \mid Init$$

We have loads (Ld), stores (St), updates (U), and thread initialization events. A load event is generated from a load or an unsuccessful CAS access, a store is generated from a store access, and update events result from a successful CAS operation. Load, store, and update labels record the location accessed (ℓ), the values read and/or written (v, v'), and the annotated memory order (o).

A *read event* is either a load or an update, whereas a *write event* is either a store or an update. Let \mathcal{R} denote the set of all read events and \mathcal{W} the set of all write events. Given an event e , we write $e.lab$ to return the type of the event (Ld, St, U, Init), and $e.loc$ (resp. $e.ord$) to return its location (resp. memory order), whenever applicable. If e is a read event, $e.rval$ returns the value read by e . Similarly, if e is a write event, $e.wval$ denotes the value written by e .

Event Structures A *prime event structure* [27] consists of a set V of events equipped with two relations: the program order and the conflict relation. The program order, po , is a strict partial order on events recording when an event precedes another one in the program, whereas the conflict relation, cf , is a symmetric irreflexive relation denoting that two events cannot belong to the same execution. The conflict relation is assumed to be forward-closed with respect to the program order (i.e., $cf; po \subseteq cf$), which intuitively means that if two events conflict, then so are all their future events.

In this paper, we use an extension of prime event structures, which we call *memory event structures* (or event structures, for short). A memory event structure is a prime event structure with an additional component, the *reads-from relation*, rf , that relates a write to the read events that read from it. We require that whenever $rf(w, r)$, then $w.wval \preceq r.rval$ (racy reads may return \mathbf{u}). We denote an event structure G as a tuple $\langle V, po, rf \rangle$ where $G.V$, $G.po$, and $G.rf$ denote the respective components of G . (We often omit the “ G .” when G is clear from the context.) In our setting, the conflict relation, $G.cf$ is a derived relation:

$$G.cf \triangleq (G.po^{-1}; G.po) \setminus (G.po^? \cup G.po^{-1})$$

It relates all events that are unordered by the program order, but have a common po -ancestor in the same thread. Immediate conflicts are created by read events corresponding to the same program command, but which return different values.

Auxiliary Definitions We define the sets of non-atomic accesses (NA), sequentially consistent accesses (SC), acquire (or stronger) accesses (Acq), and release (or stronger) accesses (Rel).

$$\begin{aligned} NA &\triangleq \{e \mid e.ord = NA\} & SC &\triangleq \{e \mid e.ord = SC\} \\ Acq &\triangleq \{e \mid e.ord \sqsupseteq ACQ\} & Rel &\triangleq \{e \mid e.ord \sqsupseteq REL\} \end{aligned}$$

We say that a release write *synchronizes* (sw) with an acquire read that reads from it. An event a *happens before* an event b , if a reaches b by a path of po or sw edges.

$$G.sw \triangleq [Rel]; rf; [Acq] \quad G.hb \triangleq (po \cup sw)^+$$

Two events race with one another (race) if they are concurrent accesses to the same location (i.e., neither happens before the other), at least one of them is non-atomic and at least one of them is a write.

$$locs \triangleq \{(e, e') \mid e.loc = e'.loc\}$$

$$G.race \triangleq (locs \cap \mathbf{one}(\mathcal{W}) \cap \mathbf{one}(NA)) \setminus (hb^? \cup hb^{-1})$$

An event structure is write-write racy if it contains two write events that race with each other.

$$WWrace(G) \triangleq \exists w, w' \in \mathcal{W}. G.race(w, w')$$

Next, $G.hbW(e)$ checks if the location read by e is initialized in the event structure G ; that is, if there exists a write to it that happens before e .

$$G.hbW(e) \triangleq \exists w \in \mathcal{W}. G.hb(w, e) \wedge e.loc = w.loc$$

Finally, $AddRF$ creates an rf edge between two events in an event structure, G , and returns the updated event structure.

$$AddRF(G, e, e') \triangleq \langle G.V, G.po, G.rf \cup \{(e, e')\} \rangle$$

3.1 Event Structure Construction

The event structures of a program are constructed incrementally with the operational semantics shown in Figure 14. For a program $P = (X_1 = v_1; \dots; X_k = v_k; \{C_1 \parallel \dots \parallel C_n\})$, we define the program’s initial event structure, $G_{init}(P)$, to be $\langle A \cup B, A \times B, \emptyset \rangle$ where $A = \{a_1, \dots, a_k\}$ with each a_i having label $St(X_i, v_i)$ and empty continuation and $B = \{b_1, \dots, b_n\}$ with each b_i having label $Init$ and continuation C_i .

Each rule from Figure 14 then takes an event structure G and typically extends it by adding one more event to it. An exception is the WW -RACE rule, which checks whether the event structure has a *write-write race*. If so, the program behavior is *undefined* and thus the program can produce any arbitrary event structure G' .

The other rules first call the helper rule $BASIC$, which selects an event e from the event structure and executes its continuation $e.code$ to get a possible next event e' to be added. The new event has a fresh identifier (i.e., $e'.id$ does not exist in the event structure), and must have not already been added to the event structure (i.e., there does not exist another event e'' with the same label as e' immediately after the previous event, e). Finally, the new event has its label and continuation determined by the thread semantics (which is a parameter to the memory model). Assuming all these premises hold, the rule adds e' to the event structure, recording that is immediately after e in program order.

$$\begin{array}{c}
\frac{e \in V \quad \forall e'' \in V. e'.id \neq e''.id \quad e.code \xrightarrow{e'.lab} e'.code \quad \forall e''. po(e, e'') \implies e'.lab \neq e''.lab}{\langle V, po, rf \rangle \xrightarrow{e'} \langle V \cup \{e'\}, (po \cup \{(e, e')\})^+, rf \rangle} \text{ (BASIC)} \\
\\
\frac{WWrace(G)}{G \rightsquigarrow G'} \text{ (WW-RACE)} \quad \frac{G \xrightarrow{e'} G' \quad e' \notin \mathcal{R}}{G \rightsquigarrow G'} \text{ (NON-READ)} \\
\\
\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.rval = \mathbf{u} \quad \neg G'.hbW(e')}{G \rightsquigarrow G'} \text{ (R-UNINIT)} \\
\\
\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.rval = \mathbf{u} \quad G'.hbW(e') \quad \exists w. G'.race(w, e')}{G \rightsquigarrow G'} \text{ (R-RACE)} \\
\\
\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.rval = w.wval \quad G'.hbW(e') \quad \neg G'.race(w, e') \quad G'' = \text{AddRF}(G', w, e') \quad \text{isCons}(G'')}{G \rightsquigarrow G''} \text{ (R-NORACE)}
\end{array}$$

Figure 5. Event structure reduction steps.

If the new event is a store, it is simply added to the event structure (NON-READ). If, however, it is a read (i.e., a load or an update), we need also need to check that the value read is correct. There are three cases to consider:

- The location is *uninitialized*, namely there does not exist a write to that location that happens before e . In this case, the read must return the undefined value \mathbf{u} . (R-UNINIT)
- The location is initialized, but the access *rac*es with some other write w . In this case, the read again returns the special undefined value \mathbf{u} . (R-RACE)
- The location is initialized and there exists a *non-racy* write w from which the new event can read. Here, we extend $G'.rf$ with the edge (w, e') to record that e' reads from w , and insist that $e'.rval = w.wval$. We then check that the resulting event structure is consistent (see §3.2) and discard it otherwise. (R-NORACE)

The rules can be applied in any sequence, which may result in multiple event structures. We define $\llbracket P \rrbracket_{\text{LLVM}}$ to return the set of all event structures generated from $G_{\text{init}}(P)$; that is, $\llbracket P \rrbracket_{\text{LLVM}} \triangleq \{G \mid G_{\text{init}}(P) \rightsquigarrow^* G\}$.

3.2 Event Structure Consistency Checking

We now move on to the consistency checking of an event structure G by $\text{isCons}(G)$. A key constraint to check is that each write justifying a read in G is indeed *visible* to the read.

Overwritten Writes We first consider cases of writes that are *not* visible. One simple case is that of overwritten writes. For example, in the program $X = 1; X = 2; t = X$; the




Figure 6. Behavior of Coh forbidden due to a wb cycle.

load of X should only be able to read from the second store and return the value 2.

More generally, we must rule out executions breaking *coherence* (a.k.a. “SC per location” [1]). Consider the program:

$$\begin{array}{l}
X_{\text{REL}} = 1; \parallel X_{\text{REL}} = 2; \\
t = X_{\text{ACQ}}; \parallel t' = X_{\text{ACQ}};
\end{array} \quad \text{(Coh)}$$

Here, we ought to prohibit the $t = 2 \wedge t' = 1$ outcome, because it violates coherence. In terms of the C11 memory model, all stores to the same location have to be totally ordered by the modification order, mo . The WR-coherence axiom says that a read cannot read from write earlier in the modification order than some other write that happened before it. Formally, the rf^{-1} ; mo ; hb should be irreflexive.

In our model, we do not record a modification order, mo . Instead we base our model on a derived partial order over writes to the same location, which is called the *writes-before* order, wb , by Lahav and Vafeiadis [14]. We have already seen two cases of the writes-before order: (1) A write w happening before another write w' to the same location induces the writes-before relation $wb(w, w')$. (2) A write w' happens before a same location read r and r reads from a write w induces the writes-before relation $wb(w', w)$. There is one additional case that arises because of the atomicity of update instructions. Consider the following program:

$$\begin{array}{l}
X_{\text{REL}} = 1; \parallel X_{\text{REL}} = 2; \\
t = X_{\text{ACQ}}; \parallel t' = X_{\text{ACQ}}; \parallel \text{CAS}_{\text{ACQ_REL}}(X, 2, 3);
\end{array} \quad \text{(UCoh)}$$

Here, the outcome $t = 3 \wedge t' = 1$ should again be forbidden, because it violates coherence and/or the atomicity of updates. In C11 terms, since $t = 3$, this means that the $X = 1$ store must precede the CAS in modification order. Similarly, since $t' = 1$, the $X = 2$ store must precede the $X = 1$ in modification order. The atomicity of updates, however, says the CAS must immediately follow the $X_{\text{REL}} = 2$ store in modification order; so the $X = 1$ store must precede the $X = 2$ store, which leads to a contradiction.

We therefore define writes-before as follows:

$$G.wb \triangleq ([\mathcal{W}]; ((brf; (hb \cap locs); brf) \setminus brf); [\mathcal{W}])^+$$

where $brf \triangleq (rf^{-1})^*$. To prohibit coherence violations, we insist that wb be irreflexive. This rules out the inconsistent behaviors of Coh (see Figure 6) and UCoh.




Figure 7. Returning $t = 1$ in Rconflict should be forbidden.




Figure 8. Conflicting write forbids behavior of IncLoop.

Conflicting Writes A second class of writes that a read cannot possibly read from are conflicting writes. For example, in the program

$$X_{REL} = 0; t = X_{ACQ}; X_{REL} = 1; \quad (\text{Rconflict})$$

t should only be able to read 0, and not 1 from the conflicting $X = 1$ store as shown in Figure 7.

Slightly generalizing this condition, we disallow conflicts between hb-related events, i.e. require $cf; hb$ to be irreflexive. This rules out strange behaviors of examples like the program IncLoop. where initially $X = Y = 0$.

$$X_{REL} = Y_{ACQ} + 1; \parallel Y_{REL} = X_{ACQ}; \quad (\text{IncLoop})$$

Returning $X > 1$ is forbidden according to release-acquire semantics. Figure 8 shows that indeed reading $Y = 1$ in the first thread is impossible. If it were allowed, then b would have taken place in conflict with a and in turn would result in $X = 2$ as a possible outcome. Clearly, this would be an *out-of-thin-air* value in any consistent execution.

Non-Conflicting Justifications To avoid weird behaviors reminiscent of “out-of-thin-air” values, we forbid conflicting writes to justify non-conflicting hb-related reads. For example, in the following program $Z = 2$ should not be a valid outcome, where all locations are initialized to zero.

$$Z_{REL} = 1; \parallel \begin{cases} \text{if}(Z_{ACQ}) Y_{REL} = 1; \\ \text{else } X_{REL} = 1; \end{cases} \parallel \begin{cases} \text{if}(X_{ACQ} \ \&\& \ Y_{ACQ}) \\ Z_{REL} = 2; \end{cases} \quad (\text{Cwrites})$$

As shown in Figure 9, if a justifies c and b justifies d then $St_{REL}(Z, 2)$ would be an event in the event structure which is actually should not have happened. To restrict such cases, we check that $rf; hb^{-1}; rf^{-1}; cf$ is irreflexive.

Preserving the Order of Sequential Consistent Accesses The LLVM specification states that the semantics of SC accesses follows that of C11, which puts all the SC events in a total order and applies certain constraints on SC-reads. However, it has recently been shown that the C11 semantics of SC accesses is broken in that the expected compilation




Figure 9. Conflicting justifier forbids behavior of Cwrites.




Figure 10. SC constraint forbids $t = t' = 0$ in SB.

schemes to Power and ARM are unsound [17, 21]. (The recent strengthening of Batty et al. [4] is also broken for the same reason.)

Therefore, instead of the C11 semantics, we adapt the solution of Lahav et al. [17]. In their solution, one checks for the acyclicity of a union of relations on the SC events. We say that a read a reads before a (different) write b if it reads from a write that was written before b . Formally,

$$G.fr \triangleq (rf^{-1}; wb) \setminus [G.V] \quad (\text{reads before})$$

The “[$G.V$]” takes care of updates because updates are $(rf^{-1}; mo)$ -before themselves. (This definition is a slight adaptation of the definition in Lahav et al. [17] that uses wb instead of the modification order.)

We next define hb_{sc} to denote a restricted hb path between two SC events that either (i) does not change location, or (ii) starts and ends with a program order edge that changes location, which in turn ensures that the compilation to Power and ARM will have an appropriate fence along the path.

$$G.po|_{\neq locs} \triangleq po \setminus locs$$

$$G.hb_{sc} \triangleq [SC]; \left((hb \cap locs) \cup po|_{\neq locs} \cup (po|_{\neq locs}; hb; po|_{\neq locs}) \right); [SC]$$

We then require $(hb_{sc} \cup wb \cup fr); [SC]$ to be acyclic in the event structure. To illustrate this condition, consider the store buffer program with SC accesses, where the outcome $t = t' = 0$ should be forbidden.

$$X_{SC} = 1; \parallel Y_{SC} = 1; \quad (\text{SB})$$

$$t = Y_{SC}; \parallel t' = X_{SC};$$

The weak behavior is ruled out because the relevant event structure (see Figure 10) contains a $(po \cup fr); [SC]$ cycle.

A slightly more complex example is the following, where the outcome $t = 2 \wedge t' = 0$ should also be forbidden.

$$X_{SC} = 1; \parallel Y_{SC} = 2; \quad (\text{SCR})$$

$$Y_{SC} = 1; \parallel t' = X_{SC};$$

$$t = Y_{ACQ};$$




Figure 11. SC constraint forbids $t = 2 \wedge t' = 0$ in SCR.

Again this is so because of a $(\text{po} \cup \text{wb} \cup \text{fr})$; [SC] cycle in the relevant event structure (see Figure 11).

Definition of isCons We say that an event structure G is *consistent* if it satisfies all the aforementioned constraints:

$$\begin{aligned} \text{isCons}(G) \triangleq & \text{irreflexive}(\text{wb}) \wedge \text{irreflexive}(\text{cf}; \text{hb}) \\ & \wedge \text{irreflexive}(\text{rf}; \text{hb}^{-1}; \text{rf}^{-1}; \text{cf}) \\ & \wedge \text{acyclic}((\text{hb} \cup \text{wb} \cup \text{fr}); [\text{SC}]) \end{aligned}$$

3.3 Consistent Executions

So far, we have discussed the construction of the event structures of a program. Since an event structure may contain events arising from multiple executions of a program, we use the function $\text{exec}(G)$ to extract individual (consistent) executions from a fully constructed event structure.

$$\text{exec}(G) \triangleq \left\{ \begin{array}{l} \mathbf{E} = \langle A, G.\text{po} \cap (A \times A), \text{rf} \rangle \mid \text{isCons}(\mathbf{E}) \\ \wedge A \subseteq G.V \wedge G.\text{hb}; [A] \subseteq (A \times A) \\ \wedge [A]; \text{cf}; [A] = \emptyset \wedge [A]; G.\text{rf}; [A] \subseteq \text{rf} \\ \wedge \forall r \in A. (\exists w. \text{rf}(w, r)) \iff G.\text{hbW}(r) \end{array} \right\}$$

Each execution, $\mathbf{E} \in \text{exec}(G)$, is a consistent conflict-free hb-prefix of the event structure, whose reads-from relation has been extended to provide a justification for each initialized read in the prefix. Note that executions of an event structure constructed by the operational semantics, unlike the event structure itself, may have $(\text{po} \cup \text{rf})$ cycles.

The consistency check together with the requirement that all initialized read events be justified removes some strange behaviors that would otherwise be allowed. Consider the following program, in which the outcome $t = 0 \wedge t' = 1$ should be forbidden.

$$\begin{array}{l} t = Z_{\text{ACQ}}; \\ \text{if}(t) \\ \quad X_{\text{REL}} = 1; \end{array} \parallel \begin{array}{l} \text{if}(Y_{\text{ACQ}}) \\ \quad t' = X_{\text{ACQ}}; \end{array} \parallel \begin{array}{l} Z_{\text{REL}} = 1; X_{\text{REL}} = 1; \\ X_{\text{REL}} = 2; Y_{\text{REL}} = 1; \end{array} \quad (\text{CEX})$$

Consider, for example, the candidate execution of this program highlighted in Figure 12. The highlighted execution is, however, cannot be made consistent because it has to provide a justification for the $\text{Ld}_{\text{ACQ}}(X, 1)$ event. Its only possible justification is the $\text{St}_{\text{REL}}(X, 1)$ event, but this would violate coherence.

3.4 Observable Behaviors

We take the observable behavior of an execution (a conflict-free event structure) to be the set of the last values written to




Figure 12. Forbidden execution $t = 0 \wedge t' = 1$ in CEX.

each variable, namely by writes that were not written before some other write (to the same location).

$$\begin{aligned} \text{Behavior}(G) \triangleq \\ \{ (\ell, v) \mid \exists e \in G.V. v \preceq e.\text{wval} \wedge \nexists e'. G.\text{wb}(e, e') \} \end{aligned}$$

The \preceq in the definition above allows treating \mathbf{u} as any possible value. The LLVM-behaviors of a program P are those of its consistent executions.

$$\text{Behavior}_{\text{LLVM}}(P) \triangleq \{ \text{Behavior}(G) \mid G \in \text{exec}(\llbracket P \rrbracket_{\text{LLVM}}) \}$$

This definition can straightforwardly be extended to other memory models such as (O)SC, RA, C11, TSO, and Power by substituting the appropriate event structure consistency definition. In these execution-based axiomatic models each execution can be considered as conflict-free event structure.

4. DRF Theorems

In this section, we prove that the proposed LLVM model satisfies two expected DRF theorems: DRF-RA and DRF-OSC. These theorems enable developers programming over the model to follow a defensive style of programming and avoid the need of understanding the model. Although our model's intended use is for an intermediate language, our DRF theorems can be seen as sanity checks that the model is not overly weak. For readability, we opted for an informal high-level presentation of the theorems in this section; formal statements and proofs can be found in [8].

4.1 LLVM and Release-Acquire Consistency

Release-acquire (RA) consistency is a strengthening of our model, where during the event structure construction each initialized read *reads from* some *happens-before* write. In other words, the event structure is constructed without using the R-RACE rule. Our first theorem says that if a program does not have any read-write races under RA semantics, then RA-consistency and LLVM-consistency produce the same set of event structures. A trivial corollary of DRF-RA is that if a program contains only atomic operations, then its behaviors under LLVM and RA coincide.

Theorem 1 (DRF-RA). *If under RA consistency, a program has no read-write races, then its LLVM-consistent behaviors coincide with its RA-consistent ones.*

Considering the scenarios based on the read-write races, Theorem 1 follows from the following lemmas.

Lemma 1. *If an LLVM-consistent event structure G has no read-write races, then G is also RA-consistent.*

Lemma 2. *Given a program P and an LLVM-consistent event structure of P with a read-write race, there exists some RA-consistent event structure of P with a read-write race.*

4.2 LLVM and (Observable) Sequential Consistency

We model *observable sequential consistency* (OSC) by dropping the WW-RACE rule in the event structure construction and treating all the operations as having SC memory order in the graph consistency checks, $\text{isCons}(G)$. We note that this definition is technically slightly different from the standard SC definition [18]; in particular, writes that are never observed (i.e. read) may appear somewhat loosely ordered. Concretely, consider the following program:

$$\begin{array}{l} X_{\text{sc}} = 1; \parallel Y_{\text{sc}} = 1; \\ Y_{\text{sc}} = 2; \parallel X_{\text{sc}} = 2; \end{array} \quad (2+2W)$$

Under SC, the final result cannot be $X = Y = 1$. Yet it is a OSC-behavior: since there are no reads in the program, $\text{wb} = \emptyset$. Intuitively, however, OSC and SC are essentially the same. We conjecture that this is also formally so if we restrict the program behaviors to observations made by the program itself (e.g., recorded in thread-local variables).

We say that a program is RA-race-free under a memory model M , if all concurrent accesses to the same location are either both loads or both SC. Our main theorem states that LLVM-consistency and sequential consistency coincide for programs that are RA-race-free under OSC.

Theorem 2 (DRF-OSC). *If a program is RA-race-free under OSC, then OSC and LLVM consistency coincide.*

Note that the RA-race-freedom implies lack of read-write races. Hence, it suffices to prove correspondence between the RA-consistency and OSC in absence of RA races. To do so, we prove the following lemmas.

Lemma 3. *If an RA-consistent event structure G is RA-race-free, then G is OSC-consistent.*

Lemma 4. *For each RA-racy LLVM-consistent event structure there is a RA-racy OSC-consistent event structure.*

5. Compilation from C11 to LLVM

In this section, we discuss the correctness of the LLVM’s front-end, `clang`, as far as concurrency is concerned. Specifically, we show that the straightforward mapping from C11 memory accesses to LLVM memory accesses used by `clang` is correct with respect to our model.²

Compilation correctness means that the behaviors of the target program are included in those of the source program.

²Recall that LLVM’s memory orders include those of C11’s except for *consume* reads, which `clang` converts to the strictly stronger *acquire* reads.

$\downarrow a \setminus b \rightarrow$	St_{REL}	Ld_{ACQ}	U	$(\text{St} \text{Ld})_{\text{NA}}$	Ld_{SC}
St_{REL}	×	✓	×	✓	✓
St_{SC}	×	✓	×	✓	×
Ld_{ACQ}	×	×	×	×	×
U	×	×	×	×	×
$(\text{St} \text{Ld})_{\text{NA}}$	×	✓	×	✓	✓

Table 1. Safe and unsafe memory access reorderings.

Theorem 3. *For every program P , $\text{Behavior}_{\text{LLVM}}(P) \subseteq \text{Behavior}_{\text{C11}}(P)$.*

First, for the theorem to even hold, we assume that C11’s treatment of SC accesses is changed to match the one in our LLVM model (cf. §3.2). Then, the proof of this theorem is straightforward because the remaining difference between the models is only in the treatment of read-write races: in LLVM, such reads may return `u` (using R-RACE), whereas in C11, they produce completely arbitrary behavior.

6. Program Transformations/Optimizations

Next, we consider program transformations that may be performed as part of the LLVM optimization passes, and prove their correctness. Generally, a transformation is correct if it does not introduce any new behaviors; namely, if every behavior of the target program is also a behavior of the source program. For space reasons, we just give an overview of the results in this section: the actual proofs can be found in our technical appendix [8].

6.1 Reordering of Independent Memory Accesses

One standard transformation is to reorder two independent instructions, which may improve the locality of computation or enable further optimizations. Two independent adjacent instructions a and b are *safely reorderable* if the following conditions hold: (i) do not access the same location, (ii) they are not both SC, (iii) the first instruction, a , is not an acquire read, and (iv) the second instruction, b , is not a release write. Table 1 summarizes the safe (✓) and unsafe (×) reordering transformations. For the unsafe ones, we provide counterexamples in our technical appendix [8].

Theorem 4. *Reordering two adjacent safely reorderable instructions is a correct transformation.*

To prove this theorem, we show that the safe reordering conditions ensure that the target event structure has fewer behaviors compared to the source event structure. Since, however, each memory access instruction in a program may produce multiple conflicting events in the event structure, reordering two such instructions may result in quite different every structure that are not obvious to relate at first.

Let A and B be the sets of conflicting events generated by instructions a and b respectively. In general, the events in these two sets do not have one-to-one po-correspondence, since an event in A may have multiple immediate po-

successors (all must be in B). To relate the two event structures, we first *expand* the predecessor event set such that the events in these two sets have one-to-one po-relation and then perform the reordering and finally *collapse* the *expanded* event set to create the target event structure.

6.2 Elimination Optimizations

Next, we consider the possible safe deletions of the redundant shared memory accesses in the proposed LLVM model. Some of these transformations are performed in the common subexpression elimination (CSE) or common subexpression elimination optimization passes.

We begin with the deletion of accesses because of another adjacent access to the same location. We have three cases:

Overwritten Write (OW) Deletes the first of two same location adjacent store operations:

$$\text{St}_{o'}(\ell, v'); \text{St}_o(\ell, v) \rightsquigarrow \text{St}_o(\ell, v) \quad \text{when } o' \sqsubseteq o.$$

We prove this transformation correct by simulation: we match the target execution steps by the exact same source execution steps, except when the target inserts the event corresponding to the $\text{St}_o(\ell, v)$ store. In this case, we perform two execution steps in the source, adding events for both the eliminated and the remaining store.

Read after Write (RAW) Deletes a load that immediately follows a store to the same location:

$$\text{St}_o(\ell, v); \text{Ld}_{o'}(\ell) \rightsquigarrow \text{St}_o(\ell, v) \quad \text{when } o' \sqsubseteq o.$$

We prove this transformation correct again by simulation: when the target inserts the event corresponding to the $\text{St}_o(\ell, v)$ store, the source also inserts a $\text{Ld}_{o'}(\ell, v)$ event reading from that store.

Read after Read (RAR) Deletes the second of two adjacent same location reads.

$$\text{Ld}_o(\ell); \text{Ld}_{o'}(\ell) \rightsquigarrow \text{Ld}_o(\ell) \quad \text{when } o' \sqsubseteq o.$$

This transformation is correct and in the proof, we simulate the respective load step by two load steps reading from the same write.

In all these cases, LLVM actually performs these transformations only if the deleted instruction is non-atomic.

An access may also be deleted if the access justifying the deletion is non-adjacent, but can be made adjacent by reordering the access to be deleted over all the intermediate accesses. For example, in the program below, the second load of ℓ can be eliminated as follows:

$$\begin{array}{l} a : \text{Ld}_{\text{NA}}(\ell); \quad a : \text{Ld}_{\text{NA}}(\ell); \\ b : \text{St}_{\text{REL}}(X, 1); \rightsquigarrow c : \text{Ld}_{\text{NA}}(\ell); \quad \rightsquigarrow a : \text{Ld}_{\text{NA}}(\ell); \\ c : \text{Ld}_{\text{NA}}(\ell); \quad b : \text{St}_{\text{REL}}(X, 1); \quad b : \text{St}_{\text{REL}}(X, 1); \end{array}$$

Note that if instead of the release store, there were an acquire load of X between the two loads of ℓ , then the second load

cannot be eliminated in this way, because reordering it with the acquire load is unsafe. It may, however, be eliminated, if the first load is moved after the acquire:

$$\begin{array}{l} a : \text{Ld}_{\text{NA}}(\ell); \quad b : \text{Ld}_{\text{ACQ}}(X); \quad b : \text{Ld}_{\text{ACQ}}(X); \\ b : \text{Ld}_{\text{ACQ}}(X); \rightsquigarrow a : \text{Ld}_{\text{NA}}(\ell); \quad \rightsquigarrow a : \text{Ld}_{\text{NA}}(\ell); \\ c : \text{Ld}_{\text{NA}}(\ell); \quad c : \text{Ld}_{\text{NA}}(\ell); \end{array}$$

Note that in this case the remaining load of ℓ cannot be moved back in its original place.

There are two more cases of eliminating non-adjacent non-atomic accesses, which have a somewhat weaker correctness condition, namely that there has to be no *release-acquire pair* between the eliminated access and the justifying access, which has to be a non-atomic store. A *release-acquire pair* is formed by a release write followed by an acquire read access. The transformations are:

$$\text{St}_{\text{NA}}(\ell, v'); C; \text{St}_{\text{NA}}(\ell, v) \rightsquigarrow C; \text{St}_{\text{NA}}(\ell, v) \quad (\text{NA-OW})$$

$$\text{St}_{\text{NA}}(\ell, v); C; \text{Ld}_{\text{NA}}(\ell) \rightsquigarrow \text{St}_{\text{NA}}(\ell, v); C \quad (\text{NA-RAW})$$

where C does not have any *release-acquire* pair nor any accesses of location ℓ . For NA-OW, these conditions ensure that the only reads reading from the eliminated store are actually racy, and hence do not need to read from there. Similarly, for NA-RAW, one can show that the eliminated load can always read from the preceding store because there can be no other same-location write in between.

6.3 Speculative Load Introduction

One sound and occasionally useful transformation is to insert a load, whose value is just never used. Assuming that the address, whence the inserted load reads, is valid and allocated, this transformation is trivially correct according to our LLVM model. The load introduction is, however, incorrect according to the C11 model, because it may introduce a data race (and hence undefined behavior in C11). LLVM frequently performs such load introductions in the “simplify CFG” pass; e.g., when hoisting loads outside of conditionals.

6.4 Access Strengthening

Finally, a sound, but not so useful, transformation is strengthening the memory order of memory accesses, i.e., converting some access_o of the program into $\text{access}_{o'}$ where $o \sqsubseteq o'$. For example, a non-atomic write may be changed into a release write. Soundness of access strengthening follows directly from the *monotonicity* property of the LLVM model. Given $P \xrightarrow{\text{strengthen}} P'$, we show that that for each $G' \in \llbracket P' \rrbracket_{\text{LLVM}}$, then we can construct a similar event structure in $\llbracket P \rrbracket_{\text{LLVM}}$ (actually, the same modulo memory orders).

7. Compilation to x86-TSO and Power

The LLVM compiler generates target code for various architectures including x86 and Power. In this section, we show that the standard compilation schemes from LLVM to x86-TSO and Power shown in Figure 13 are correct.

LLVM	x86-TSO	Power
store _{NA}	mov	stw
store _{REL}	mov	lwsync; stw
store _{SC}	mov; mfence	hwsync; stw
load _{NA}	mov	lwz
load _{ACQ}	mov	lwz; lwsync
load _{SC}	mov	hwsync; lwz; lwsync
CAS _{ACQ_REL}	lock cmpxchg	lwsync; M ; lwsync
CAS _{SC}	lock cmpxchg	hwsync; M ; lwsync

where $M \triangleq L$: lwarx; cmpw; bne L' ; stwcx.; bne L ; L' :

Figure 13. Compilation schemes to x86 and Power.

7.1 Compilation to x86-TSO

In x86, almost all shared accesses get compiled down to plain mov instructions, which serve as both loads and stores. The two exceptions are: (1) sequentially consistent stores, whose compilation includes a memory fence (mfence) after the actual store, and (2) atomic updates, such as compare-swap, which get compiled to special ‘locked’ instructions.

The correctness of compilation to TSO follows easily from our existing results. First, by monotonicity, we can strengthen all non-atomic accesses of a program to become release stores or acquire loads. (Note that the compilation schemes for non-atomic accesses and the corresponding release/acquire accesses are identical, so we do not incur any performance penalty by this strengthening.) Next, by the DRF-RA theorem, since there are no non-atomic accesses left in the program (and thus no races), the semantics of the program according to the LLVM model corresponds exactly to that according to the RA model. Further, it is well-known that aforementioned compilation scheme to TSO is correct for release/acquire and SC accesses [2]. Putting everything together, we get that the compilation mappings from LLVM to x86-TSO [22] is correct.

7.2 Compilation to Power

The shared memory access instructions in the Power architecture are load (lwz), store (stw), load-linked (lwarx) and store-conditional (stwcx) along with various fence instructions such as hwsync (hardware sync), lwsync (lightweight sync), isync (instruction sync).

As mappings in the Figure 13 show, the non-atomic accesses do not require any fence, release and SC stores have lightweight and full fence respectively before the store. The acquire and SC loads place an lwsync after the load.³ In addition, SC loads also place a full fence before the load. Updates are implemented using a loop with load-linked and store-conditional instructions, and have fences at both ends.

For the correctness proof, we use the empirically validated axiomatic memory model of Power by Alglave et al.

³ An alternative—possibly more efficient—compilation scheme for acquire loads places a control-dependency to an isync fence instead of an lwsync fence. LLVM, however, does not yet implement this other scheme.

[1] and a recent result by Lahav and Vafeiadis [15]. This result reduces the correctness of compilation to Power to the correctness of reordering independent plain memory accesses of different locations and the correctness of compilation to a stronger model, SPower, which strengthens the Power model of Alglave et al. [1] with the acyclic ($po \cup rf$) requirement. The full definition of SPower and the LLVM to SPower compilation correctness proof can be found in [8].

8. Related Work and Conclusion

Our work is the first to formalize a non-trivial fragment of the concurrency model of LLVM. Earlier work [20, 28] has studied LLVM’s memory model only for sequential programs. Our formal concurrency model is based on the informal descriptions in the LLVM language reference manual [19]. These descriptions often refer to the C/C++11 memory model [10, 11], which they follow quite closely, both in terms of concurrency primitives and reordering constraints. The major difference, as discussed in Chakraborty and Vafeiadis [7], is the treatment of the read-write data races, whose modelling has been the main subject of this paper. The LLVM informal specification also provides constructs corresponding to C/C++11 relaxed and Java ordinary accesses, which are beyond the purview of this work.

Prior to our work, Pichon-Pharabod and Sewell [23] and Jeffrey and Riely [12] introduce memory models based on event structures as attempts to solve the *out-of-thin-air* problem of relaxed memory models (see [3, 5] for discussions about the problem and some of its implications). Both models attempt to capture all the executions of a program in a single event structure and start with an event structure recording all the potential executions of a program.

Pichon-Pharabod and Sewell [23] have an operational semantics that gradually simplifies the event structure, by either committing an event and pruning the event structure, or transforming it as part of an optimization step. Our operational semantics is of a very different flavor to theirs; we do not attempt to prune a big event structure, but rather enlarge a small event structure.

Jeffrey and Riely [12] target Java and therefore do not guarantee read-read coherence. Their model, however, fails to validate the reordering of independent read events, as shown in §8 of their paper. LLVM’s semantics for read-write races is weaker than that of Java, and thus we do not have any problems with read-read reorderings.

Finally, Kang et al. [13] very recently proposed a promising solution to the out-of-thin-air problem, which is based on operational semantics with timestamps and a special reduction step that allows a thread to make a locally certifiable promise to perform a write. While the set up there is quite different from our event structures, it would be extremely useful if one can combine the approaches to extend our semantics to handle LLVM’s monotonic accesses. We leave this as future work.

References

- [1] J. Alglave, L. Maranget, and M. Tautschnig. Herding cats: modelling, simulation, testing, and data-mining for weak memory. *ACM Trans. Program. Lang. Syst.*, 36(2):7:1–7:74, 2014. doi: 10.1145/2627752.
- [2] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL'11*, pages 55–66. ACM, 2011. doi: 10.1145/1926385.1926394.
- [3] M. Batty, K. Memarian, K. Nienhuis, J. Pichon-Pharabod, and P. Sewell. The problem of programming language concurrency semantics. In *ESOP'15*, pages 283–307, 2015. doi: 10.1007/978-3-662-46669-8_12.
- [4] M. Batty, A. F. Donaldson, and J. Wickerson. Overhauling SC atomics in C11 and OpenCL. In *POPL '16*, pages 634–648. ACM, 2016. doi: 10.1145/2837614.2837637.
- [5] H.-J. Boehm and B. Demsky. Outlawing ghosts: avoiding out-of-thin-air results. In *MSPC'14*. ACM, 2014. doi: 10.1145/2618128.2618134.
- [6] Bug #22514. Wrong transformation due to semantic gap between C11 and LLVM semantics. https://llvm.org/bugs/show_bug.cgi?id=22514.
- [7] S. Chakraborty and V. Vafeiadis. Validating optimizations of concurrent C/C++ programs. In *CGO'16*, pages 216–226. ACM, 2016. doi: 10.1145/2854038.2854051.
- [8] S. Chakraborty and V. Vafeiadis. Technical appendix, 2016. Available at <http://plv.mpi-sws.org/llvmcs/>.
- [9] S. Flur, K. E. Gray, C. Pulte, S. Sarkar, A. Sezgin, L. Maranget, W. Deacon, and P. Sewell. Modelling the ARMv8 architecture, operationally: concurrency and ISA. In *POPL'16*, pages 608–621. ACM, 2016. doi: 10.1145/2837614.2837615.
- [10] ISO/IEC 14882:2011. Programming language C++.
- [11] ISO/IEC 9899:2011. Programming language C.
- [12] A. Jeffrey and J. Riely. On thin air reads: Towards an event structures model of relaxed memory. In *LICS'16*. ACM/IEEE, 2016. doi: 10.1145/2933575.2934536.
- [13] J. Kang, C.-K. Hur, O. Lahav, V. Vafeiadis, and D. Dreyer. A promising semantics for relaxed-memory concurrency. In *POPL'17*. ACM, 2017.
- [14] O. Lahav and V. Vafeiadis. Owicki-Gries reasoning for weak memory models. In *ICALP'15*, pages 311–323, 2015. doi: 10.1007/978-3-662-47666-6_25.
- [15] O. Lahav and V. Vafeiadis. Explaining relaxed memory models with program transformations. In *FM'16*, pages 479–495, 2016. doi: 10.1007/978-3-319-48989-6_29.
- [16] O. Lahav, N. Giannarakis, and V. Vafeiadis. Taming release-acquire consistency. In *POPL'16*, pages 649–662. ACM, 2016. doi: 10.1145/2837614.2837643.
- [17] O. Lahav, V. Vafeiadis, J. Kang, C.-K. Hur, and D. Dreyer. Repairing sequential consistency in C/C++11. Technical Report MPI-SWS-2016-011, MPI-SWS, 2016.
- [18] L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979. doi: 10.1109/TC.1979.1675439.
- [19] LLVM documentation. LLVM atomic instructions and concurrency guide. <http://llvm.org/docs/Atomics.html>.
- [20] N. P. Lopes, D. Menendez, S. Nagarakatte, and J. Regehr. Provably correct peephole optimizations with Alive. In *PLDI*, pages 22–32. ACM, 2015. doi: 10.1145/2737924.2737965.
- [21] Y. A. Manerkar, C. Trippel, D. Lustig, M. Pellauer, and M. Martonosi. Counterexamples and proof loophole for the C/C++ to POWER and ARMv7 trailing-sync compiler mappings, 2016. arXiv:1611.01507.
- [22] S. Owens, S. Sarkar, and P. Sewell. A better x86 memory model: x86-TSO. In *TPHOLS*, pages 391–407, 2009. doi: 10.1007/978-3-642-03359-9_27.
- [23] J. Pichon-Pharabod and P. Sewell. A concurrency semantics for relaxed atomics that permits optimisation and avoids thin-air executions. In *POPL'16*, pages 622–633. ACM, 2016. doi: 10.1145/2676726.2676995.
- [24] D. E. Shasha and M. Snir. Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.*, 10(2):282–312, 1988. doi: 10.1145/42190.42277.
- [25] R. C. Steinke and G. J. Nutt. A unified theory of shared memory consistency. *J. ACM*, 51(5):800–849, 2004. doi: 10.1145/1017460.1017464.
- [26] V. Vafeiadis and C. Narayan. Relaxed separation logic: A program logic for C11 concurrency. In *OOPSLA'13*. ACM, 2013. doi: 10.1145/2509136.2509532.
- [27] G. Winskel. Event structures. In *Advances in Petri Nets*, pages 325–392. Springer, 1986. doi: 10.1007/3-540-17906-2_31.
- [28] J. Zhao, S. Nagarakatte, M. M. K. Martin, and S. Zdancewic. Formalizing the LLVM intermediate representation for verified program transformations. In *POPL'12*, pages 427–440. ACM, 2012. doi: 10.1145/2103656.2103709.

A. Definitions and Rules

A.1 Definitions

We summarize the auxiliary definitions used in the paper.

Access Types

$$\text{Acq} \triangleq \{e \mid e.\text{ord} \sqsupseteq \text{ACQ}\}$$

$$\text{NA} \triangleq \{e \mid e.\text{ord} = \text{NA}\}$$

$$\text{Rel} \triangleq \{e \mid e.\text{ord} \sqsupseteq \text{REL}\}$$

$$\text{SC} \triangleq \{e \mid e.\text{ord} = \text{SC}\}$$

Preliminary Relations

$$\text{locs} \triangleq \{(e, e') \mid e.\text{loc} = e'.\text{loc}\}$$

$$[A](x, y) \triangleq x = y \wedge x \in A$$

$$\text{one}(A) \triangleq \exists x, y. (x \in A \vee y \in A)$$

$$(R; S)(x, y) \triangleq \exists z. R(x, z) \wedge S(z, y)$$

Relations on Event Structure

$$\text{brf} \triangleq (\text{rf}^{-1})^*$$

$$G.\text{cf} \triangleq (G.\text{po}^{-1}; G.\text{po}) \setminus (G.\text{po}^? \cup G.\text{po}^{-1})$$

$$G.\text{sw} \triangleq [\text{Rel}]; G.\text{rf}; [\text{Acq}]$$

$$G.\text{hb} \triangleq (G.\text{po} \cup G.\text{sw})^+$$

$$G.\text{race} \triangleq (\text{locs} \cap \text{one}(\mathcal{W}) \cap \text{one}(\text{NA})) \setminus (G.\text{hb}^? \cup G.\text{hb}^{-1})$$

$$G.\text{po}|_{\neq \text{locs}} \triangleq \text{po} \setminus \text{locs}$$

$$G.\text{fr} \triangleq (G.\text{rf}^{-1}; G.\text{wb}) \setminus [G.\mathcal{V}]$$

$$G.\text{hbsc} \triangleq [\text{SC}]; \left((\text{hb} \cap \text{locs}) \cup \text{po}|_{\neq \text{locs}} \cup \left(\text{po}|_{\neq \text{locs}}; \text{hb}; \text{po}|_{\neq \text{locs}} \right) \right); [\text{SC}]$$

$$G.\text{wb} \triangleq ([\mathcal{W}]; ((G.\text{brf}; (G.\text{hb} \cap \text{locs}); G.\text{brf}) \setminus G.\text{brf}); [\mathcal{W}])^+$$

Functions

$$\text{WWrace}(G) \triangleq \exists w, w' \in \mathcal{W}. G.\text{race}(w, w')$$

$$\text{hbW}(G, e) \triangleq \exists w \in \mathcal{W}. (w, e) \in (G.\text{hb} \cap \text{locs})$$

$$\text{AddRF}(G, e, e') \triangleq \langle G.\mathcal{V}, G.\text{po}, G.\text{rf} \cup \{(e, e')\} \rangle$$

A.2 Event Structure Construction Rules

We also repeat the event structure rules in Figure 14.

$$\frac{e \in \mathcal{V} \quad \forall e'' \in \mathcal{V}. e'.\text{id} \neq e''.\text{id} \quad e.\text{code} \xrightarrow{e'.\text{lab}} e'.\text{code} \quad \forall e''. \text{po}(e, e'') \implies e'.\text{lab} \neq e''.\text{lab}}{\langle \mathcal{V}, \text{po}, \text{rf} \rangle \xrightarrow{e'} \langle \mathcal{V} \cup \{e'\}, (\text{po} \cup \{(e, e')\})^+, \text{rf} \rangle} \text{ (BASIC)}$$

$$\frac{\text{WWrace}(G)}{G \rightsquigarrow G'} \text{ (WW-RACE)} \quad \frac{G \xrightarrow{e'} G' \quad e' \notin \mathcal{R}}{G \rightsquigarrow G'} \text{ (NON-READ)}$$

$$\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.\text{rval} = \mathbf{u} \quad \neg G'.\text{hbW}(e')}{G \rightsquigarrow G'} \text{ (R-UNINIT)}$$

$$\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.\text{rval} = \mathbf{u} \quad G'.\text{hbW}(e') \quad \exists w. G'.\text{race}(w, e')}{G \rightsquigarrow G'} \text{ (R-RACE)}$$

$$\frac{G \xrightarrow{e'} G' \quad e' \in \mathcal{R} \quad e'.\text{rval} = w.\text{wval} \quad G'.\text{hbW}(e') \quad \neg G'.\text{race}(w, e') \quad G'' = \text{AddRF}(G', w, e') \quad \text{isCons}(G'')}{G \rightsquigarrow G''} \text{ (R-NORACE)}$$

Figure 14. Event structure reduction steps.

B. Proofs of DRF Theorems

RA-consistency.

The formal definition:

- Let $G \rightsquigarrow_{\text{RA}} G'$ and $G \longrightarrow_{\text{RA}} G'$ be defined exactly as $G \rightsquigarrow G'$ and $G \longrightarrow G'$ but without the rule R-RACE.
- $\llbracket P \rrbracket_{\text{RA}} \triangleq \{G \mid G_{\text{init}}(P) \rightsquigarrow_{\text{RA}}^* G\}$.

Proposition 1. *If $G \in \llbracket P \rrbracket_{\text{RA}}$, then $G.\text{rf} \subseteq G.\text{hb}$.*

Proof. We apply induction on the length of the event structure construction for G .

We consider two cases

- (i) case $G = G_{\text{init}}(P)$. In this case $G.\text{rf} = \emptyset$ and $G.\text{rf} \subseteq G.\text{hb}$ trivially holds.
- (ii) case $G \neq G_{\text{init}}(P)$. In this case let $G_{\text{init}}(P) = G_0 \rightsquigarrow^* G_i \rightsquigarrow^* G_{i+1} \rightsquigarrow \dots \rightsquigarrow G$ be the construction of the event structure G .

Consider the event structure G_i is RA-consistent and $\forall 0 \leq k \leq i G_k \in \llbracket P \rrbracket_{\text{RA}}$. As a result, $G_i.\text{rf} \subseteq G_i.\text{hb}$.

We append an event e in the event structure G_i and construct G_{i+1} .

The possible cases are

- case (1) $e \in [\text{St}]$. Apply (NON-READ) rule. It does not result in any new rf edge. As a result $G_{i+1}.\text{rf} \subseteq G_{i+1}.\text{hb}$.
- case (2) $e \in [\mathcal{R}]$.

- case (2.1) $\neg G_{i+1}.\text{hbW}(e)$. Apply (R-UNINIT). It does not result in any new rf edge. As a result $G_{i+1}.\text{rf} \subseteq G_{i+1}.\text{hb}$.
- case (2.2) $G_{i+1}.\text{hbW}(e)$. Identify the write $w \in (G_i.\mathcal{V} \cap \mathcal{W})$ such that $w.\text{loc} = r.\text{loc}$ and w is wb-maximal in $\in G_{i+1}.\mathcal{V}$ i.e. there is no write w' such that $w' \in (G_{i+1}.\mathcal{V} \cap \mathcal{W}) \wedge w'.\text{loc} = r.\text{loc} \wedge G_i.\text{wb}(w, w')$.

The read event e reads from w by (R-NORACE) rule. Note that, it does not create any fr edge starting from e and hence appending e reading from w is not constrained by any consistency rule.

Hence, $G_{i+1}.\text{rf} = G_i.\text{rf} \uplus \{(w, r)\}$ and $G_{i+1}.\text{hb} = G_i.\text{hb} \uplus \{(w, r)\}$.

As a result, $G_{i+1}.\text{rf} \subseteq G_{i+1}.\text{hb}$. \square

RW-race (RWrace). The formal definition:

$$\begin{aligned} \text{RWrace}(G) &\triangleq \exists w, r \in G.\mathcal{V}. r \in \mathcal{R} \\ &\quad \wedge w \in \mathcal{W} \wedge G.\text{race}(w, r) \end{aligned}$$

Lemma 5. *If $\text{RWrace}(G)$ and $G \rightsquigarrow G'$, then $\text{RWrace}(G')$.*

Proof. Let r be the first racy read in G i.e. there is no r' in the $(\text{po} \cup \text{rf})^+$ -prefix of r which is racy. Such an r exists as G is $(\text{po} \cup \text{rf})$ acyclic by construction.

In this case there exists a write w in G such that (w, r) races and there is no $(\text{po} \cup \text{rf})$ path from w to r .

Considering all rules of event structure construction, we cannot create an event structure G' such that $(w, r) \in (\text{po} \cup \text{rf})^+$ and hence in (w, r) races in G' . As a result, $\text{RWrace}(G')$ holds. \square

We relate read-write races and the reads which return \mathbf{u} in the following lemma.

Lemma 6. $\forall P. \forall G \in \llbracket P \rrbracket_{\text{LLVM}}. \text{RWrace}(G) \implies \exists r \in G.\mathcal{V}. r.\text{rval} = \mathbf{u}$

Proof. Follows from the event structure construction rules that there is a racy read write event pair (w, r) .

$$\exists w, r \in G.\mathcal{V}. G.\text{race}(w, r)$$

The possible cases are:

case (i) the racy read r returns **undef** i.e. $r.\text{rval} = \mathbf{u}$.

case (ii) the read r races with w but reads from some other write i.e. $\neg \text{rf}(w, r)$. In this case by the event construction rules, there is another read r' in conflict with r which is in race with w and returns \mathbf{u}

$$\exists r' \in G.\mathcal{V}. G.\text{cf}(r, r') \wedge r'.\text{rval} = \mathbf{u}$$

\square

Lemma 1. *If an LLVM-consistent event structure G has no read-write races, then G is also RA-consistent.*

Formal statement:

$$\forall P. \forall G \in \llbracket P \rrbracket_{\text{LLVM}}. \neg \text{RWrace}(G) \implies G \in \llbracket P \rrbracket_{\text{RA}}$$

Proof. From $G \in \llbracket P \rrbracket_{\text{LLVM}}$, we know that $G_{\text{init}}(P) = G_0 \rightsquigarrow G_1 \rightsquigarrow \dots \rightsquigarrow G_k = G$ for some $k \geq 0$.

From Lemma 5, we get that $\neg \text{RWrace}(G_i)$, for all $0 \leq i \leq k$. So all the $G_i \rightsquigarrow G_{i+1}$ steps are also $G \rightsquigarrow_{\text{RA}} G'$ steps, and hence $G \in \llbracket P \rrbracket_{\text{RA}}$. \square

Lemma 2. *Given a program P and an LLVM-consistent event structure of P with a read-write race, there exists some RA-consistent event structure of P with a read-write race.*

Formal statement:

$$\begin{aligned} \forall P. \forall G \in \llbracket P \rrbracket_{\text{LLVM}}. \text{RWrace}(G) \\ \implies \exists G' \in \llbracket P \rrbracket_{\text{RA}}. \text{RWrace}(G') \end{aligned}$$

Proof Outline. To prove this lemma, we consider the construction of the racy event structure: there must be a point during its construction where a read-write race was first introduced; i.e., there was a transition $G \rightsquigarrow G'$, where G has no read-write races, but G' has a read-write race. Clearly, the race must involve the newly added event. If the rule applied was not R-RACE, the same rule can be applied according to the RA-semantics. If, however, R-RACE was applied, then instead we add a read event (possibly returning a different value) that reads from some visible happens-before write or (in case of an SC read, of the last global write if it is visible). In the constructed event structure, the read races with the write operation which was its earlier justifier.

Proof. By induction on the size of G

case (i) No read returns \mathbf{u} i.e. $\forall r \in G.V. r.rval \neq \mathbf{u}$.

In this case a read r reads from a write w such that $G.hb(w, r) \wedge G.rf(w, r)$ holds.

Take $G' = G$; G' is RA-consistent because $G'.rf \subseteq G'.hb$. If $RWrace(G)$ holds due to r then $RWrace(G')$ also holds.

case (ii) otherwise i.e. $\exists r \in G.V. r.rval = \mathbf{u}$.

We identify a read r such that it returns \mathbf{u} and its $(G.po \cup G.rf)^+$ execution prefix has no read which returns \mathbf{u} .

$$r.rval = \mathbf{u} \wedge \nexists r'. (r', r) \in (G.po \cup G.rf)^+ \wedge r'.rval = \mathbf{u}$$

(1) r is uninitialized. Take $G' = G$; G' is RA-consistent because $G'.rf \subseteq G'.hb$. If $RWrace(G)$ holds due to r then $RWrace(G')$ also holds.

(2) r is initialized and $r.rval = \mathbf{u}$.

Which means $\exists w \in G.V. G.race(w, r)$.

We discard r and all the events e such that $(r, e) \in (G.po \cup G.rf)^+$ along with the $G.po$ and $G.rf$ edges and obtain G'' . More formally,

let $V_1 = \{e \mid (e = r) \vee (r, e) \in (G.po \cup G.rf)^+\}$ and

$$G''.V = G.V \setminus V_1,$$

$$G''.po = G.po \setminus \{(_, e) \mid e \in G.V_1\},$$

$$G''.rf = G.rf \setminus \{(_, e) \mid e \in G.V_1\}.$$

Then we append back r' in place of r which reads from a immediate happens before visible write and races with w . We observe that $w, r' \in G'.V$ since $acyclic(G'.po \cup G'.rf)^+$. As a result we can apply the induction hypothesis and construct G' .

Therefore, G' is RA-consistent and $RWrace(G')$ holds. \square

Read-Write-DRF-RA ($RWDRF_{RA}$). The formal definition:

$$RWDRF_{RA}(P) \stackrel{\text{def}}{=} \nexists G. G \in \llbracket P \rrbracket_{RA} \wedge RWrace(G)$$

Theorem 1 (DRF-RA). *If under RA consistency, a program has no read-write races, then its LLVM-consistent behaviors coincide with its RA-consistent ones.*

The formal description of the DRF-RA theorem:

$$\forall P. RWDRF_{RA}(P) \implies \llbracket P \rrbracket_{LLVM} = \llbracket P \rrbracket_{RA}$$

Proof. Assume $RWDRF_{RA}(P)$ holds on a program P . In that case we show,

$$\begin{aligned} \llbracket P \rrbracket_{LLVM} &\subseteq \llbracket P \rrbracket_{RA} \\ \wedge \llbracket P \rrbracket_{RA} &\subseteq \llbracket P \rrbracket_{LLVM} \end{aligned}$$

First, we show that $\llbracket P \rrbracket_{LLVM} \subseteq \llbracket P \rrbracket_{RA}$.

We show that for all LLVM-consistent event structure G i.e. $\forall G \in \llbracket P \rrbracket_{LLVM}$, either G has no read-write race or there is another RA-consistent event structure G' which has read-write race.

$$\begin{aligned} \forall G \in \llbracket P \rrbracket_{LLVM}. \neg RWrace(G) \\ \vee (\exists G' \in \llbracket P \rrbracket_{RA}. RWrace(G')) \end{aligned}$$

Now we consider the two cases:

case (i) $\neg RWrace(G)$.

From Lemma 1, G is also RA-consistent; $G \in \llbracket P \rrbracket_{RA}$.

case (ii) $RWrace(G)$.

From Lemma 2, there exists an RA-consistent read-write race event structure G' i.e.

$$\exists G' \in \llbracket P \rrbracket_{RA}. RWrace(G')$$

But it contradicts the assumption that $RWDRF_{RA}(P)$. Hence $\llbracket P \rrbracket_{LLVM} \subseteq \llbracket P \rrbracket_{RA}$ holds.

The other direction; $\llbracket P \rrbracket_{RA} \subseteq \llbracket P \rrbracket_{LLVM}$ trivially holds. \square

Release-Acquire-Race (RArace). The formal definition:

$$\begin{aligned} RArace(G) \stackrel{\text{def}}{=} \exists a, b \in G.V. G.conc(a, b) \wedge a.loc = b.loc \\ \wedge \neg(a.ord = SC \wedge b.ord = SC) \wedge (a \in \mathcal{W} \vee b \in \mathcal{W}) \end{aligned}$$

Lemma 3. *If an RA-consistent event structure G is RA-race-free, then G is OSC-consistent.*

The formal description:

$$\begin{aligned} \forall P. \forall G \in \llbracket P \rrbracket_{LLVM}. \\ \neg RArace(G) \implies G \in \llbracket P \rrbracket_{SC} \end{aligned}$$

Proof Outline Essentially, the lack of RA races means that any two same-location writes are either both SC-writes or they are ordered either by happens-before and therefore also by writes-before.

Proof. We prove this by contradiction.

Assume G is an LLVM-consistent event structure without RA-race; but G is not SC-consistent.

$$\exists G. G \in \llbracket P \rrbracket_{LLVM} \wedge \neg RArace(G) \wedge G \notin \llbracket P \rrbracket_{SC}$$

We identify a read r on location X in G which has not read from the latest write on X during the construction of G ; and there exists no such r' in the $(G.po \cup G.rf)^+$ prefix of r .

$$\exists X, r, w, w'. G.rf(w, r) \wedge iswrite_X(w') \wedge \neg G.hb(r, w')$$

Following cases are possible:

case (i) $w'.ord = SC \wedge r.ord = SC$.

In this case r does not violate SC-semantics. But the assumption $G \notin \llbracket P \rrbracket_{SC}$ is not satisfied.

case (ii) $G.\text{RArace}(w', r)$.

It violates the antecedent $\neg\text{RArace}(G)$.

Hence the lemma holds. \square

Lemma 4. *For each RA-racy LLVM-consistent event structure there is a RA-racy OSC-consistent event structure.*

The formal description:

$$\begin{aligned} \forall P. \forall G \in \llbracket P \rrbracket_{\text{LLVM}}. \text{RArace}(G) \\ \implies \exists G' \in \llbracket P \rrbracket_{\text{SC}} \wedge \text{RArace}(G') \end{aligned}$$

Proof Outline. We identify the prefix in the event structure which contains no RA-race and the first RA-race. Next, we check from which accesses the events of the first race were created. We create the events from those two accesses and append them in the event structure. In this case if a generated event is a read then it is justified by some happens-before write from the non-RA racy event structure prefix. At the same time these events have RA-race with the generated events from the other access. As a result the constructed event structure is SC-consistent and at the same time RA-racy.

Proof. In the LLVM-consistent event structure G , we identify the first race $G.\text{RArace}(w, r)$ on location X ; such that there exists no w_1 and r_1 in $(G.\text{po} \cup G.\text{rf})^+$ prefix of w and r respectively such that $G.\text{RArace}(w_1, _)$ or $G.\text{RArace}(_, r_1)$ holds.

Let e_w and e_r be the events such that $G.\vec{\text{po}}(e_w, w)$ and $G.\vec{\text{po}}(e_r, r)$.

Identify G_1 excluding w, r and all the following events in $(\text{po} \cup \text{rf})^+$.

Let $V_1 = \{w, r\} \cup \{e \mid (w, e) \in (G.\text{po} \cup G.\text{rf})^+ \vee (r, e) \in (G.\text{po} \cup G.\text{rf})^+\}$ in

$$G_1.V = G.V \setminus V_1$$

$$G_1.\text{po} = G.\text{po} \setminus \{(_, e) \mid e \in V_1\}$$

$$G_1.\text{rf} = G.\text{rf} \setminus \{(_, e) \mid e \in V_1\}$$

Following cases are possible.

case (i) $\text{iswrite}(w) \wedge \text{isread}(r)$.

Let w and r is generated from the commands C_w and C_r respectively.

Now append $\text{Ev}(C_r)$ in G_1 after e_r to create G_2 i.e. $G_2.\text{po}(e_r, \text{Ev}(C_r))$ where $\text{Ev}(C_r)$ is justified by some existing write w' i.e. $w' \in G_1.V \wedge G_2.\text{rf}(w', e_r)$.

If no such w' exists then $\text{Ev}(C_r)$ is uninitialized and returns **u**.

Next, we append w in G_2 to create G_3 i.e. $G_3.\text{po}(e_w, w)$.

If w is an update event then it is justified by same $w' \in G.V$.

Note that $G_3 \in \llbracket P \rrbracket_{\text{SC}} \wedge \text{RArace}(G_3)$.

We construct G' from G_3 by continuing with the event structure construction rules.

Hence $G' \in \llbracket P \rrbracket_{\text{SC}} \wedge \text{RArace}(G')$.

case (ii) $\text{case } \text{iswrite}(w) \wedge \text{iswrite}(r) \wedge \neg(\text{isrmw}(w) \vee \text{isrmw}(r))$.

we append w and r in G_1 to create G_2 i.e. $\text{po}(\sigma_w, w) \wedge \text{po}(\sigma_r, r)$.

Note that $G_2 \in \llbracket P \rrbracket_{\text{SC}} \wedge \text{RArace}(G_2)$.

We construct G' from G_2 by continuing with the event structure construction rules.

Hence $G' \in \llbracket P \rrbracket_{\text{SC}} \wedge \text{RArace}(G')$. \square

Finally, we show that the LLVM memory model follows the DRF_{SC} property.

DRF-SC (DRF_{SC}). The formal definition:

$$\text{DRF}_{\text{SC}}(P) \stackrel{\text{def}}{=} \nexists G. G \in \llbracket P \rrbracket_{\text{SC}} \wedge \text{RArace}(G)$$

Theorem 2 (DRF-OSC). *If a program is RA-race-free under OSC, then OSC and LLVM consistency coincide.*

The formal description of the DRF-SC theorem:

$$\begin{aligned} \forall P. \text{DRF}_{\text{SC}}(P) \implies \\ \llbracket P \rrbracket_{\text{LLVM}} = \llbracket P \rrbracket_{\text{RA}} = \llbracket P \rrbracket_{\text{SC}} \end{aligned}$$

Proof. Two cases.

case(i) $\forall P. \text{DRF}_{\text{SC}}(P) \implies \llbracket P \rrbracket_{\text{LLVM}} = \llbracket P \rrbracket_{\text{RA}}$

Note that $\text{DRF}_{\text{SC}}(P) \implies \text{RWDRF}_{\text{RA}}(P)$.

In that case $\llbracket P \rrbracket_{\text{LLVM}} = \llbracket P \rrbracket_{\text{RA}}$ from Theorem 1.

case(ii) $\forall P. \text{DRF}_{\text{SC}}(P) \implies \llbracket P \rrbracket_{\text{LLVM}} = \llbracket P \rrbracket_{\text{SC}}$

Assume $\text{DRF}_{\text{SC}}(P)$ holds on a program P .

First, we show that $\llbracket P \rrbracket_{\text{LLVM}} \subseteq \llbracket P \rrbracket_{\text{SC}}$.

Let $\forall G \in \llbracket P \rrbracket_{\text{LLVM}}$.

case (1) $\neg\text{RArace}(G)$.

From Lemma 3, $G \in \llbracket P \rrbracket_{\text{SC}}$.

case (2) $\text{RArace}(G)$.

From Lemma 4, there exists SC-consistent RA-consistent event structure G' .

$$\exists G' \in \llbracket P \rrbracket_{\text{SC}}. \text{RArace}(G')$$

It contradicts the assumption that $\text{RWDRF}_{\text{SC}}(P)$.

The other direction; $\llbracket P \rrbracket_{\text{SC}} \subseteq \llbracket P \rrbracket_{\text{LLVM}}$ trivially holds.

As a result the DRF-SC theorem holds. \square

C. Reordering Optimizations

Given a relation R , we write \vec{R} for the immediate subrelation of R ; that is, $\vec{R}(a, b) \triangleq R(a, b) \wedge \nexists c. R(a, c) \wedge R(c, b)$.

Let C_a and C_b be the two independent adjacent accesses to be reordered. Let A (resp. B) be the set of events generated from C_a (resp. C_b).

Expand & Collapse Since C_a and C_b are adjacent, all the immediate successors of events in A belong to the set B . There may, however, be events in A that have multiple successors (in B). To simplify the definition of reordering at the level of event structures, we will first expand the source event structure so that each event in A has a unique successor.

Let an event a from the set A have po-successor events b and b' in B . We define $\text{expand}(G, a, b')$ to extend G with a new event a' with same label as a and make b' be the po-successors of the new event a' instead of a .

$$\begin{aligned} \text{expand}(G, a, b') &\triangleq G' \text{ where} \\ G'.V &= G.V \uplus \{a'\} \\ G'.\text{po} &= \left(\begin{array}{l} G.\text{po} \setminus \{(a, c) \mid (b', c) \in G.\text{po}^*\} \\ \cup \{(a', b')\} \cup \{(c, a') \mid (c, a) \in G.\text{po}\} \end{array} \right)^+ \\ G'.\text{rf} &= G.\text{rf} \cup \{(x, a') \mid (x, a) \in G.\text{rf}\} \\ a'.\text{lab} &= a.\text{lab} \end{aligned}$$

We repeatedly apply $\text{expand}(G, a, b')$ to *expand* the set A to A' until each event in A' has a unique immediate successor (in B).

We also define the inverse transformation that collapses to events a and a' that have the same label (lab) and the same po-predecessor.

$$\begin{aligned} \text{collapse}(G, a, a') &\triangleq G' \text{ where} \\ G'.V &= G.V \setminus \{a'\} \\ G'.\text{po} &= \left(\begin{array}{l} G.\text{po} \cup \{(a, x) \mid (a', x) \in G.\text{po}\} \\ \setminus \{(a', x), (x, a') \mid (a', x) \in G.\text{po}\} \end{array} \right) \\ G'.\text{rf} &= G.\text{rf} \setminus \{(x, a') \mid (x, a') \in G.\text{rf}\} \end{aligned}$$

Definitions We define the *successor* and *predecessor* and *adjacent* (Adj) for the set of events A and B with respect to a relation R .

Definition 1. $\text{succ}(R, A, B) \stackrel{\text{def}}{=} \forall a \in A, c \in V, \exists b \in B. (R(a, c) \wedge (b \neq c)) \implies R(b, c)$

Definition 2. $\text{pred}(R, A, B) \stackrel{\text{def}}{=} \forall b \in B, c \in V, \exists a \in A. (R(c, b) \wedge a \neq c) \implies R(c, a)$

Definition 3. $\text{Adj}(R, A, B) \stackrel{\text{def}}{=} \forall a \in A, b \in B, c \in V. a.\text{tid} = b.\text{tid}$

$$\begin{aligned} \wedge \text{succ}(R, A, B) \wedge R(b, c) &\implies R(a, c) \\ \wedge \text{pred}(R, A, B) \wedge R(c, a) &\implies R(c, b) \end{aligned}$$

Definition 4. $\text{Reord}(A, B) \stackrel{\text{def}}{=} (A \cap B = \emptyset) \wedge \text{Adj}(\text{po}, A, B) \wedge (\nexists a, b. a \in A \wedge b \in B. \wedge a.\text{loc} = b.\text{loc}) \wedge \neg(A.\text{ord} = \text{SC} \wedge B.\text{ord} = \text{SC})$

$$\begin{aligned} \wedge (\forall a \in A. \neg(\text{isread}(a) \wedge a \in \text{Acq})) \\ \wedge (\forall b \in B. \neg(\text{iswrite}(b) \wedge b \in \text{Rel})) \end{aligned}$$

Proofs We now proceed to the safe reordering proofs on the expanded event structure.

We first show that a reordering following Reord results in more synchronization in the event structure.

Lemma 7. $(\text{pred}(R, B, A) \wedge \text{succ}(R, B, A) \wedge (A \cap B = \emptyset)) \implies (R \setminus (B \times A) \cup (A \times B))^+ \subseteq (R^+ \setminus (B \times A) \cup (A \times B))$

Proof. We prove by induction on transitive closure of R . Base Case.

$$R \setminus (B \times A) \cup (A \times B) \subseteq (R^+ \setminus (B \times A) \cup (A \times B))$$

The base case holds by monotonicity.

The induction step.

$$(R \setminus (B \times A) \cup (A \times B)) \circ (R^+ \setminus (B \times A) \cup (A \times B)) \subseteq (R^+ \setminus (B \times A) \cup (A \times B))$$

\implies It is sufficient to show

- (i) $(R \setminus (B \times A)) \circ (R^+ \setminus (B \times A)) \subseteq R^+ \setminus (B \times A) \cup (A \times B)$
- (ii) $(R \setminus (B \times A)) \circ (A \times B) \subseteq R^+ \setminus (B \times A) \cup (A \times B)$
- (iii) $(A \times B) \circ (R^+ \setminus (B \times A)) \subseteq R^+ \setminus (B \times A) \cup (A \times B)$
- (iv) $(A \times B) \circ (A \times B) \subseteq R^+ \setminus (B \times A) \cup (A \times B)$

case (i) It is sufficient to show

$$(R \setminus (B \times A)) \circ (R^+ \setminus (B \times A)) \subseteq (R^+ \setminus (B \times A))$$

It is sufficient to show

$$\begin{aligned} (R \setminus (B \times A)) \circ (R^+ \setminus (B \times A)) &\subseteq R^+ \\ \wedge (B \times A) &\notin (R \setminus (B \times A)) \circ (R^+ \setminus (B \times A)) \end{aligned}$$

$$(1) (R \setminus (B \times A)) \circ (R^+ \setminus (B \times A)) \subseteq R^+$$

holds by monotonicity.

(2) We show by contradiction that

$$(B \times A) \notin (R \setminus (B \times A)) \circ (R^+ \setminus (B \times A))$$

Assume $\exists c \in V. (b, c) \in (R \setminus (B \times A)) \wedge (c, a) \in (R^+ \setminus (B \times A))$ where $b \in B \wedge a \in A$.

From the definition of $\text{succ}(R, B, A)$ we know $c \in A$.

Hence $(b, c) \notin (R \setminus (B \times A))$ which is a contradiction.

As a result $(B \times A) \notin (R \setminus (B \times A)) \circ (R^+ \setminus (B \times A))$.

case (ii) $(R \setminus (B \times A)) \circ ((A \times B)) = \emptyset$

since $\nexists b \in B \wedge a \in A. (b, a) \in (R \setminus (B \times A))$ from the definition of $\text{succ}(R, B, A)$.

case (iii) $(A \times B) \circ (R^+ \setminus (B \times A)) = \emptyset$ since $\nexists c, a. c \notin B \wedge a \in A$ from the definition of $\text{pred}(R, B, A)$.

case (iv) $(A \times B) \circ (A \times B) = \emptyset$ since $A \cap B = \emptyset$. \square

Let R^S and R^T be the set of R relations in the source and the target event structures G^S and G^T respectively. Also e^S and e^T refers to an event e in the the event structure before and after the reordering respectively.

Next, we choose the hb relation and show that a safe reordering of adjacent event sets A and B would result in more hb relations in the target compared to the source event structure.

Lemma 8. $\text{pred}(\text{hb}^T, B, A) \wedge \text{succ}(\text{hb}^T, B, A) \implies \text{hb}^S \subseteq (\text{hb}^T \setminus (B \times A) \cup (A \times B))$

Proof. From Lemma 7 $\text{po}^S = \text{po}^T \setminus (B \times A) \cup (A \times B)$. Also we choose $R = (\text{po}^T \cup R')$ where R' is some other relation independent of po^T . Hence

$$\begin{aligned} & (R \setminus (B \times A) \cup (A \times B))^+ \subseteq (R^+ \setminus (B \times A) \cup (A \times B)) \\ \implies & ((\text{po}^T \cup R') \setminus (B \times A) \cup (A \times B))^+ \subseteq ((\text{po}^T \cup R')^+ \setminus (B \times A) \cup (A \times B)) \\ \implies & ((\text{po}^T \setminus (B \times A) \cup (A \times B)) \cup R')^+ \subseteq ((\text{po}^T \cup R')^+ \setminus (B \times A) \cup (A \times B)) \\ \implies & (\text{po}^S \cup R')^+ \subseteq ((\text{po}^T \cup R')^+ \setminus (B \times A) \cup (A \times B)) \\ \implies & (\overrightarrow{\text{po}}^S \cup R')^+ \subseteq ((\overrightarrow{\text{po}}^T \cup R')^+ \setminus (B \times A) \cup (A \times B)) \\ & \text{since } (\text{po}^S \cup R')^+ = (\overrightarrow{\text{po}}^S \cup R')^+ \text{ and } (\text{po}^T \cup R')^+ = (\overrightarrow{\text{po}}^T \cup R')^+; \text{ substituting } R' = \text{sw} \text{ we get} \\ & (\overrightarrow{\text{po}}^S \cup \text{sw})^+ \subseteq ((\text{po}^T \cup \text{sw})^+ \setminus (B \times A) \cup (A \times B)) \\ \implies & \text{hb}^S \subseteq (\text{hb}^T \setminus (B \times A) \cup (A \times B)) \\ & \text{since } \text{hb}^S = (\overrightarrow{\text{po}}^S \cup \text{sw})^+ \text{ and } \text{hb}^T = (\overrightarrow{\text{po}}^T \cup \text{sw})^+. \quad \square \end{aligned}$$

In the event structure any two events on the same location are related by *happens-before* or are *concurrent* or in *conflict* with each other. The *conflict* relations do not change before and after reordering. As a result, the union of *happens-before* and *concurrent* relations on a particular location remains same before and after reordering.

Lemma 9. $(\text{hb}^S \subseteq \text{hb}^T \setminus (B \times A) \cup (A \times B)) \implies \text{hb}_\ell^S \cup \text{conc}_\ell^S = \text{hb}_\ell^T \cup \text{conc}_\ell^T$

Proof. In the event structure two events on the same location ℓ either *conflict* or are *concurrent* or one *happens-before* the other. So $\text{conc}_\ell^S \cup \text{cf}_\ell^S \cup \text{hb}_\ell^S = \text{conc}_\ell^T \cup \text{cf}_\ell^T \cup \text{hb}_\ell^T$.

We also know that reordering does not introduce any conflict relation i.e. $\text{cf}_\ell^S = \text{cf}_\ell^T$.
 $\implies \text{conc}_\ell^S \cup \text{hb}_\ell^S = \text{conc}_\ell^T \cup \text{hb}_\ell^T \quad \square$

A *safe* reordering results in more *happens-before* relations and reduces *concurrent* event pairs on the same location in the target event structure compared to the source event structure.

Lemma 10. $(\text{hb}^S \subseteq \text{hb}^T \setminus (B \times A) \cup (A \times B)) \implies (\text{hb}_\ell^S \subseteq \text{hb}_\ell^T) \wedge (\text{conc}_\ell^T \subseteq \text{conc}_\ell^S)$

Proof. from Lemma 8 we know $\text{hb}^S \subseteq \text{hb}^T \setminus (B \times A) \cup (A \times B)$ where $\forall a \in A, b \in B. a.\text{loc} \neq b.\text{loc}$.
As a result $\text{hb}_\ell^S \subseteq \text{hb}_\ell^T$.
From Lemma 9 $(\text{hb}_\ell^S \cup \text{conc}_\ell^S = \text{hb}_\ell^T \cup \text{conc}_\ell^T)$.
As a result $\text{conc}_\ell^T \subseteq \text{conc}_\ell^S$. \square

In a *safe* reordering the same location *happens-before* relations of the source event structure remain in the target event structure and some *concurrent* events on the same

location in the source event structure turns into *happens-before* relations in the target.

Lemma 11. $(\text{hb}^S \subseteq \text{hb}^T \setminus (B \times A) \cup (A \times B)) \implies \text{hb}_\ell^T \subseteq \text{hb}_\ell^S \cup \text{conc}_\ell^S$

Proof. From Lemma 9 $(\text{hb}_\ell^T \cup \text{conc}_\ell^T = \text{hb}_\ell^S \cup \text{conc}_\ell^S)$.
From Lemma 10 $(\text{conc}_\ell^T \subseteq \text{conc}_\ell^S)$.
As a result, $\text{hb}_\ell^T \subseteq \text{hb}_\ell^S \cup \text{conc}_\ell^S$. \square

A reordering does not change the direction of *happens-before* between a pair of events on the same location.

Lemma 12. $(\text{hb}^S \subseteq \text{hb}^T \setminus (B \times A) \cup (A \times B)) \implies \nexists a, b. \text{hb}_\ell^S(a, b) \wedge \text{hb}_\ell^T(b, a)$.

Proof. We prove this by contradiction.
Assume such $\exists a, b. \text{hb}_\ell^S(a, b) \wedge \text{hb}_\ell^T(b, a)$.
In that case $a \in A \wedge b \in B$.
However, we already know that $a \in A \wedge b \in B \implies a.\text{loc} \neq b.\text{loc}$.
Hence the contradiction. \square

We show that in a *safe* reordering a read operation reads subset of values in the target compared to the source. Let r^S and r^T represents an event r before and after the reordering.

Lemma 13. $(\text{hb}^S \subseteq \text{hb}^T \setminus (B \times A) \cup (A \times B)) \implies \forall r^T \exists r^S. r^T.\text{rval} \preceq r^S.\text{rval}$

Proof. A read can be initialized or uninitialized.

case (i) r^T is initialized.

We know $\text{conc}_\ell^T \subseteq \text{conc}_\ell^S$ (Lemma 10) and $\text{hb}_\ell^T \subseteq \text{hb}_\ell^S \cup \text{conc}_\ell^S$ (Lemma 11).

Let $\text{rf}^T(w^T, r^T)$ and $\text{rf}^S(w^S, r^S)$ holds. Considering both the source and target event structure following cases are possible:

- $\text{hb}^T(w^T, r^T) \wedge \text{hb}^S(w^S, r^S)$.
In this case $r^T.\text{rval} = r^S.\text{rval}$.
- $\text{hb}^T(w^T, r^T) \wedge \text{conc}^S(w^S, r^S)$.
In this case $r^S.\text{rval} = \mathbf{u}$ and hence $r^T.\text{rval} \preceq r^S.\text{rval}$ holds.
- $\text{conc}^T(w^T, r^T) \wedge \text{conc}^S(w^S, r^S)$.
In this case $r^T.\text{rval} = r^S.\text{rval} = \mathbf{u}$ and hence $r^T.\text{rval} \preceq r^S.\text{rval}$ holds.

Note that r^T does not read from any write w^T in target such that

(1) $\text{conc}^T(w^T, r^T) \wedge \text{hb}^S(r^S, w^S)$
since $\text{conc}_\ell^T \subseteq \text{conc}_\ell^S$ (Lemma 10) or

(2) $\text{hb}^T(w^T, r^T) \wedge \text{hb}^S(r^S, w^S)$
since $\text{hb}_\ell^T \subseteq \text{hb}_\ell^S \cup \text{conc}_\ell^S$ (Lemma 12).




Figure 17. Event structure before Rna;Wrel reordering.




Figure 18. Event structure after Rna;Wrel reordering.




Figure 19. Event structure before Wna;Wrel reordering.




Figure 20. Event structure after Wna;Wrel reordering.




Figure 15. Event structure before Wna;Wrel reordering.




Figure 16. Event structure after Wna;Wrel reordering.

case (ii) r^T is uninitialized. In this case r^S is also uninitialized and hence $(r^T.rval \preceq r^S.rval)$. \square

The Lemma 13 signifies that for each target execution branch in the event structure there exists some execution branch in the source. We prove, the observable target behaviors is the subset of the observable source behaviors.

Lemma 14. $\forall r^T, r^S. r^S.rval \preceq r^S.rval \implies \text{Behavior}_{\text{LLVM}}(P^T) \subseteq \text{Behavior}_{\text{LLVM}}(P^S)$

Proof. The antecedent signifies that in each thread for each target execution branch there exists an execution branch in the source since each read event initiates an individual execution branch. In turn for each target write event there exists a source write event generated from the same instruction;

$\forall w^T \exists w^S. w^T.\text{lab} = w^S.\text{lab}$

$\implies \forall \text{Behavior}(G^T) \exists \text{Behavior}(G^S)$

$\implies \text{Behavior}_{\text{LLVM}}(P^T) \subseteq \text{Behavior}_{\text{LLVM}}(P^S)$ \square

D. Examples of Unsafe Reorderings

$St_k; St_{\sqsupseteq REL} \rightsquigarrow St_{\sqsupseteq REL}; St_k$ Consider the reordering in the following example:

$$\begin{array}{l} X_k := 4; \\ Y_{\sqsupseteq REL} := 1; \end{array} \rightsquigarrow \begin{array}{l} Y_{\sqsupseteq REL} := 1; \\ X_k := 4; \end{array} \left[\begin{array}{l} - \parallel \text{if}(Y_{\sqsupseteq ACQ}) \\ X_{NA} = 1; \end{array} \right]$$

The event structure before the transformation is Figure 15: The event structure after the transformation is Figure 16:

$Ld_k; St_{\sqsupseteq REL} \rightsquigarrow St_{\sqsupseteq REL}; Ld_k$ Consider the reordering in the following example:

$$\begin{array}{l} t := X_k; \\ Y_{\sqsupseteq REL} := 1; \end{array} \rightsquigarrow \begin{array}{l} Y_{\sqsupseteq REL} := 1; \\ t := X_k; \end{array} \left[\begin{array}{l} - \parallel \text{if}(Y_{\sqsupseteq ACQ}) \\ X_{NA} = 1; \end{array} \right]$$

The event structure before the transformation is Figure 17: The event structure after the transformation is Figure 18 where the event a is created which returns **undef**:

$St_k; U_{\sqsupseteq REL} \rightsquigarrow U_{\sqsupseteq REL}; St_k$ Consider the reordering in the following example:

$$\begin{array}{l} X_k := 4; \\ \text{CAS}(Y, 0, 1); \end{array} \rightsquigarrow \begin{array}{l} \text{CAS}(Y, 0, 1); \\ X_k := 4; \end{array} \left[\begin{array}{l} - \parallel \text{if}(Y_{\sqsupseteq ACQ}) \\ X_{NA} = 1; \end{array} \right]$$

The event structure before the transformation is Figure 19: The event structure after the transformation is Figure 20:




Figure 24. Event structure after Racq;Wk reordering.




Figure 23. Event structure before Racq;Wk reordering.




Figure 25. Event structure before Racq;Rk reordering.




Figure 26. Event structure after Racq;Rk reordering.




Figure 21. Event structure before Rna;Wrel reordering.




Figure 22. Event structure after Rna;Wrel reordering.

$Ld_k; U_{\square_{ACQ_REL}} \rightsquigarrow U_{\square_{ACQ_REL}}; Ld_k$ Consider the reordering in the following example:

$$\begin{array}{l}
 t := X_k; \\
 CAS(Y, 0, 1);
 \end{array}
 \rightsquigarrow
 \begin{array}{l}
 CAS(Y, 0, 1); \\
 t := X_k;
 \end{array}
 \left[\begin{array}{l}
 \text{Context:} \\
 - \parallel \text{if}(Y_{\square_{ACQ}}) \\
 \quad X_{NA} := 1;
 \end{array} \right]$$

The event structure before the transformation is Figure 21: The event structure after the transformation is Figure 22 where the event a is created which returns **undef**.

$Ld_{\square_{ACQ}}; St_k \rightsquigarrow St_k; Ld_{\square_{ACQ}}$ Consider the reordering in the following example:

$$\begin{array}{l}
 \text{while}(!Y_{\square_{ACQ}}); \\
 X_k := 4;
 \end{array}
 \rightsquigarrow
 \begin{array}{l}
 X_k := 4; \\
 \text{while}(!Y_{\square_{ACQ}});
 \end{array}
 \left[\begin{array}{l}
 \text{Context:} \\
 - \parallel X_{NA} := 1; \\
 \quad Y_{\square_{REL}} := 1;
 \end{array} \right]$$

The event structure before the transformation is Figure 23: The event structure after the transformation is Figure 24 where the event a is created which returns **undef**:

$Ld_{\square_{ACQ}}; Ld_k \rightsquigarrow Ld_k; Ld_{\square_{ACQ}}$ Consider the reordering in the following example:

$$\begin{array}{l}
 \text{while}(!Y_{\square_{ACQ}}); \\
 t := X_k;
 \end{array}
 \rightsquigarrow
 \begin{array}{l}
 t := X_k; \\
 \text{while}(!Y_{\square_{ACQ}});
 \end{array}
 \left[\begin{array}{l}
 \text{Context:} \\
 - \parallel X_{NA} := 1; \\
 \quad Y_{\square_{REL}} := 1;
 \end{array} \right]$$

The event structure before the transformation is Figure 25: The event structure after the transformation is Figure 26 where the event a is created which returns **undef**.

E. Elimination Optimizations

The elimination transformation (elim) generates a target program P^T by eliminating some redundant shared memory access from the source program P^S .

$$P^S \xrightarrow{\text{elim}} P^T$$

For the safety proofs we show that the elimination optimizations preserves the program behavior. To prove this, it suffices to show :

- For each consistent event structure G^T of the program P^T , there exists a consistent event structure G^S of the program P^S .

$$\begin{aligned} \forall G^T \in \text{exec}(\llbracket P^T \rrbracket_{\text{LLVM}}) &\implies \\ \exists G^S \in \text{exec}(\llbracket P^S \rrbracket_{\text{LLVM}}) & \\ \wedge \text{Behavior}(G^T) \subseteq \text{Behavior}(G^S) & \end{aligned}$$

- G^S is constructible.

In order to do so, given a consistent target event structure G^T we construct a source event structure G^S following the consistency rules.

We first show the safety of adjacent elimination transformations and then we will consider the elimination of non-adjacent events.

In all these cases we consider a target event structure G and then we construct consistent event structure G' from the source following the construction steps of G except certain steps. In the construction, given an event e in G , the corresponding event in G' is $\mathcal{M}(e)$. If the labels of the e and $\mathcal{M}(e)$ are same then the events are equivalent and we use e instead of $\mathcal{M}(e)$ in G' .

In these proofs let $\text{wbsc} \triangleq [\text{SC}]; \text{wb}; [\text{SC}]$ and $\text{frsc} \triangleq [\text{SC}]; \text{fr}; [\text{SC}]$.

Adjacent Overwritten Write (AOW)

$$\text{store}_{o'}(\ell, v'); \text{store}_o(\ell, v) \rightsquigarrow \text{store}_o(\ell, v) \text{ when } o' \sqsubseteq o$$

Proof. Let $w = \text{St}_o(\ell, v)$ be the event in the target event structure G .

We follow the construction steps of target event structure G and construct the source event structure G' .

Let w is appended in the i^{th} step and event e is appended in $(i-1)^{\text{th}}$ step to G .

The construction steps of G' is of the form

$$G'_{\text{init}} \rightsquigarrow_e^* G'_{i-1} \rightsquigarrow_w G'_i \rightsquigarrow \dots$$

We follow the $(i-1)^{\text{th}}$ step of G and construct G' by applying same construction rule and hence appending the equivalent event to G' .

$$G'_{\text{init}} \rightsquigarrow_e^* G'_{i-1}$$

Let $w' = \text{St}_{o'}(\ell, v')$ be an event generated from the store $o'(\ell, v')$ of the source program.

We append w' to G' in the i^{th} step.

$$G'_{\text{init}} \rightsquigarrow_e^* G'_{i-1} \rightsquigarrow_{w'} G'_i$$

As w' is a store, G'_i is consistent.

We continue to follow the i^{th} and onwards steps of G and append the equivalent events to G' .

The overall construction of G' following the G looks as follows.

$$G'_{\text{init}} \rightsquigarrow_e^* G'_{i-1} \rightsquigarrow_w G'_i \dots$$

$$G'_{\text{init}} \rightsquigarrow_e^* G'_{i-1} \rightsquigarrow_{w'} G'_i \rightsquigarrow_w G'_{i+1} \dots$$

In the construction of G' preserves consistency.

When we follow a step in G where (NON-READ) or (R-UNINIT) is applied then we apply same rule in G' and append equivalent event.

If we apply (R-RACE) or (R-NORACE) in G then we can always apply same rule in G' as the corresponding writes are already been appended in G' similar to G .

Also, considering the derived relations in G and G'

- $G'.\text{wb}(e, w') \implies G'.\text{wb}(e, w) \wedge G.\text{wb}(e, w)$
and $G'.\text{wb}(w', e) \implies G'.\text{wb}(w, e) \wedge G.\text{wb}(w, e) \vee w = e$.
- $G'.\text{cf}(e, w') \implies G'.\text{cf}(e, w) \wedge G.\text{cf}(e, w)$
- $G'.\text{hb}(e, w') \implies G'.\text{hb}(e, w) \wedge G.\text{hb}(e, w)$ and
 $G'.\text{hb}(w', e) \implies G'.\text{hb}(w, e) \wedge G.\text{hb}(w, e) \vee e = w$.
- Also if $o' = o = \text{SC}$ then
 - $G'.\text{wbsc}(e, w') \implies G'.\text{wbsc}(e, w) \wedge G.\text{wbsc}(e, w)$
 - $G'.\text{wbsc}(w', e) \implies G'.\text{wbsc}(w, e) \wedge G.\text{wbsc}(w, e) \vee w = e$
 - $G'.\text{hbsc}(e, w') \implies G'.\text{hbsc}(e, w) \wedge G.\text{hbsc}(e, w)$
 - $G'.\text{hbsc}(w', e) \implies G'.\text{hbsc}(w, e) \wedge G.\text{hbsc}(w, e) \vee w = e$.
 - $G'.\text{frsc}(e, w') \implies G'.\text{frsc}(e, w) \wedge G.\text{frsc}(e, w)$.

Thus the introduction of w' does not introduce any new cycle in G' and G' is consistent.

The behaviors of the event structures G and G' are same and the transformation is safe. \square

Read after Write (RAW)

$$\text{store}_o(\ell, v); \text{load}_{o'}(\ell) \rightsquigarrow \text{store}_o(\ell, v) \text{ when } o' \sqsubseteq o$$

Proof. Let $w = \text{St}_o(\ell, v)$ be the event in the target event structure G from the access $\text{store}_o(\ell, v)$.

We follow the construction steps of target event structure G and construct the source event structure G' .

Let w is appended in the $(i-1)^{\text{th}}$ step and event e is appended in $(i^{\text{th}}$ step to G .

The construction steps of G' is of the form

$$G_{init} \rightsquigarrow_w^* G_{i-1} \rightsquigarrow_e G_i \dots$$

We follow the $(i-1)^{th}$ step of G and construct G' by applying same construction rule and hence appending the equivalent event to G' .

$$G'_{init} \rightsquigarrow_w^* G'_{i-1}$$

Let $r = \text{Ld}_{o'}(\ell, v)$ be an event generated from the $\text{load}_{o'}(\ell)$ of the source program. We append r to G' in the i^{th} step which reads from the write w applying the (R-NORACE) rule.

$$G'_{init} \rightsquigarrow_w^* G'_{i-1} \rightsquigarrow_r G'_i$$

We continue to follow the i^{th} and onwards steps of G and append the equivalent events to G' .

The overall construction of G' following the G looks as follows.

$$G_{init} \rightsquigarrow_w^* G_{i-1} \rightsquigarrow_e G_i \dots$$

$$G'_{init} \rightsquigarrow_w^* G'_{i-1} \rightsquigarrow_r G'_i \rightsquigarrow_e G'_{i+1} \dots$$

Also, considering the derived relations in G and G'

- $G'.\text{fr}(r, e) \implies G'.\text{wb}(w, e) \wedge G.\text{wb}(w, e)$
- The event r in G' does not affect the any derived relation R among other events in G' compared to $G \forall e, e' \in G.V \setminus \{r\}$. $G'.R(e, e') \implies G.R(e, e')$ These also holds when $o' = o = \text{SC}$.

Thus the introduction of r does not introduce any new cycle in G' compared to G and G' is consistent at every construction step.

As a result, the behaviors of the event structures G and G' are same and the transformation is safe. \square

Read after Read (RAR)

$$\text{load}_o(\ell); \text{load}_{o'}(\ell) \rightsquigarrow \text{store}_o(\ell) \text{ when } o' \sqsubseteq o$$

Proof. Let $r = \text{Ld}_o(\ell, v)$ be the event in the target event structure G resulting from memory access $\text{load}_o(\ell)$.

We follow the construction steps of target event structure G and construct the source event structure G' .

Let r is appended in the $(i-1)^{th}$ step and event e is appended in (i^{th}) step to G .

The construction steps of G' is of the form

$$G'_{init} \rightsquigarrow_r^* G'_{i-1} \rightsquigarrow_e G'_i \dots$$

We follow the $(i-1)^{th}$ step of G and construct G' by applying same construction rule and hence appending the equivalent event to G' .

$$G'_{init} \rightsquigarrow_r^* G'_{i-1} \rightsquigarrow_{r'} G'_i$$

Let $r' = \text{Ld}_{o'}(\ell, r.\text{val})$ be an event generated from the $\text{load}_{o'}(\ell)$ of the source program. We append r to G' in the i^{th} step. Cases

- Read r reads from a write w and appended to G' by (R-NORACE). Then r' also reads from w .
- Read r is uninitialized and appended by rule (R-UNINIT). In this case r returns \mathbf{u} . r' also returns \mathbf{u} and appended to G' by (R-UNINIT) rule.
- r races with a write w , returns \mathbf{u} and appended to G' by (R-RACE) rule. The read r' also races with w and returns \mathbf{u} and also appended to G' by (R-RACE) rule.

We continue to follow the i^{th} and onwards steps of G and append the equivalent events to G' .

The overall construction of G' following the G looks as follows.

$$G_{init} \rightsquigarrow_r^* G_{i-1} \rightsquigarrow_e G_i \dots$$

$$G'_{init} \rightsquigarrow_r^* G'_{i-1} \rightsquigarrow_{r'} G'_i \rightsquigarrow_e G'_{i+1} \dots$$

Also, considering the derived relations in G and G'

- If r' results in any derived relation with any event e in G' then r already have the same relation with the event e in G' and G . Hence, $G'.R(r', e) \implies G'.R(r, e) \wedge G.\text{fr}(r, e) \vee e = r'$.
- The event r in G' does not affect the any derived relation R among other events in G' compared to $G \forall e, e' \in G.V \setminus \{r\}$. $G'.R(e, e') \implies G.R(e, e')$ These also holds when $o' = o = \text{SC}$.

Thus the introduction of r' does not introduce any new cycle in G' compared to G and G' is consistent at every construction step.

As a result, the behaviors of the event structures G and G' are same and the transformation is safe. \square

E.1 Non-Adjacent Eliminations

Non-Adjacent Overwritten Write (NAOW)

$$\text{store}_{\text{NA}}(\ell, v'); \text{cmds}; \text{store}_{\text{NA}}(\ell, v) \rightsquigarrow \text{cmds}; \text{store}_{\text{NA}}(\ell, v)$$

where C is a set of commands without any any *release-acquire-pair*.

Let G be the target event structure where

(i) $w' = \text{St}_{\text{NA}}(\ell, v')$ is a write event generated from $\text{store}_{\text{NA}}(\ell, v')$.

(ii) For each $w' \in G$, let C be a po-successor event sequence created from cmd and C is followed by a write $w = \text{St}_{\text{NA}}(\ell, v)$ created from $\text{store}_{\text{NA}}(\ell, v)$. As a result, C does not have *release-acquire-pair* of events.

In the event structure G there can be a set of such event sequence and a set of following writes following w' in a tree structure. Let \mathbb{C} be the set of such event sequences C and \mathbb{W}

be the set of writes where each write $w \in \mathbb{W}$ follows a event sequence $C \in \mathbb{C}$. Thus $w'; \mathbb{C}; \mathbb{W}$ denotes the set of event sequences

$$\{w'; C; w \mid C \in \mathbb{C} \text{ and } w \in \mathbb{W}\}$$

(iii) All the events in C are in total po-order. Let $C[0], C[1] \dots$ be those total po-ordered events in C .

Proof. To prove correctness of this transformation, given a consistent target event structure G (i.e. $G \in \text{exec}[[P^T]_{\text{LLVM}}]$), we construct a consistent source event structure G' ($G' \in \text{exec}[[P^S]_{\text{LLVM}}]$) such that $\text{Behavior}(G) \subseteq \text{Behavior}(G')$. The construction of G' is as follows.

Phase 1. Append w' in G' .

(i) Let c be the first event in \mathbb{C} which is appended in G in i^{th} step immediately after the event a such that $a \notin C$, $c \in C$ and $C \in \mathbb{C}$. Considering this step the construction of G has the form

$$G_{\text{init}} \rightsquigarrow_a^* G_{i-1} \rightsquigarrow_c G_i \dots$$

We follow the construction steps of G in the construction of G' upto $(i-1)^{\text{th}}$ step where we have appended equivalent event of a in G' following the same construction rule. Hence $G'_{i-1} = G_{i-1}$.

(ii) We apply (NON-READ) and append w' to G' in the i^{th} step. Formally

$$G'_{\text{init}} \rightsquigarrow_a^* G'_{i-1} \rightsquigarrow_{w'} G'_i$$

At this point the relations in G'_i are

- $G'_i.V = G_{i-1}.V \uplus \{w'\}$
- $G'_i.\text{po} = (G_{i-1}.\text{po} \cup \{(a, w')\})^+$
- $G'_i.\text{rf} = G_{i-1}.\text{rf}$ and
- $G'_i.R = G_{i-1}.R$ where R are the set of derived relations. as w' is a non-atomic store.

As a result, G'_i is also consistent.

Phase 2. Follow G steps again.

(iii) We follow the i^{th} and onwards steps of G and construct G' with corresponding events. We consider following cases of event e which are appended in G .

(a) case $e \in \text{St}$

(b) case $e \in \mathcal{R}$ In this case we consider two subcases subcase (b.1) $e = r$ such that

- $r.\text{loc} = w'.\text{loc}$
- There exists a write w'' in G and therefore also in G' and *happend-before* w' .
- w'' *happend-before* r in G .
- There exists a release write rel in an event sequence $C \in \mathbb{C}$ such that w'' *happens-before* rel which *happens-before* r in G .

subcase (b.2) otherwise.

case (a) and (b.2)

We append e to G' . The program order and the read-from relations are also updated following the relations in G . Hence the constructed G' remains consistent. In these cases the construction steps of G and G' are as follows.

$$G_{\text{init}} \rightsquigarrow_a^* G_{i-1} \rightsquigarrow_c G_i \rightsquigarrow_{\forall e}^* G$$

$$G'_{\text{init}} \rightsquigarrow_a^* G'_{i-1} \rightsquigarrow_{w'} G'_i \rightsquigarrow_{\forall e}^* G'$$

We apply the same rule in G' construction as the rules applied in G at each step.

case (b.1)

We observe that \mathbb{C} does not contain any event which accesses ℓ . Hence such r may occur only if the *rel* synchronizes with some acquire read acq and acq *happens-before* r . The construction of G is of the form

$$G_{\text{init}} \rightsquigarrow_{w'}^* G_1 \rightsquigarrow_{rel}^* G_2 \rightsquigarrow_{acq}^* G_3 \rightsquigarrow_b^* G_4 \rightsquigarrow_r G_5 \dots$$

Phase 3. Handle case (b.1).

We append the events in G' following the steps of G applying the same rules until the next event appended to G is r . The construction steps in G and G' are

$$G_{\text{init}} \rightsquigarrow_b^* G_4 \rightsquigarrow_r G_5 \dots$$

$$G'_{\text{init}} \rightsquigarrow_b^* G'_4$$

The G' is consistent as G is consistent.

(iv) Next, we consider appending $\mathcal{M}(r)$ in G' . Two cases

• **case (b.1.1)**

There is a write w in G' (specifically G'_4) such that $G'.\text{race}(w, r)$.

Instead of using the (R-NORACE), we apply (R-RACE) and create $\mathcal{M}(r)$ in G' which returns \mathbf{u} . In this case also G' has less number of relations compared to G . hence G' is less constrained compared to G and as a result, we can create all behaviors which are present in G .

• **case (b.1.2)**

We cannot append r directly in G' as in G' the write w'' is over-written by w' in G' . Hence we want to append $\mathcal{M}(r)$ in G' .

Let \mathbb{W}_1 be the set of writes such that $\mathbb{W}_1 \subseteq \mathbb{W}$ and all the writes in \mathbb{W}_1 are po-successor to the *rel*.

Let \mathbb{R} be the set of reads in G between the *rel* and \mathbb{W}_1 . For any value a read $r' \in \mathbb{R}$ returns, there exists a write in \mathbb{W}_1 .

Starting next from the *rel*, we follow the steps of G to construct G' until the next event is in \mathbb{W}_1 . Let event e is appended in G between *rel* and the first event of \mathbb{W}_1 .

Considering the cases of event e

- **case (b.1.2.1)** Event e is a write event which is appended to G using (NON-READ) rule. We append the same event in G' using (NON-READ) rule. The G' remains consistent.
- **case (b.1.2.2)** e is a read event. Subcases are as follows
 - **case (b.1.2.2.1)** e is appended to G by (R-NORACE) rule and reads from a write w whose equivalent event is already appended in G' . We append e in G' which reads from the corresponding write $\mathcal{M}(w)$ using (R-NORACE) rule.
 - **case (b.1.2.2.2)** e is appended to G by (R-RACE) rule, returns \mathbf{u} and races with a write w in G and $\mathcal{M}(w)$ is already appended in G' . The event e is also appended to G' using (R-RACE) rule.
 - **case (b.1.2.2.3)** The read event e is appended to G by (R-NORACE) or (R-RACE) rule, reads from or races with (and returns \mathbf{u}) some write w_2 in G respectively. The corresponding event of w_2 is not yet appended in G' . Two possible subcases
 - **case (b.1.2.2.3.1)** It is possible to append the corresponding prefix and $\mathcal{M}(w_2)$ to G' following the steps in G . We do so, and then create and append the e in G' following the corresponding rule applied in G .
 - **case (b.1.2.2.3.2)** We need to create $\mathcal{M}(e)$ which is different from e label. It is possible in certain cases, for example, if a write w_2 is po-successor of r and $G.\text{race}(w_2, e)$. As w_2 is not yet present in G' , we must create $\mathcal{M}(e)$ which is different e .
In this case we append $\mathcal{M}(e)$ to G' using (R-NORACE) rule and $\mathcal{M}(e)$ reads from some wb-maximal write event in G' such that G' is consistent.
Note that, we later create an appropriate write event from which we will create a read event e'' in immediate conflict with $\mathcal{M}(e)$ such that $e''.\text{rval} = e.\text{rval}$ or $e''.\text{rval} = \mathbf{u}$.
The po-successor paths of e'' would also have some write in \mathbb{W}_1 . As a result, we can preserve all the behaviors of G in G' .

In all cases e is a non-atomic read and for all cases G' is consistent. To elaborate, either we read from the write which does not affect consistency as in G or we read from the wb-maximal write and hence no frsc is created which results in inconsistent G' .

Thus we will be able to follow the construction steps of G and reach a write in \mathbb{W}_1 as all po-successor paths of rel have a write in \mathbb{W}_1 . At the same time, we append corresponding events to G' .

(v) Next, we consider r of G and append $\mathcal{M}(r)$ to G' using (R-RACE) rule. As we have already created $\mathcal{M}(w_1)$ where $w_1 \in \mathbb{W}_1$ with which $\mathcal{M}(r)$ races in G' , the $\mathcal{M}(r)$ returns \mathbf{u} .

Phase 4. Follow the rest of G' .

As $r.\text{rval}$ is a specialized value of $\mathcal{M}(r).\text{rval} = \mathbf{u}$, we append more events to G' which

- corresponds to the $(\text{po} \cup \text{rf})^+$ -successor of r in G or
- corresponds to the events which races with $(\text{po} \cup \text{rf})^+$ -successor of r in G .

(vi) After appending r or $\mathcal{M}(r)$ in G' , we follow the steps of G in construction of G' . In a step of G , if G' already contains a corresponding event then we skip the step. Otherwise we create an equivalent event and continue from there. Especially for the read events we create immediately conflicting corresponding events.

(vii) Thus cmds in the source results in \mathbb{C}' in G' and in the target results in \mathbb{C} in G . If e in \mathbb{C} has an equivalent event $\mathcal{M}(e)$ in \mathbb{C}' then

- $e.\text{rval} = \mathcal{M}(e).\text{rval}$ or
- $\mathcal{M}(e).\text{rval} = \mathbf{u}$ or
- $e.\text{wval} = \mathcal{M}(e).\text{wval}$ or
- $\mathcal{M}(e).\text{wval} = \mathbf{u}$

Thus event structure G' has all the behaviors of G and as a result the transformation is safe. \square

Non-Adjacent Read After Write (NARAW)

$$\text{store}_{\text{NA}}(\ell, v); \text{cmds}; \text{load}_{\text{NA}}(\ell) \rightsquigarrow \text{store}_{\text{NA}}(\ell, v); \text{cmds}$$

where cmds has no *release-acquire-pair*.

Proof. Let $w = \text{St}_{\text{NA}}(\ell, v)$ be an event generated from $\text{store}_{\text{NA}}(\ell, v)$ in the target event structure G . Let C be a sequence of po-ordered sequence of events generated from the cmds and event c is the last event in c in G . Let w , c and e be the events appended to G in k^{th} , $(i-1)^{\text{th}}$ and i^{th} step respectively where $k < i-1$. The construction of G is of the form

$$G_{\text{init}} \rightsquigarrow_c^* G_{i-1} \rightsquigarrow_e G_i \dots$$

We follow the $(i-1)^{\text{th}}$ step of construction of G and we introduce a read event $r = \text{Ld}_{\text{NA}}(\ell, v)$ in the i^{th} step which reads from the w and then continue following G for i^{th} and onwards steps. Thus the construction of G and G' are as follows.

$$G_{\text{init}} \rightsquigarrow_w^* G_k \rightsquigarrow_c G_{i-1} \rightsquigarrow_e^* G_i \dots$$

$$G'_{\text{init}} \rightsquigarrow_w^* G'_k \rightsquigarrow_c G'_i \rightsquigarrow_e^* G' \dots$$

We apply the same rule in G' construction as the rules applied in G at each step. The r is appended to G' using the (R-NORACE) rule.

The consistency of G' is preserved though similar to the RAW elimination, the introduction of r results in new relation.

Thus behavior of the target event structure G is preserved in the source event structure G' and hence the transformation is safe. \square

F. Access Strengthening

Definition 5. Given two event structures G and G' , the $G \sqsubseteq G'$ denotes that the G' is a sub-structure of G . The G' contains a subset of events of G i.e. $G'.\mathcal{V} \subseteq G.\mathcal{V}$ and the relations and global memory events are also subsets of G respectively i.e. $G'.\text{po} \subseteq G.\text{po}$, $G'.\text{rf} \subseteq G.\text{rf}$.

The access strengthening theorem is as follows:

Theorem 5. $\forall P \xrightarrow{\text{strengthen}} P'$.

$\text{Behavior}_{\text{LLVM}}(P') \subseteq \text{Behavior}_{\text{LLVM}}(P)$

Proof outline. We prove this by induction over the construction of the event structure G of the program P . In the weakened program P' , we construct an event structure G' which contains all the events and relations of G i.e. $G \sqsubseteq G'$ (cf. Definition 5).

The set of visible writes for a particular sequentially consistent read is the subset of the visible writes if the read were acquire or non-atomic. Let $\text{events}_{\text{sc}}(C_r)$ be the set of read events created from SC-read access C_r and $\text{events}_{\sqsubseteq\text{sc}}(C_r)$ be the set of read events when we weaken the memory or-

der of C_r . Following the event structure construction rules $\text{events}_{\text{sc}}(C_r) \subseteq \text{events}_{\sqsubseteq\text{sc}}(C_r)$. In this case $G \sqsubseteq G'$ where G is the event structure when C_r is sequentially consistent and G' is the event structure when C_r is acquire or non-atomic.

When we weaken a release write to a non-atomic write or an acquire read to a non-atomic read then the value of a set of read events are affected in the event structure since certain synchronizations may be removed. If a read event in G returns a concrete value then after the weakening the corresponding read event in G' would return the same value or an \mathbf{u} expression. If a read returns \mathbf{u} in G then after the weakening the corresponding read in G' returns \mathbf{u} . Since \mathbf{u} is a generalization of any value all the events and relations of G are preserved in G' . As a result $G \sqsubseteq G'$.

$$\begin{aligned} \forall P \xrightarrow{\text{strengthen}} P'. \forall G' \in \llbracket P' \rrbracket'_{\text{LLVM}}. \\ \exists G \in \llbracket P \rrbracket_{\text{LLVM}}. G' \sqsubseteq G \end{aligned}$$

As a result $\text{Behavior}(G') \subseteq \text{Behavior}(G)$ and in turn $\text{Behavior}_{\text{LLVM}}(P') \subseteq \text{Behavior}_{\text{LLVM}}(P)$ holds. \square

G. Compilation: C11 to LLVM

Theorem 6. *For every program P , $\text{Behavior}_{\text{LLVM}}(P) \subseteq \text{Behavior}_{\text{C11}}(P)$.*

To prove this theorem, consider an LLVM behavior, \mathcal{X} , generated by an event structure G . If G contains no **undef** reads because of read-write races, then it is easy to show that \mathcal{X} is also a consistent C11 execution.

If the event structure G contains an **undef** read due to a read-write race, then we will show that P has undefined semantics according to C11. Let G' be the graph in the derivation of G right before the first application of the R-

RACE rule, which is about to add the new read event e' after the event $e \in G'.\mathcal{V}$, where e' races with the existing write event $w \in G'.\mathcal{V}$ on the memory location ℓ . Since the events e and w are not in conflict, we can take an execution \mathcal{X}' of G' containing the events e and w . We can further extend \mathcal{X}' by adding a read event e'' after e reading a happens-before maximum value for the memory location ℓ . (We know such a value exists because $\text{hbW}(\ell, e'')$.) It is easy to show that the constructed execution is an prefix execution of the program P , that it is consistent according to C11, and has a race between e'' and w . Thus, P has undefined semantics according to C11, as required.

H. Compilation: LLVM to PowerPC

H.1 Definitions

- *dependency* (deps). The set of data, address and control dependencies from reads/updates to other events is represented by $\text{deps} (\subseteq \text{po})$ i.e. $\text{deps} = \{(a, b) \mid \text{po}(a, b) \wedge a \in \mathcal{R} \wedge b \in \mathcal{RW}\}$.
- *preserved program order* (ppo). If deps is the dependency relation among events and must be preserved then $\text{ppo} \subseteq (\text{deps} \cup \text{po}_\ell)^+$.
- *read-from* (rf) is same as the read from (rf) in Section 3 where $(\text{rf}(r) = w) \implies (r.\text{rval} = w.\text{wval} = c)$ where c is some constant value. Also external read from $\text{rf} = \text{rf} \setminus \text{po}$.
- *modification order* (mo). It is a binary relation among the same location writes or updates and imposes a total order in the execution. The mo_x is the mo on some location ℓ . It is also known as the *coherence order*.
- *from-read* (frp). This is similar to the from-read defined in Section 3.2. If a read r reads from a write w_1 and w_2 is another write mo ordered after w_1 then $(r, w_2) \in \text{frp}$ i.e. $\text{frp} \triangleq (\text{rf}^{-1}; \text{mo}) \setminus [\mathcal{X}.A]$ where $\mathcal{X}.A$ is the set of power execution events. Also external from read $\text{fre} = \text{frp} \setminus \text{po}$.
- *fence orders*:

$$\begin{aligned} \text{lwsync} &\triangleq \text{po}; [\text{lwsync}]; \text{po} \\ \text{hwsync} &\triangleq \text{po}; [\text{hwsync}]; \text{po} \\ \text{fence} &\triangleq \text{hwsync} \cup ([\mathcal{R}]; \text{lwsync}; [\mathcal{RW}]) \\ &\quad \cup ([\text{St}]; \text{lwsync}; [\mathcal{W}]) \end{aligned}$$

- Power's *happens-before*. $\text{hbp} \triangleq \text{ppo} \cup \text{fence} \cup \text{rfe}$
- *Propagation* (prop) denotes the fence enforced ordering of the write events.
 $\text{prop} = \text{prop}_1 \cup \text{prop}_2$ where
 $\text{prop}_1 = ([\mathcal{W}]; \text{rfe}^?; \text{fence}; \text{hbp}^*; [\mathcal{W}])$ and
 $\text{prop}_2 = (((\text{mo} \cup \text{frp}) \setminus \text{po})^?; \text{rfe}^?; (\text{fence}; \text{hbp}^*)^?; \text{hwsync}; \text{hbp}^*)$

H.2 Axioms

The SPower axioms are as follows.

1. $\text{order}(\text{iswrite}, \text{mo}) \wedge \forall \ell. \text{total}(\text{iswrite}_\ell, \text{mo})$ (total-MO)
2. $\nexists a. \text{hbp}(a, a)$ (no-thin-air)
3. $\forall \ell \in \text{Loc}. \text{acyclic}(\text{po}_\ell \cup \text{rf} \cup \text{frp} \cup \text{mo})$ (SC-per-loc)
4. $\text{fre}; \text{prop}; \text{hbp}^*$ is irreflexive. (observation)
5. $(\text{mo} \cup \text{prop})$ is acyclic. (propagation)
6. $\text{frp}; \text{mo}$ is irreflexive (atomicity)
7. $\text{mo}; [\mathcal{U}]; \text{po}; [\mathcal{U}]$ is acyclic.
8. $\text{acyclic}(\text{po} \cup \text{rf})$ (Aporf)

H.3 LLVM and SPower Relations

We relate the LLVM relations with the SPower relations so that we can relate the respective events in LLVM event structure which is constructed following a given SPower consistent execution.

We first relate the hb relation in LLVM with relations in SPower.

Lemma 15. $\text{hb} \subseteq \text{po} \cup \text{prop}; \text{hbp}^*$

Proof. Therefore two cases for $\text{hb}(a, b)$

(1) $a.\text{tid} = b.\text{tid}$.

In this case $\text{po}(a, b)$.

(2) $a.\text{tid} \neq b.\text{tid}$.

In this case $(a, b) \in \text{po}; ([\text{Rel}]; \text{rfe}; [\text{Acq}]; \text{po})^*$.

From the compilation scheme in Figure 13 we know that $\text{po}; [\text{Rel}]$ creates a fence in SPower. Similarly $[\text{Acq}]; \text{po}$ also creates a fence in SPower.

Thus in SPower $(a, b) \in (\text{rfe} \cup \text{fence})^+$.

As a result, $\text{iswrite}(b) \implies ((a, b) \in \text{prop})$ and

$\text{isread}(b) \implies ((a, b) \in \text{prop}; \text{hbp}^+)$.

Which means $(a, b) \in \text{prop}; \text{hbp}^*$.

Hence $\text{hb} \subseteq \text{po} \cup \text{prop}; \text{hbp}^*$. \square

The wb in LLVM and mo in SPower relate the writes on a particular location. We show that the resulting wb relations in LLVM correspond to a subset of mo relations in SPower.

Lemma 16. $[\mathcal{W}]; \text{brf} \subseteq \text{mo}^?$

Proof. Recall that $\text{brf} = (\text{rf}^{-1})^*$. The lemma follows directly from (SC-per-loc) and (total-MO). \square

Lemma 17. $\text{wb} \subseteq \text{mo}$

Proof. We first define an update sequence $\text{Useq}(w, u)$ to denote a sequence of rf edges starting from the write w and ending at the update u , i.e. $w \xrightarrow{\text{rf}} u_1 \dots \xrightarrow{\text{rf}} u$.

$\text{brf} \subseteq \text{mo}^?$

Note that the $\text{Useq} \subseteq \text{mo}$ because $\text{acyclic}(\text{mo} \cup \text{rf})$ holds (SC-per-loc).

For $\text{wb}(e, w)$ consider the most general case

$$\begin{aligned} \exists w', r. \text{brf}(e, w') \wedge \text{hb}(w', r) \wedge \text{locs}(w, r) \\ \wedge \text{brf}(r, w) \wedge \neg \text{brf}(e, w) \end{aligned}$$

From Lemma 16, we know that $\text{mo}^?(e, w')$. It therefore suffices to show $\text{mo}(w', w)$.

By Lemma 15, we have $\text{po}(w', r) \vee (\text{prop}; \text{hbp}^*)(w', r)$.

Case $r = w$, $\text{po}(w', r)$. Then $\text{mo}(w', w)$ because of (SC-per-loc).

Case $r = w$, $(\text{prop}; \text{hbp}^*)(w', r)$. Then $\text{prop}(w', w)$, and thus $\text{mo}(w', w)$ by (propagation)

Case $r \neq w$. Then, there exists $u \in \mathcal{U}$ such that $\text{rf}(u, r)$ and $\text{brf}(u, w)$. From Lemma 16, we know that $\text{mo}^?(u, w)$. It therefore suffices to show $\text{mo}(w', u)$.

Assume that $\text{mo}(u, w')$ and so $\text{fr}(r, w')$.

Subcase $\text{po}(w', r)$. This contradicts (SC-per-loc).

Subcase $(\text{prop}; \text{hbp}^*)(w', r)$.

If $r.\text{tid} = w'.\text{tid}$, then it contradicts (SC-per-loc).

If $r.\text{tid} \neq w'.\text{tid}$, then it contradicts (observation). \square

Lemma 18. $fr \subseteq frp$

Proof. $fr = (rf^{-1}; wb) \subseteq (rf^{-1}; mo) = frp$ by Lemma 17. \square

H.4 Compilation Correctness

Let P be a LLVM program and P' be its mapping to Power. We have to show that every behavior of P' under the SPower memory model is also a behavior of P under the LLVM model. So, consider a SPower-consistent execution \mathcal{X} . It suffices to show that the LLVM semantics can construct an event structure of P with the same memory accesses. To construct the event structure, we will add one by one the events of \mathcal{X} . Since $(po \cup rf)$ is acyclic, a natural order to consider is the partial order $(po \cup rf)^+$. In the case of events unordered by $(po \cup rf)^+$, we order writes before loads.

Suppose we have already constructed an event structure G corresponding to a prefix of \mathcal{X} and the next event to add is e_p in \mathcal{X} . If e_p is a store, we simply add the corresponding store e' to G using the NON-READ rule. The corresponding store e will be a SC-store if there is hwsync immediately before e_p in \mathcal{X} .po, a REL-store if there is lwsync immediately before e_p in \mathcal{X} .po, and a non-atomic store otherwise. If the resulting event structure has any *write-write* race then we abort the construction since the program behavior is undefined and the compilation is correct.

Now, consider the case when e_p is a load. Let w_p be the event in \mathcal{X} such that $\mathcal{X}.rf(w_p, e_p)$. Note that since we are adding events in (an extension of the) $(po \cup rf)^+$ order, there exists already an event w in G that corresponds to w_p . Our task is to extend G with a load e corresponding to e_p returning either the same value or \mathbf{u} . This will be a SC-load if there is a preceding hwsync fence and a successor lwsync fence, a ACQ-load if there is only a successor lwsync fence, and a non-atomic read otherwise. If the location is uninitialized in G , then we can use R-UNINIT rule and make e' return \mathbf{u} . If, however, the location is initialized, we have to show that the addition is allowed by the R-RACE or R-NORACE rules, and the resulting G would be consistent. The final case, where the e_p is an update is similar to the case of a read. We prove the consistency of G in by showing that w is indeed visible to e .

Essentially we prove that the $\text{isCons}(G)$ conditions holds by considering the conditions.

Lemma 19.

- (i) $\text{irreflexive}(cf; hb)$
- (ii) $\text{irreflexive}(rf; hb^{-1}; rf^{-1}; cf)$
- (iii) $\text{irreflexive}(wb)$

Proof. (i) and (ii) hold trivially because $cf = \emptyset$.

(iii) We know that mo is irreflexive and $wb \subseteq mo$ (Lemma 17). Hence $\text{irreflexive}(wb)$ holds. \square

Lemma 20. $\text{acyclic}((hbsc \cup fr \cup wb); [SC])$

Proof. We prove this by contradiction.

Assume that $[SC]; (hbsc \cup fr \cup wb); [SC]$ sequence is a cycle.

$wb; fr \implies mo$ and $fr; wb \implies fr$.

We also know $wb \subseteq mo$ and hence $fr \subseteq frp$.

Thus in the sequence all the mo and frp edges are separated by the $hbsc$.

The final edge in the $hbsc$ can be either po or rf . Hence call them $pohbsc$ and $rfhbsc$ respectively i.e.

Note that,

- (i) $[SC]; po|_{\neq locs}; [SC] \subseteq pohbsc$,
- (ii) $[SC]; po|_{\neq locs}; hb; po|_{\neq locs}; [SC] \subseteq pohbsc$ and
- (iii) $[SC]; (hb \cap loc); [SC] = pohbsc \cup rfhbsc$.

Without loss of generality the sequence is $mo; (pohbsc/rfhbsc); frp; (pohbsc/rfhbsc) \dots$

Now $mo; rfhbsc \implies mo; rf$ and $frp; rfhbsc \implies frp; rf$.

Thus the sequence is transformed to

$(mo; pohbsc/mo; rf); (frp; pohbsc/frp; rf) \dots$

Now $rf; mo \implies mo$ and $rf; frp \implies mo$.

Hence the sequence is of $mo, frp, pohbsc$ and the mo and the frp are separated by the $pohbsc$.

Also $mo; pohbsc \implies \text{prop}_2$

$fr; pohbsc \implies \text{prop}_2$.

As a result, the sequence consists of a sequence of prop and cannot form a cycle. (propagation).

Hence a contradiction and

$\text{acyclic}((hbsc \cup fr \cup wb); [SC])$ holds. \square

Lemmas 19 and 20 ensure that the constructed LLVM event structure G is consistent and thus the LLVM to PowerPC mappings are correct.




Figure 27. $t == 0, t' == 3$ is Forbidden in C11.

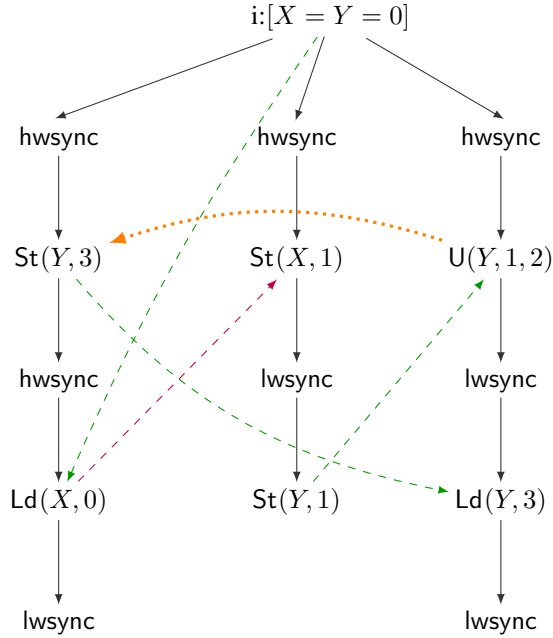


Figure 28. Allowed $t == 0, t' == 3$ in PowerPC.

I. SC Counter-Example

Consider the following program:

$$Y_{SC} = 3; \parallel X_{SC} = 1; \parallel \text{if}(\text{CAS}_{SC}(Y, 1, 2)) \\ t = X_{SC}; \parallel Y_{REL} = 1; \parallel t' = Y_{ACQ};$$

Behavior in question is $t = 0, t' = 3$. This execution is disallowed in C11 as a, b, c, d, e form a $(\text{hb} \cup \text{fr} \cup \text{mo}); [\text{SC}]$ cycle in the Figure 27. However, the compiled program to PowerPC may have this execution as shown in Figure 28. As a result, the C11 model is stronger than the PowerPC model and the C11 to PowerPC compilation is unsound.