Grounding Thin-Air Reads with Event Structures

SOHAM CHAKRABORTY, MPI-SWS, Germany

VIKTOR VAFEIADIS, MPI-SWS, Germany

The key challenge in defining the concurrency semantics of a programming language is how to enable the most efficient compilation to existing hardware architectures, and yet forbid programs from reading thin-air values, i.e., ones that do not appear in the program. At POPL'17, Kang et al. achieved a major breakthrough by introducing the 'promising' semantics that came with results showing that it was a good candidate solution to the problem. Unfortunately, however, the promising semantics is rather complicated, and due to its complexity it contains some flaws and limitations that are very hard to address.

In response, we present an alternative solution to this problem based on event structures. We show that it is indeed a solution by establishing the standard results about the semantics (DRF theorems, implementation and optimization correctness) as well as a formal connection to the semantics of Kang et al. Further, we show that it is easier to adapt, by extending the semantics to cover features (such as SC accesses) that are not supported by Kang et al. and to rule out some dubious behaviors admitted by the promising semantics.

CCS Concepts: • **Software and its engineering** → **Semantics; Compilers**; • **Theory of computation** → **Concurrency**;

Additional Key Words and Phrases: Weak memory consistency, concurrency, promising semantics, C/C++11

ACM Reference Format:

Soham Chakraborty and Viktor Vafeiadis. 2019. Grounding Thin-Air Reads with Event Structures. *Proc. ACM Program. Lang.* 3, POPL, Article 70 (January 2019), 28 pages. <https://doi.org/10.1145/3290383>

1 INTRODUCTION

By now, relaxed memory consistency has become the de facto concurrency paradigm for architectures, languages, and compilers. Yet, there is still no agreed adequate definition of the concurrency semantics of a language like C or Java. Developing such a definition is very challenging because it must be strong enough to eliminate undesirable behaviors and provide certain programming guarantees, and yet weak enough to allow efficient compilation—to support compiler optimizations and optimal mappings to the mainstream architectures.

We demonstrate the conflicting requirements with an example. Assume that initially $X = Y = 0$ and all shared memory accesses are C11 'relaxed' atomic accesses [ISO/IEC 14882 2011; ISO/IEC 9899 2011]. The question is whether the outcome $a = b = 1$ should be allowed.

$a = X; \quad \parallel \quad b = Y;$ $\text{if}(a) \quad \parallel \quad \text{if}(b)$ $Y = 1; \quad \parallel \quad X = 1;$	(CYC)	$a = X; \quad \parallel \quad b = Y;$ $Y = 1; \quad \parallel \quad \text{if}(b)$ $X = 1;$	(LB)	$a = X; \quad \parallel \quad b = Y;$ $\text{if}(a) \, Y = 1; \quad \parallel \quad \text{if}(b)$ $\text{else } Y = 1; \quad \parallel \quad X = 1;$	(LBfd)
---	-------	--	------	--	--------

In the **CYC** program, this outcome is considered an “out-of-thin-air” (OOTA) behavior that should be forbidden. Allowing it would violate the data-race-freedom (DRF) property [Adve and Hill 1990],

Authors' addresses: Soham Chakraborty, MPI-SWS, Germany, sohachak@mpi-sws.org; Viktor Vafeiadis, MPI-SWS, Saarland Informatics Campus, Germany, viktor@mpi-sws.org.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/1-ART70

<https://doi.org/10.1145/3290383>

which states that programs without races under *sequential consistency* (SC) [Lamport 1979] should not have any non-SC behaviors. Indeed, under SC, **CYC** is race-free and has only the outcome $a = b = 0$. In contrast, the outcome $a = b = 1$ should be allowed for **LB** because it arises when the compiler or the hardware reorders the (clearly independent) instructions of the first thread. This outcome should also be allowed for **LBfd** because it may be converted into **LB** by an optimizing compiler. As noted by Batty et al. [2015], these programs cannot be correctly distinguished by “per execution” models (e.g., C/C++11 [Batty et al. 2011], Linux kernel model [Alglave et al. 2018], and all the hardware models).

To solve this problem, a number of alternative modeling approaches have emerged [Jeffrey and Riely 2016; Kang et al. 2017; Manson et al. 2005; Pichon-Pharabod and Sewell 2016]. The most complete among these approaches is the *promising semantics* (PS) of Kang et al. [2017], who also claim to have resolved the conflicting requirements. Their model forbids the mentioned OOTA behavior in **CYC**, provides DRF guarantees, and supports the expected program transformations and mappings to x86-TSO and PowerPC. Although PS is a major step towards providing a formal model for C11 concurrency, as with the other approaches, it suffers from certain problems, which we discuss in detail in §7. For instance, PS’s intended compilation scheme for read-modify-write instructions to ARMv8 is unsound because PS forbids the weak behavior of the **FADD** program shown in Fig. 6, while it is allowed by ARMv8. To address these problems, we tried to adapt the PS definition but found its definition style to be inflexible: any small change affected nearly all the results about the semantics.

In this paper, we develop an alternative approach to formalize the semantics of concurrent programs. Following Jeffrey and Riely [2016], we use *event structures* [Winskel 1986], a formalism that captures multiple executions of a program in one structure without referring to the concrete program syntax. We adapt them to the weak consistency setting by introducing in §2 the notions of *justified* and *visible* events. Event structures differentiate between the **CYC** and **LB** programs, but not between **LB** and **LBfd**, which enables us to give the right semantics to these programs.

In our experience, event structures are equally powerful (in that we establish essentially the same results as Kang et al. [2017]) but much more flexible than PS. To demonstrate their flexibility, we present two variants of our model, both of which allow the expected weak behavior of **FADD**.

The first variant is the **WEAKEST** (“*weak event structure*”) model. In §4, we prove that it is weaker than PS, which allows us to leverage the existing compilation results about PS. Because of this, however, **WEAKEST** allows certain dubious program behaviors (e.g., that of **Coh-CYC** in Fig. 3).

The second variant, **WEAKESTM0**, avoids such undesirable behaviors by maintaining a per-location *modification order*, which totally orders any two non-conflicting writes to the same location. Consequently, **WEAKESTM0** is strictly stronger than **WEAKEST** and incomparable to PS. Moreover, it meets the established criteria for a good memory model:

- It disallows *OOA behavior*—at least for the established Java causality tests (§3.5).
- It provides *DRF guarantees*, which allow programmers to write concurrent code defensively and reason about their code without needing to understand the **WEAKESTM0** definition (§5).
- It supports *efficient compilation*. We prove the correctness of the mapping from C11 to **WEAKESTM0** (§6.1), source-to-source optimizations within **WEAKESTM0** (§6.2), and mappings to x86 and PowerPC (§6.3). For simplicity, in the case of PowerPC, we establish the correctness of a suboptimal mapping with a spurious conditional branch after each read.
- It supports *all* types of memory accesses, fences, and memory orders in C/C++11 except ‘consume’ atomics, which are known to be problematic.

In the following, we start in §2 with an informal overview of our approach. From §3 onwards, we give the formal details. The proofs can be found in our appendix [Chakraborty and Vafeiadis 2018].

Fig. 1. One event structure of **LB**, two executions extracted from it, and the only event structure of **CYC**.

2 MAIN IDEAS

In this overview section, we gradually explain our event structure models (§2.1-§2.3) and compare them with the promising semantics of Kang et al. [2017] (§2.4).

2.1 Event Structures as a Solution to the OOTA Problem

We define the semantics of programs in two steps. First, we iteratively construct a “consistent” event structure from the program and then we extract its “consistent” executions.

An event structure is a representation of multiple executions of a program. It comprises a set of events representing the program’s individual memory accesses and two relations over these events:

- The *program order* (*po*) represents the control flow order of the program; it relates two events a and b if they belong to the same thread and a is executed before b .
- The *conflict relation* (*cf*) relates events of the same thread that *cannot* belong to the same execution. A conflict, for example, occurs between two load events correspond to the same program instruction that return different values. By definition, *cf* is a symmetric relation. Moreover, when a and b conflict, and c is after b in program order, then a and c also conflict. That is, a conflict results in two different branches of the event structure, after which all events conflict with one another. For conciseness, in diagrams (e.g., in Fig. 1a), we display only the first conflicting edges (a.k.a. immediate conflicts) with $\sim$ and omit the induced ones.

To these basic event structures, we add an extra relation, which we call the *justified-from* relation (*jf*). For each read event in the event structure, *jf* associates a (prior) write event that *justifies* the read returning its value. Namely, the write should have written the same value to the same location that the read is reading from.

An event structure is constructed by appending one event at a time following the program text. For a write event to be added to the structure, all its *po*-previous events must have been added; for a read event, there must additionally already exist a justifying write event. Thus, by construction ($\text{po} \cup \text{jf}$) is acyclic. The construction terminates when we can no longer add a new event to the structure. In general, whenever we add an event to an event structure, we check that the resulting structure is *consistent* (i.e., whether it satisfies a few consistency conditions that will be defined later) and discard any event structures that do not satisfy those conditions. Among other things, consistency ensures that sequential programs have their expected semantics by forbidding, for instance, a read to be justified by a stale write—one overwritten by a more recent write to the same location *po*-before the read.

As an example, Fig. 1a displays one consistent event structure of **LB**. Note that we cannot construct this event structure for program **CYC**, because once we add the $\text{Ld}(X, 0)$ and $\text{Ld}(Y, 0)$ events as shown in Fig. 1d, no further events can be added.

Fig. 2. A naively constructed event structure of RNG and C11-consistent extracted execution from it.

Next, given a consistent event structure, we extract the set of its maximal conflict-free subsets, which we call *executions*. (These are called configurations by Winskel [1986].) As an example (ignoring the *rf* edges for the moment), Figs. 1b and 1c show the two executions that can be extracted from the event structure in Fig. 1a. We filter the set of extracted executions by a certain execution consistency predicate, discarding those executions that are inconsistent. The allowed outcomes of a program are determined by these consistent executions.

One standard consistency requirement for executions is that every read should get its value from some write to the same location. This is formalized by a *reads-from* (*rf*) edge from the write to the read. Other consistency axioms may further constrain the allowed reads-from edges.

In this paper, we take execution consistency predicate to be that of RC11 (i.e., the repaired definition of C11 by Lahav et al. [2017]) without its $\text{po} \cup \text{rf}$ acyclicity requirement. We thus get a model that is *stronger* than C11 (which allows OOTA) and *weaker* than RC11 (which forbids load-store reordering). In particular, it assigns the right meaning to *CYC*, *LB*, and *LBfd*.

In general, due to different ways in which events may be appended to an event structure, a program may have multiple consistent event structures, and each such event structure may yield one or more consistent executions.

2.2 A Problem with the Simple Construction Scheme and How to Solve it

Although the simple scheme presented so far works well for *LB* and its variants, it suffers from a subtle problem demonstrated by the following “random number generator” litmus test.

$$Y = X + 1; \parallel X = Y; \parallel \begin{cases} \text{if } (X == 100) \\ X = 99; \end{cases} \quad (\text{RNG})$$

First, as shown in Fig. 2, even for this obviously terminating program, the event structure construction can go on for ever. Thread 1 reads $X = 0$ and writes $Y = 1$; thread 2 then reads $Y = 1$ and writes $X = 1$. Then, in a conflicting execution branch, thread 1 reads $X = 1$ and writes $Y = 2$; thread 2 reads this value and writes $X = 2$. Then, thread 1 can read $X = 2$ and the process can be repeated indefinitely. This problem is not so serious because one can always stop the construction at some arbitrary point. The more serious problem is that the construction can generate an event structure containing the $\text{St}(X, 99)$ event of the last thread. From there, we can extract the execution shown at the right part of Fig. 2, which is consistent according to C11 (as witnessed by the depicted *rf* edges) and leads to a very counterintuitive outcome.

Clearly, the straightforward approach of constructing an event structure does not work here. In this example, we might blame the addition of $\text{Ld}(X, 1)$, as it is the first event that will never appear in any reasonable run of the program. From then on all other events added (shown in red in Fig. 2) are equally bogus. We may observe that $\text{Ld}(X, 1)$ is $(\text{po} \cup \text{jf})^+$ -reachable from $\text{Ld}(X, 0)$ in the event structure and thus the creation of $\text{Ld}(X, 1)$ *causally depends* on its conflicting event $\text{Ld}(X, 0)$. As the two events are in conflict, they cannot both occur in a single execution, so perhaps a possible constraint is that *no event should causally depend on a conflicting event*.

Unfortunately, this restriction turns out to be too strong because it also rules out the $a = b = 1$ outcome of **LB** (cf. Fig. 1a). We therefore relax the restriction by introducing the notion of *visibility*, and adapt the extraction of executions from event structures to *discard any executions containing invisible events*. As we will shortly see, all events of Fig. 2 drawn in red are invisible, and so the problematic extracted execution shown to right of the figure is discarded.

To define the notion of visible events, we introduce the *equal-write* (**ew**) relation, which relates two writes on the same location with the same written value that are in conflicting execution branches. For instance, in Fig. 1a, the two $\text{St}(Y, 1)$ events are **ew**-related, whereas in Fig. 2 no events are **ew**-related.

Given this relation, we say that an event e is *invisible* whenever there is a conflicting write event in its $(\text{po} \cup \text{jf})^+$ -prefix that does not have an equal write in the same execution branch as e (i.e., $(\text{po}^? \cup \text{po}^{-1})$ -related to e). For example, the $\text{Ld}(X, 1)$ event is invisible in Fig. 2 because of the $\text{St}(Y, 1)$ store, whereas $\text{Ld}(X, 1)$ is visible in Fig. 1a. Note that by definition if an event is invisible in an event structure, then so are all its *po*-later events.

The equal-write relation also allows to make a formal connection between the *justification* relation (**jf**) at the level of event structures and the *reads-from* relation (**rf**) at the level of extracted executions. The idea is to also define **rf** at the level of event structures in terms of **jf** and **ew**. We say that a read event r *reads from* a write w if it is justified by w or by one of its equal writes. When extracting an execution from an event structure, we take their **rf** relations to match for the extracted events. (We invite the reader to check that this is indeed so for the executions in Fig. 1).

2.3 Two Options Regarding Coherence: The **WEAKEST** and **WEAKESTM0** Models

We move on to *coherence*, a basic property enforced by almost all memory models.¹ Coherence states that programs containing only one shared variable exhibit only SC behaviors. To ensure this property, axiomatic memory models require the existence of a (strict) total order among all writes to a given location. We call the union of all these per-location orders the *modification order* (**mo**).

With our event structure model, a natural question arises: should we enforce coherence at the level of event structures or only at the level of extracted executions? In the following, we will consider both options. We call **WEAKEST** the plain model that checks coherence only at the execution level, while **WEAKESTM0** the model that records a modification order in its event structures and discards any event structures containing an incoherent execution.

To make this concrete, consider the following standard coherence litmus test program, whose annotated outcome should be forbidden.

$$\begin{array}{l} X = 1; \\ a = X; // 2 \end{array} \parallel \begin{array}{l} X = 2; \\ b = X; // 1 \end{array} \quad (\text{Coh})$$

At the execution level, the outcome $a = 2 \wedge b = 1$ is forbidden because **mo** must order two writes to X . Suppose without loss of generality that $\text{St}(X, 1)$ is ordered before $\text{St}(X, 2)$. Then, returning $b = 1$ is inconsistent because it reads from a write that is overwritten by a write before the read.

At the event structure level, **WEAKEST** allows the following event structure:

This, on its own, is not a problem because we cannot extract from it a consistent execution, and so **WEAKEST** forbids the incoherent outcome.

¹The exceptions are Itanium and Java, which enforce a slightly weaker coherence property for their plain accesses.

Fig. 3. A program and its WEAKEST event structure yielding the outcome $a = 3 \wedge b = 2 \wedge c = 1$.

WEAKESTMO goes further and disallows the incoherent event structure as well. Similar to executions, WEAKESTMO event structures have a modification order, **mo**. The model requires **mo** to totally order non-conflicting writes to the same location, and declares event structures containing coherence violations as inconsistent.

We move on to a more complex example that distinguishes the two models. Consider the program shown in Fig. 3 and the outcome $a = 3 \wedge b = 2 \wedge c = 1$. We note that this outcome is rather dubious: one can see it as arising due to a violation of coherence or a circular dependency. Indeed, if $b = 2$ then $X = 1$ must be **mo**-before $X = 2$, and so the first thread cannot read $a = 1$. Likewise, it should not be able to read $a = 3$ because that depends on $c = Y$ reading 1, which circularly depends on the first thread reading $a = 3$. We further note that this outcome is not observable on any machine and the program cannot be transformed in any reasonable fashion so as to enable that outcome. (In particular, read-after-write elimination ruins the outcome, and the only reordering possible is moving the $c = Y$ instruction earlier, which does not enable any further optimizations.)

This outcome is, however, allowed by the WEAKEST model because of the event structure shown in Fig. 3. In contrast, it is not allowed by WEAKESTMO because the displayed event structure is incoherent. Recall that under WEAKESTMO, non-conflicting stores will have to be ordered by **mo**. If $\text{St}(X, 2)$ is ordered before $\text{St}(X, 1)$, then the $\text{Ld}(X, 2)$ load violates coherence. If, instead, $\text{St}(X, 2)$ is ordered after $\text{St}(X, 1)$, then the $\text{Ld}(X, 1)$ load of the first thread is incoherent.

As we will shortly see, this outcome is also allowed by the promising semantics (PS) [Kang et al. 2017]. In fact, we will show that WEAKEST can indeed simulate PS, which allows us to leverage the existing results about PS. In contrast, since the WEAKESTMO model fails to simulate PS, we will separately establish its compilation correctness.

2.4 Relating WEAKEST and WEAKESTMO to the Promising Semantics

We start with a simplified informal presentation of PS sufficient for programs containing only relaxed loads and stores. (We refer the reader to Kang et al. [2017] for further details.) PS is defined operationally by a machine, whose state consists of the states of the threads and the shared memory. The shared *memory* is modeled as a set of messages, representing the writes that the program has performed. Messages are of the form $\langle x : v@t \rangle$ denoting a write of value v to location x at timestamp t . Timestamps enforce coherence by totally ordering the writes to each location.

Each thread maintains a *view*, V , mapping each memory location to a timestamp, indicating the message with the maximum timestamp that the thread is aware of. When a thread performs a store $\text{St}(x, v)$, it chooses an unused timestamp $t > V(x)$, adds the message $\langle x : v@t \rangle$ to memory, and updates its thread view of x to timestamp t . When it performs a load of x , it reads from a message $\langle x : v@t \rangle$ with $t \geq V(x)$ and updates its thread view of x to timestamp t .

Fig. 4. A promise machine execution of LB with its corresponding event structure and extracted execution.

Besides executing its next instruction, a thread can also *promise* to write value v at some location x at some future timestamp $t > V(x)$. It does so by appending the message $\langle x : v@t \rangle$ to memory as if the write were actually performed, and keeping track of the message as being an outstanding promise in its local state. To make such a promise, the respective thread must be able to *certify* its promise: namely, to be able to perform a write fulfilling the promise in a finite number of thread-local steps. More generally, PS requires that after each execution step all outstanding promises be certifiable. When a write step later fulfills an outstanding promise, it marks it as fulfilled.

Example 1 (Load buffering). Figure 4a displays an execution of the promise machine for LB resulting in the outcome $a = b = 1$. The steps are as follows.

- (\$1, \$2) Initially, X and Y are zero and the memory is $\{\langle X : 0@0 \rangle, \langle Y : 0@0 \rangle\}$.
- (\$3) Thread T_1 promises to write $Y = 1$ at timestamp 2 and appends a message $\langle Y : 1@2 \rangle$ to the memory. The certification steps read from message $\langle X : 0@0 \rangle$ ($\$C_1$) and fulfill the promise ($\$C_2$).
- (\$4) T_2 executes $b = Y$; and reads from message $\langle Y : 1@2 \rangle$.
- (\$5) T_2 executes $X = 1$; adding message $\langle X : 1@1 \rangle$ to the memory.
- (\$6) T_1 executes $a = X$; reading message $\langle X : 1@1 \rangle$. Note that T_1 can still fulfill its promise by $\$C_3$.
- (\$7) Finally, T_1 executes the $Y = 1$; statement and fulfills its promise. $\square$

We now demonstrate on the load buffering example how we can simulate the promise machine execution with our WEAKEST model. Figure 4b displays the corresponding event structure following the steps of the promise machine execution.

- (\$1, \$2) We add the two initialization stores $\text{St}(X, 0)$ and $\text{St}(Y, 0)$ at the beginning of the graph.
- (\$3) For a promise step, we cannot immediately add a corresponding store to the event structure. We therefore look at its promise-free certificate execution (i.e., steps $\$C_1$ and $\$C_2$) and add the corresponding events to the structure. This adds events $\text{Ld}(X, 0)$ and $\text{St}(Y, 1)$.

$T_0: \text{St}(X, 0)@0$	\$1
$T_0: \text{St}(Y, 0)@0$	\$2
$T_2: \text{St}(X, 1)@2$	\$3
$T_1: \text{promise}\langle Y : 1@5 \rangle$	\$4
$T_1: \text{St}(X, 2)@1$	$\$C_1$
$T_1: \text{Ld}(X, 1)@2$	$\$C_2$
$T_1: \text{St}(Y, 1)@5$	$\$C_3$
$T_2: \text{promise}\langle X : 3@10 \rangle$	\$5
$T_2: \text{Ld}(X, 1)@2$	$\$C_4$
$T_2: \text{Ld}(Y, 1)@5$	$\$C_5$
$T_2: \text{St}(X, 3)@10$	$\$C_6$
$T_1: \text{St}(X, 2)@7$	\$6
$T_2: \text{Ld}(X, 2)@7$	\$7
$T_2: \text{Ld}(Y, 1)@5$	\$8
$T_1: \text{Ld}(X, 3)@10$	\$9
$T_2: \text{fulfill}\langle X : 3@10 \rangle$	\$10
$T_1: \text{fulfill}\langle Y : 1@5 \rangle$	\$11

(a) Promise execution steps

(b) Constructed WEAKEST event structure

(c) Extracted execution

Fig. 5. A promise execution of **Coh-CYC** with its corresponding WEAKEST event structure and execution.

($\$4$) In the promise machine, this step reads from the promised message $\langle Y : 1@2 \rangle$. Recall that the corresponding event in the event structure (i.e., $\$C_2$) was created by the certificate run. So we append a $\text{Ld}(Y, 1)$ event justifying from $\$C_2$.

($\$5$) We simply append $\text{St}(X, 1)$ to the event structure.

($\$6$) We construct event $\text{Ld}(X, 1)$ justifying from $\$5$ (i.e., the event corresponding to message $\langle X : 1@1 \rangle$), which is appended immediately after $\text{St}(Y, 0)$ in the first thread. As $\text{St}(Y, 0)$ already has an immediate po-following event in the first thread, $\text{Ld}(X, 1)$ is in conflict with $\text{Ld}(X, 0)$.

Note that at this point, the newly added event $\$6$ is *invisible* because of the conflicting store $\$C_2$ that is justified from thread T_2 . To make $\$6$ visible, we look at T_1 's certification run and add the corresponding events to the event structure (i.e., event $\$C_3$). Since the certificate must fulfill all outstanding promises, it must contain a write producing the same message as that of step $\$C_2$ (namely, the write $\$C_3$). We relate the two writes by *ew*, which makes events $\$6$ and $\$C_3$ visible.

($\$7$) At this step, we would normally append a $\text{St}(Y, 1)$ event immediately after event $\$6$, but we notice that such an event already exists because of the previous step. We therefore do nothing and simply record that $\$7$ corresponds to the same event as $\$C_3$.

By selecting the non-promised events, we can then extract the execution shown in Fig. 4c. The constraints about the messages' timestamps ensure that the execution is consistent.

In summary, we have shown that WEAKEST is strictly weaker than both WEAKESTM0 and PS, while WEAKESTM0 and PS are incomparable. Although our models further correct the aforementioned

counterexample of PS's compilation to ARMv8, we do not have proof of soundness of compilation from WEAKESTMO to ARMv8.

3 FORMALIZATION: THE WEAKEST AND WEAKESTMO MODELS

In this section, we formally define the WEAKEST and WEAKESTMO models that were introduced in the previous section.

We take programs, $\mathbb{P}$, to be top-level parallel compositions of some number of threads, $T_1 \parallel \dots \parallel T_N$. To make our semantics parametric over the programming language syntax, we represent each thread T_i as some non-empty set of traces denoting the possible sequences of memory accesses it can produce. Formally, a trace, τ , is a (finite) sequence of labels given by the following grammar:

$$lab ::= Ld_{o_r}(x, v) \mid St_{o_w}(x, v) \mid U_{o_u}(x, v, v') \mid F_{o_u} \quad (\text{Event labels})$$

$$o_r ::= NA \mid RLX \mid ACQ \mid SC \quad (\text{Memory orders for reads})$$

$$o_w ::= NA \mid RLX \mid REL \mid SC \quad (\text{Memory orders for writes})$$

$$o_u ::= RLX \mid ACQ \mid REL \mid ACQ-REL \mid SC \quad (\text{Memory orders for updates and fences})$$

We have load (Ld), store (St), update (U), and fence (F) labels. Updates are used to model the effect of atomic read-modify-write (RMW) instructions, such as fetch-and-add and compare-and-swap (CAS). The labels record the location accessed (x), the values read and/or written (v, v'), and the corresponding C11 *memory order* ($o_r/o_w/o_u$): non-atomic (NA), relaxed (RLX), acquire (ACQ), release (REL), acquire-release (ACQ-REL), or sequentially consistent (SC). In increasing strength, these orders are: $NA \sqsubset RLX \sqsubset \{ACQ, REL\} \sqsubset ACQ-REL \sqsubset SC$.

We assume that the thread semantics is *prefix-closed*, *receptive*, and *deterministic*. Prefix-closedness requires for every $\tau \cdot \tau' \in T_i$ (where $\cdot$ denotes trace concatenation), we have $\tau \in T_i$. In particular, this means that the empty trace is always included in T_i . Receptiveness requires that whenever $\tau \cdot lab \in T_i$ and lab is a read label (a load or an update) of location x , then for every value v , there exists a read label lab' of location x reading that value v such that $\tau \cdot lab' \in T_i$. In other words, whenever a thread can read x , it can do so for any possible value. Determinism requires that whenever $\tau \cdot lab \in T_i$ and $\tau \cdot lab' \in T_i$, then the two labels agree except perhaps due to receptiveness. That is, if unequal, then they are both reads (loads or updates) of the same location.

Definition 1. An *event* is a tuple, $\langle id, tid, lab \rangle$, where $id \in \mathbb{N}$ is a unique identifier for the event, $tid \in \mathbb{N}$ identifies the thread to which the event belongs (we use $tid = 0$ for initialization events), and lab denotes the label of the event.

A *read event* is either a load or an update, whereas a *write event* is either a store or an update. Let $\mathcal{E}$ denote the set of all events, $\mathcal{R}$ the set of all read events, $\mathcal{W}$ the set of all write events, and $\mathcal{F}$ the set of fence events. We use subscripts to constraint those sets; e.g., $\mathcal{E}_{\sqsupseteq REL}$ denotes all events whose memory order is at least REL. We write $e.id$, $e.tid$, $e.lab$, $e.op$, $e.loc$, $e.ord$, $e.rval$, and $e.wval$, to return (when applicable) the identifier, thread, label, type (Ld, St, U, F), location, memory order, read value, and written value respectively.

Relational Notation. Let $R, S \subseteq \mathcal{E} \times \mathcal{E}$ be binary relations on events. We write $R ; S$ for the relational composition of R and S . Let R^{-1} denote the inverse of R , $\text{dom}(R)$ its domain and $\text{codom}(R)$ its range. We write $R^?$, R^+ , R^* , $R^=$ for the reflexive, the transitive, the reflexive-transitive, and the reflexive-symmetric closures of R respectively. (Reflexive closure is with respect to the set $\mathcal{E}$, while reflexive-symmetric closure means $R^= \triangleq (R \cup R^{-1})^?$.) Further, $R|_{\text{loc}} \triangleq \{(e, e') \in R \mid e.loc = e'.loc\}$ restricts R on pairs of events of the same location. Similarly, $R|_{\neq \text{loc}} \triangleq R \setminus R|_{\text{loc}}$ restricts R on pairs of events of different locations.

The $[A]$ notation denotes an identity relation on set A , i.e. $[A](x, y) \triangleq x = y \wedge x \in A$. For a finite set A that is totally ordered by R , we write $\text{sequence}_R(A)$ to denote the sequence of its elements ordered according to R .

Definition 2. An event structure, $G \triangleq \langle E, \text{po}, \text{jf}, \text{ew}, \text{mo} \rangle$, contains the following components:

- $E \subseteq \mathcal{E}$ is the *set of events* of the event structure containing the set of initialization events, E_0 , which has exactly one event with label $\text{St}_{\text{NA}}(x, 0)$ for every location x . For a thread i , we write E_i to refer to the events of the event structure with thread identifier i .
- $\text{po} \subseteq E \times E$ denotes the *program order*. It is a strict partial order on events recording when an event precedes another one in the control flow of the program. Initialization events are po-before all non-initialization events (i.e., $E_0 \times (E \setminus E_0) \subseteq \text{po}$), and that non-initialization events related by po have the same thread identifier.

Further, we assume that non-initialization po-predecessors of an event are totally ordered by po. That is, we never have distinct $a, b \in E \setminus E_0$ that have a common po-successor and are not po-related to one another. In relational notation, $[E \setminus E_0]; \text{po}; \text{po}^{-1}; [E \setminus E_0] \subseteq \text{po}^=$.

We note that po is not total over events of the same thread. In particular, we say that events of the same thread not related by the program order are in *conflict* with one another.

$$G.\text{cf} \triangleq \{(a, b) \in G.E \times G.E \mid a.\text{tid} = b.\text{tid} \neq 0\} \setminus \text{po}^= \quad (\text{Conflict relation})$$

By definition, the conflict relation, cf, is symmetric and irreflexive. Moreover, it is forward-closed with respect to the program order (i.e., $\text{cf} ; \text{po} \subseteq \text{cf}$): if two events conflict, then so are all their future events.

- $\text{jf} \subseteq E \times E$ is the *justified-from* relation, which relates a write event to the reads it justifies. That is, whenever $\text{jf}(w, r)$, then $w \in \mathcal{W}$, $r \in \mathcal{R}$, $w.\text{loc} = r.\text{loc}$ and $w.\text{wval} = r.\text{rval}$ hold. We require that every read is justified from exactly one write. That is, jf^{-1} is functional and $\text{dom}(\text{jf}) = E \cap \mathcal{R}$.
- $\text{ew} \subseteq [\mathcal{W}_{\sqsubseteq \text{RLX}}]; \text{cf}|_{\text{loc}}; [\mathcal{W}_{\sqsubseteq \text{RLX}}]$ is the *equal-writes* relation, a symmetric relation relating conflicting (relaxed) writes on the same location writing the same value.
- $\text{mo} \subseteq ((E \cap \mathcal{W}) \times (E \cap \mathcal{W}))|_{\text{loc}}$ is the *modification order*: a strict partial order that orders write operations on the same memory location. We require that mo is total on non-conflicting writes to the same location and that equal writes have the same mo -successors, i.e., $G.\text{ew}; G.\text{mo} \subseteq \text{mo}$. (In the WEAKEST model, we do not use the mo component in the event structure.)

Given an event structure G , we use notation $G.X$ to project the X component out of G . When G is clear from the context, we sometimes omit the “ G ”.

We now define a number of derived relations on event structures. First, we define *synchronizes-with* (sw) and *happens-before* (hb) following RC11 [Lahav et al. 2017] (replacing C11’s reads-from relation with justification). The sw definition is quite involved and states conditions under which a justification edge synchronizes two threads. The simplest case inducing synchronization is when an acquire read event reads from a release write. More advanced cases involve release/acquire fences and/or reads through sequences of updates (known as release sequences in C11). An event a *happens before* an event b if there is a path from a to b consisting of program order and synchronization edges.

$$G.\text{sw} \triangleq [\mathcal{E}_{\sqsubseteq \text{REL}}]; ([\mathcal{F}]; G.\text{po})^?; [\mathcal{W}]; G.\text{po}|_{\text{loc}}^?; [\mathcal{W}_{\sqsubseteq \text{RLX}}]; G.\text{jf}^+; [\mathcal{R}_{\sqsubseteq \text{RLX}}]; (G.\text{po}; [\mathcal{F}])^?; [\mathcal{E}_{\sqsubseteq \text{ACQ}}] \quad (\text{Synchronizes-with})$$

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

We say that two events are in *immediate conflict* ($\sim$) if neither of them has a po-previous event conflicting with the other event. Two events are in *extended conflict* (ecf) if they happen after

conflicting events, which means that they cannot be part of the same execution.

$$G.\sim \triangleq G.cf \setminus (G.cf ; G.po \cup G.po^{-1} ; G.cf) \quad (\text{Immediate conflict})$$

$$G.ecf \triangleq (G.hb^{-1})^? ; G.cf ; G.hb^? \quad (\text{Extended conflict})$$

Next, the *reads-from* relation, $\mathbf{rf}$, lifts the justified-from relation to all equal writes; i.e., whenever r is justified by w and w is equal to w' and w' and r do not conflict, then $\mathbf{rf}(w', r)$. This, in particular, means that $\mathbf{rf}^{-1}$ is not necessarily functional at the level of event structures.

$$G.\mathbf{rf} \triangleq (G.\mathbf{ew}^? ; G.\mathbf{jf}) \setminus G.cf \quad (\text{Reads-from relation})$$

The next two relations—*reads-before/from-read* ($\mathbf{fr}$) and the *extended coherence order* ($\mathbf{eco}$)—are also defined as in RC11 [Lahav et al. 2017]. A read event r reads before a write event w if r reads from some write w' that is $\mathbf{mo}$ -after w . (In the formal definition, we subtract the identity relation so as to avoid saying that an update reads before itself.) Finally, the extended coherence order is the transitive closure of the union of the three coherence-enforcing relations ($\mathbf{rf}$, $\mathbf{mo}$, and $\mathbf{fr}$).

$$G.\mathbf{fr} \triangleq (G.\mathbf{rf}^{-1} ; G.\mathbf{mo}) \setminus [\mathcal{E}] \quad (\text{Reads-before/from-read relation})$$

$$G.\mathbf{eco} \triangleq (G.\mathbf{rf} \cup G.\mathbf{mo} \cup G.\mathbf{fr})^+ \quad (\text{Extended coherence order})$$

Since the **WEAKEST** model does not record the modification order, we define $\mathbf{mo}_{\text{strong}}$, a stronger version of the modification order, which relates writes to the same location that are ordered by happens-before. We then use $\mathbf{mo}_{\text{strong}}$ in place of $\mathbf{mo}$ to derive stronger versions of $\mathbf{fr}$ and $\mathbf{eco}$.

$$G.\mathbf{mo}_{\text{strong}} \triangleq [\mathcal{W}] ; G.hb|_{\text{loc}} ; [\mathcal{W}] \quad (\text{Strong modification order})$$

$$G.\mathbf{fr}_{\text{strong}} \triangleq (G.\mathbf{rf}^{-1} ; G.\mathbf{mo}_{\text{strong}}) \setminus [\mathcal{E}] \quad (\text{Strong reads-before})$$

$$G.\mathbf{eco}_{\text{strong}} \triangleq (G.\mathbf{rf} \cup G.\mathbf{mo}_{\text{strong}} \cup G.\mathbf{fr}_{\text{strong}})^+ \quad (\text{Strong extended coherence order})$$

Remark 1. Note that the definitions above allow $\mathbf{ew}$ to relate a relaxed store with a conflicting relaxed update writing the same value to the same location. This flexibility is not needed for any of our proofs about **WEAKESTM0**; in fact, it would suffice for $\mathbf{ew}$ to relate only events with the same label and, in the case of updates, with the same justification ($\mathbf{jf}$). Under these restrictions, then $\mathbf{sw}$ as defined above is equivalent to the C11-definition of $\mathbf{sw}$ that uses $\mathbf{rf}$ in the place of $\mathbf{jf}$. Without these restrictions, however, the definition with $\mathbf{rf}$ is problematic because adding a new update event and making it $\mathbf{ew}$ -related to an existing store could induce synchronizations between existing events in the event structure. The only reason why we did not restrict $\mathbf{ew}$ is so that we can establish that **WEAKEST** is weaker than the promising semantics in §4. The promising semantics, in particular, allows certifying a promised message with an update event and then fulfilling it with a write, and vice versa.

3.1 Event Structure Consistency Checking

We say that an event e is visible in an event structure G if all the writes recursively used to justify e that are in conflict with e have some equal write that does not conflict with e . Formally:

Definition 3. An event e is *visible* in an event structure G if $e \in G.E$ and

$$[\mathcal{W}] ; (G.cf \cap G.\mathbf{jfe} ; (G.po \cup G.\mathbf{jf})^* ; G.\mathbf{jfe} ; G.po^?) ; [\{e\}] \subseteq G.\mathbf{ew} ; G.po^=$$

where $G.\mathbf{jfe} \triangleq G.\mathbf{jf} \setminus G.po$ denotes the *external justification edges*. We write $\text{vis}(G)$ for the set of all visible events of an event structure G .

As we will shortly see, our model essentially treats extended conflicts as normal conflicts.

Fig. 7. Constraint (CF) restricts $Z = 2$ in the **Cwrites** program.

Definition 4. An event structure G is *consistent* according to model $M \in \{\text{WEAKEST}, \text{WEAKESTMO}\}$, written $\text{isCons}_M(G)$, iff the following conditions hold:

- (CF) No event can be in extended conflict with itself: $G.\text{ecf}$ is irreflexive.
- (CFJ) A write cannot justify a read in extended conflict: $G.\text{jf} \cap G.\text{ecf} = \emptyset$.
- (VISJ) Only visible events justify reads of other threads: $\text{dom}(G.\text{jfe}) \subseteq \text{vis}(G)$.
- (ICF) Immediately conflicting events must be reads: $G.\sim \subseteq \mathcal{R} \times \mathcal{R}$.
- (ICFJ) Immediately conflicting reads cannot be justified by the same or by equal writes:
 $G.\text{jf} ; G.\sim ; G.\text{jf}^{-1} ; G.\text{ew}^?$ is irreflexive.
- (COH) if $M = \text{WEAKEST}$, then $G.\text{hb} ; G.\text{eco}_{\text{strong}}^?$ is irreflexive.
- (COH') if $M = \text{WEAKESTMO}$, then $G.\text{hb} ; G.\text{eco}^?$ is irreflexive.

We now explain these constraints in order.

The first constraint (CF) ensures that every event in the event structure could belong to some execution by checking that its **hb**-predecessors are conflict-free. This, in particular, ensures that po-related reads do not observe values from two different execution branches of some other thread in the case where both reads-from edges would result in a synchronization. We explain the scenario with the **Cwrites** program in Fig. 7 taken from [Chakraborty and Vafeiadis \[2017\]](#). Intuitively, this program should never produce the outcome $Z = 2$, as there is no execution under which both X and Y could have the value 1. With this in mind, we would like to rule out the event structure shown in Fig. 7 that contains an execution branch where both acquire loads of X and Y read the value 1. It is easy to see that the (CF) does not hold in the event structure, as $\text{St}_{\text{REL}}(Y, 1)$ conflicts with $\text{St}_{\text{REL}}(X, 1)$ which happens before $\text{Ld}_{\text{ACQ}}(Y, 1)$.

The second constraint (CFJ) forbids a read to be justified from a write event in extended conflict. Intuitively, since conflicting events can never appear in a single execution, allowing reads to read from conflicting events could generate OOTA results. This constraint rules out the problematic event structure for the variants of **Cwrites** where one of the reads of the third thread is relaxed.

Remark 2. While (CF) rules out the problematic event structure of Fig. 7, it is not strictly needed for rule out the $Z = 2$ behavior, because no RC11-consistent execution can be extracted from that event structure. Nevertheless, (CF) and (CFJ) are needed for a more complicated variant of **Cwrites** in [Kang et al. \[2017, Appendix A.2\]](#), where without them we cannot guarantee DRF-RA (Theorem 4).

The third constraint (VISJ) says that reads are either justified locally from the same thread or they are justified from some *visible* write. This constraint forbids the undesired event structure in Fig. 2 for the **RNG** program.

The next two constraints place some restrictions on immediate conflicts that rule out creating unnecessary duplication in the event structure. (ICF) requires that immediate conflicts are between read events. Because we assume that the thread semantics is deterministic, immediately conflicting events must be created by the same program instruction. If, for example, they originate from a

Fig. 8. Outcome $a = b = 1$ is not allowed in **LBra** because of a coherence violation.

load instruction, they will all be load events. If they originate for a compare-and-swap (CAS), they will either be update events (for the successful case) or load events (for the case when the CAS fails). In both of these cases, it may well make sense to have an immediate conflict in the event structure. In contrast, it does not make sense to create an immediate conflict from a store or a fence instruction, because both events will have the exact same label, which is why (ICF) rules out such events. Similarly, (ICFJ) disallows creating an immediate conflict by reading from the same write. That is, we cannot have two reads in immediate conflict that are justified by the same write or by equal writes. Again such reads will read the same value, and so will lead to duplication in the event structure.

Finally, the (COH) and (COH') constraints enforce that the event structure is coherent at various levels. Both constraints ensure that the happens-before (**hb**) is irreflexive and that **rf** does not contradict **hb**. We explain the latter with the program in Fig. 8, which is another variant of the load buffering litmus test. In this case, the outcome $a = b = 1$ must be forbidden and is actually disallowed by C/C++11 (despite thin-air problems). The reason is that if $a = 1$, then the release write of X synchronizes with the acquire read of X , and so the read of Y would have to read from an **hb**-later write. Our coherence constraints similarly enforce that the event structure in Fig. 8 is inconsistent. More specifically, the event structure without the second $\text{St}_{\text{RLX}}(Y, 1)$ event (i.e., the one displayed in red) is consistent but adding that event along with the displayed **ew** relation renders the event structure inconsistent. The equal-writes relation together with Y 's justification edge induce the displayed reads-from edge, which contradicts happens-before.

In addition, the coherence constraints also ensure that reads do not read overwritten values, as for example in the following inconsistent event structure:

$$a : \text{St}(X, 1) \longrightarrow b : \text{St}(X, 2) \longrightarrow c : \text{Ld}(X, 1) \quad (\text{Basic coherence violation})$$

The justification edge induces an **rf**-edge from a to c , and in turn an **fr_{strong}**-edge from c to b , which contradicts **hb**.

The two coherence constraints differ only in the presence of concurrent writes to the same location; i.e., of two writes to a given location that are not ordered by **hb**. If, for example, in the event structure above, the events a and b were not related by **po**, **WEAKEST** would allow the justification edge. **WEAKESTM0** would also allow it but only if the **mo**-edge went from b to a . If **mo** went from a to b , as depicted below, the event structure would be inconsistent:

$$a : \text{St}(X, 1) \xrightarrow{\text{mo}} b : \text{St}(X, 2) \longrightarrow c : \text{Ld}(X, 1) \quad (\text{WEAKESTM0 coherence violation})$$

As already explained in §2.3, the difference is evident in programs **Coh** and **Coh-CYC**. Specifically, the event structure of **Coh** shown in that section (**Coh-ES**) is consistent according to the **WEAKEST** model, but not according to **WEAKESTM0**: whichever way $\text{St}(X, 1)$ and $\text{St}(X, 2)$ are ordered by

$$\begin{array}{c}
A \subseteq G.E_{e.\text{tid}} \quad \text{dom}([E_{e.\text{tid}}]; \text{po}; [A]) \subseteq A \quad \text{labels}(\text{sequence}_{\text{po}}(A)) \cdot e.\text{lab} \in \mathbb{P}(e.\text{tid}) \\
E' = E \uplus \{e\} \quad \text{po}' = \text{po} \cup (A \times \{e\}) \quad \text{isCons}_M(\langle E', \text{po}', \text{jf}', \text{ew}', \text{mo}' \rangle) \quad CF = (E_{e.\text{tid}} \setminus A) \\
\text{if } e \in \mathcal{R} \text{ then } \exists w \in E \cap \mathcal{W}. \text{jf}' = \text{jf} \cup \{(w, e)\} \wedge w.\text{loc} = e.\text{loc} \wedge w.\text{wval} = e.\text{rval} \text{ else } \text{jf}' = \text{jf} \\
\quad \text{if } e \in \mathcal{W}_{\text{RLX}} \text{ then AddEW}(\text{ew}, \text{ew}', CF, e) \text{ else } \text{ew}' = \text{ew} \\
\quad \text{if } e \in \mathcal{W} \text{ then AddMO}(\text{mo}, \text{mo}', E, CF, \text{ew}, e) \text{ else } \text{mo}' = \text{mo} \\
\hline
\langle E, \text{po}, \text{jf}, \text{ew}, \text{mo} \rangle \rightarrow_{\mathbb{P}, M} \langle E', \text{po}', \text{jf}', \text{ew}', \text{mo}' \rangle
\end{array}$$

Fig. 9. One construction step of a program's event structure, where $\text{labels}(a_1 \cdots a_n) \triangleq a_1.\text{lab} \cdots a_n.\text{lab}$.

mo, there is a coherence violation. The event structure of **Coh-CYC** shown in Fig. 5b is similarly inconsistent according to **WEAKESTMO**, because it contains the incoherent **Coh-ES** as a substructure.

3.2 Event Structure Construction

Having defined event structures and their consistency, we move on to explain how they are constructed.

We start with the initial event structure G_{init} that contains as events only the initialization events, $E_0 = \{a_1, \dots, a_k\}$ where k is the number of variables in the program and each a_i has label $\text{St}_{\text{NA}}(X_i, 0)$ and thread identifier 0. All the other components of G_{init} are simply the empty relation.

Starting with G_{init} , we construct an event structure of a program $\mathbb{P}$ incrementally by adding one event at a time using the rule shown in Fig. 9. The reduction relation has the form $G \rightarrow_{\mathbb{P}, M} G'$ and represents a transition from event structure G to event structure G' of the program $\mathbb{P}$ according to model $M \in \{\text{WEAKEST}, \text{WEAKESTMO}\}$.

The rule adds one new event e to the event structure G . To do so, it picks the set A of its po-predecessors from the same thread. This set A includes events of thread $e.\text{tid}$ that are already in the event structure ($E_{e.\text{tid}}$), and is closed under program order prefix. That is, any po-predecessor of an event in A should also belong to A . Next, it checks that local thread semantics can produce the label of e , after the sequence of labels of A .

In the second line of the rule, the set of events is extended to include e , and the program order is similarly extended to place e after all the events in A . It checks that the resulting event structure is consistent according to model M where the updates to the event structure's remaining components will be defined shortly, and discards any inconsistent event structures. We finally let CF contain all the events conflicting with e , which are exactly those events from the same thread as e that were not placed before it.

The third line of the rule concerns the update of the justification relation. If e is a read event, then there must already be some write event w in the event structure that writes the value that e reads to the same location. In that case, the justification relation is extended with the edge (w, e) . If e is not a read (i.e., it is a store or a fence event), then **jf** remains unchanged.

The last two lines concern the update of the equal-writes relation and the modification order. The former is updated whenever the new event is a relaxed write, whereas the latter is updated for all new write events. We use the following auxiliary definitions for updating **ew** and **mo**:

$$\begin{aligned}
\text{AddEW}(\text{ew}, \text{ew}', CF, e) &\triangleq \exists W \subseteq \mathcal{W}_{\text{RLX}} \cap CF. (\forall w \in W. w.\text{loc} = e.\text{loc} \wedge w.\text{wval} = e.\text{wval}) \wedge \\
&\quad \text{ew}' = \text{ew} \cup (W \times \{e\}) \cup (\{e\} \times W) \\
\text{AddMO}(\text{mo}, \text{mo}', E, CF, \text{ew}, e) &\triangleq \exists w \in \mathcal{W} \cap E \setminus CF. w.\text{loc} = e.\text{loc} \wedge \\
&\quad \text{mo}' = \text{mo} \cup (\text{dom}(\text{mo}^?; \text{ew}^?; [\{w\}]) \times \{e\}) \\
&\quad \cup (\{e\} \times \text{codom}([\{w\}]; \text{mo}; \text{ew}^?))
\end{aligned}$$

$$\begin{aligned}
X.\text{sw}_{C11} &\triangleq [\mathcal{E}_{\sqsupset_{\text{REL}}}; ([\mathcal{F}]; X.\text{po})^?; [\mathcal{W}]; X.\text{po}|_{\text{loc}}^?; [\mathcal{W}_{\sqsupset_{\text{RLX}}}; X.\text{rf}^+; [\mathcal{R}_{\sqsupset_{\text{RLX}}}; (X.\text{po}; [\mathcal{F}])^?; [\mathcal{E}_{\sqsupset_{\text{ACQ}}}] \\
X.\text{hb}_{C11} &\triangleq (X.\text{po} \cup X.\text{sw}_{C11})^+ \\
X.\text{scb} &\triangleq X.\text{po} \cup (X.\text{po}|_{\neq \text{loc}}; X.\text{hb}_{C11}; X.\text{po}|_{\neq \text{loc}}) \cup X.\text{hb}_{C11}|_{\text{loc}} \cup X.\text{mo} \cup X.\text{fr} \\
X.\text{psc}_{\text{base}} &\triangleq [\mathcal{E}_{\text{sc}}]; ([\mathcal{F}_{\text{sc}}]; X.\text{hb}_{C11})^?; X.\text{scb}; (X.\text{hb}_{C11}; [\mathcal{F}_{\text{sc}}])^?; [\mathcal{E}_{\text{sc}}] \\
X.\text{psc}_{\text{F}} &\triangleq [\mathcal{F}_{\text{sc}}]; (X.\text{hb}_{C11} \cup X.\text{hb}_{C11}; X.\text{eco}; X.\text{hb}_{C11}); [\mathcal{F}_{\text{sc}}]
\end{aligned}$$

Fig. 10. Additional definitions needed for execution consistency (taken from Lahav et al. [2017]).

AddEW selects a set W of conflicting relaxed writes to the same location and that write the same value as e . Then, ew is extended by relating e to all the elements in W . AddMO selects a write event w in the graph that writes to the same location as e and does not conflict with it. Since the event structure is initialized, such a write always exists. It then puts e immediately after w in mo .

3.3 Execution Extraction in the WEAKEST and WEAKESTM Models

As mentioned in §2.1, after a consistent event structure is constructed, we extract from it a set of conflict-free subsets of events, which we call executions.

More formally, an execution X is a tuple of the form $\langle E, \text{po}, \text{rf}, \text{mo} \rangle$ whose components in order denote the set of events, the program order, the reads-from relation, and the modification order of the execution. As with event structures, we use the dot notation (e.g., $X.E$) to project the relevant components of an execution.

The set of events, the program order, and the reads-from relation have the same constraints as for event structures. For example, when event r reads from event w , then r is a read, w is a write, they both access the same location, and the value written by w is read by r . Unlike event structures, however, executions do have the equal write (ew) relation. (Since an execution has no conflicting events, then ew would necessarily be empty.) Similarly, the modification order of executions ($X.\text{mo}$) is total on writes to the same location.

In order to define execution consistency, we need some additional definitions, which are taken straight from Lahav et al. [2017] and are presented in Fig. 10. Here, *synchronizes-with* is defined as in $C11$ referring to the reads-from relation rather than jf , which does not exist at the level of executions. The last three definitions (scb , psc_{base} , and psc_{F}) are quite technical and concern cases when an ordering between SC events is guaranteed to be preserved by the memory model. They are not important for our examples, and are only included here for completeness.

Next, we define when an execution is consistent. As mentioned, we take the RC11 consistency constraints [Lahav et al. 2017], except that we allow $(X.\text{po} \cup X.\text{rf})^+$ -cycles. The reason for allowing cycles is that we do not rely on RC11 consistency to rule out OOTA behaviors, but rather on the construction of event structures. One further (minor) difference is that we model updates as single events, whereas Lahav et al. [2017] model them as two events (a read and a write) connected by a special read-modify-write relation. As a result, our atomicity condition is slightly different.

Definition 5. An execution X is *consistent*, denoted $\text{isCons}(X)$, if the following hold:

- (Total-RF) $\text{codom}(X.\text{rf}) = X.E \cap \mathcal{R}$;
- (Total-MO) $X.\text{mo} = ((X.E \cap \mathcal{W}) \times (X.E \cap \mathcal{W}))|_{\text{loc}}$;
- (Coherence) $X.\text{hb}_{C11}; X.\text{eco}^?$ is irreflexive;
- (Atomicity) $X.\text{fr}; X.\text{mo}$ is irreflexive; and
- (SC) $X.\text{psc}_{\text{base}} \cup X.\text{psc}_{\text{F}}$ is acyclic.

where eco and fr are defined as for event structures.

The consistency constraints require that (1) every read of the execution reads from some write; (2) the modification order, **mo**, is total over all writes to the same location; (3) the execution is coherent (i.e., the execution's restriction to accesses of only one location is SC); (4) atomic updates read from their immediate **mo**-predecessor; and (5) certain cycles going through SC atomics are disallowed. Among other cases, it ensures that (1) putting SC fences between every two accesses in each thread guarantees SC, and (2) if all accesses of a consistent execution are sequentially consistent, then the execution is SC.

We move on to define how an execution is extracted from an event structure. First, we have to select an appropriate set of events, A , from the event structure so that restricting the event structure to those events gives us an execution. We place several constraints: (1) all those events should be visible; (2) they should not conflict with one another; and (3) they should be *prefix-closed* with respect to **hb**: i.e., every **hb**-predecessor of an event in A should also be in A ; and Formally,

$$\text{GoodRestriction}(G) \triangleq \{A \mid A \subseteq \text{vis}(G) \wedge [A]; G.\text{cf}; [A] = \emptyset \wedge \text{dom}(G.\text{hb}; [A]) \subseteq A\}$$

Second, given such a conflict-free set of events, $A \subseteq G.E$, we define how to restrict G to get an execution. The following formal definition, $\text{Project}_{\text{WEAKESTMO}}(G, A)$, does this by simply restricting **po**, **rf**, and **mo** to the set of events, A . $\text{Project}_{\text{WEAKEST}}(G, A)$ is defined in a similar way. Since, however, **WEAKEST** does not record the modification order in event structures, $\text{Project}_{\text{WEAKEST}}(G, A)$ generates an arbitrary **mo**.

$$\begin{aligned} \text{Project}_{\text{WEAKESTMO}}(G, A) &\triangleq \{\langle A, G.\text{po} \cap (A \times A), G.\text{rf} \cap (A \times A), G.\text{mo} \cap (A \times A) \rangle\} \\ \text{Project}_{\text{WEAKEST}}(G, A) &\triangleq \{\langle A, G.\text{po} \cap (A \times A), G.\text{rf} \cap (A \times A), \text{mo} \rangle \mid \text{mo} \subseteq (A \times A)\} \end{aligned}$$

Putting these two components together, the set of executions of event structure G are those consistent executions that are a result of projecting G to some good restriction of its events.

$$\text{ex}_M(G) \triangleq \{X \mid \exists A \in \text{GoodRestriction}(G). X \in \text{Project}_M(G, A) \wedge \text{isCons}_M(X)\}$$

Note that the executions extracted from G are not necessarily *maximal*. Rather, for any execution extracted from G , all its consistent (**po** $\cup$ **rf**)-prefixes can also be extracted from G .

3.4 Program Behaviors

The final step is to define program behaviors. We define the behavior of an execution to be the final contents of the memory at the end of an execution, i.e., the value written by the **mo**-maximal write for each location.

$$\text{Behavior}(X) \triangleq \{(e.\text{loc}, e.\text{wval}) \mid \exists e \in X.E \cap \mathcal{W}. [\{e\}]; X.\text{mo} = \emptyset\}$$

Then, the set of behaviors of a program under a model $M \in \{\text{WEAKEST}, \text{WEAKESTMO}\}$ contains the behaviors of any *maximal* execution X that can be extracted from an event structure G that was constructed from the program by following the model.

$$\text{Behavior}_M(\mathbb{P}) \triangleq \{\text{Behavior}(X) \mid \exists G. G_{\text{init}} \rightarrow_{\mathbb{P}, M}^* G \wedge X \in \text{ex}_M(G) \wedge \text{maximal}_{\mathbb{P}}(X)\}$$

where $\text{maximal}_{\mathbb{P}}(X) \triangleq \nexists i, \text{lab}. \text{labels}(\text{sequence}_{X, \text{po}}(X.E_i)) \cdot \text{lab} \in \mathbb{P}(i)$. Maximality ensures that all threads have terminated according to the thread semantics.

3.5 Java Causality Tests

We evaluate **WEAKEST** and **WEAKESTMO** on the Java causality tests [Manson et al. 2004]. Each test consists of a program along with a particular behavior and a remark as to whether that behavior should be allowed or not. For each test, we check whether our models can yield the behavior in question.

Our two models agree with each other and with PS on all the tests and with the prescribed Java behavior on 17/20 tests. Among the remaining cases, the behavior of test 16 is a classic coherence violation: while Java allows this outcome, our models forbid it. The other two cases (tests 19 and 20) are due to thread sequentialization (i.e., $\mathbb{C}_1 \parallel \mathbb{C}_2 \rightsquigarrow \mathbb{C}_1 ; \mathbb{C}_2$): Java allows the weak outcomes, whereas our models do not.

As a side note, we remark that the surprisingly weak outcomes of tests 5 and 10 (forbidden by our models) can also be explained by thread sequentialization, yet Java forbids them.

4 THE WEAKEST MODEL AND PROMISING SEMANTICS

In this section, we relate the WEAKEST model to the promising semantics (PS) of Kang et al. [2017]. We first introduce PS and then, in §4.2, we show that WEAKEST is weaker than PS.

4.1 Overview of Promising Semantics

As already discussed in §2.4, PS is defined by an abstract machine. Its state $MS = (\mathcal{TS}, \mathcal{S}, M)$ is a tuple with three components: $\mathcal{TS}$ is the *thread state map*, a function from thread identifiers to thread local states; $\mathcal{S}$ is the *global SC view*, which places a total order on the sc fences; and M is the *shared memory*, represented as a set of messages.

The thread state of thread i , $\mathcal{TS}(i)$, is again a tuple $TS = \langle \sigma, V, P \rangle$ with three components. The first component, σ , stores the local state of the thread (e.g., the value of the program counter). As our models are parametric with respect to the syntax of the programming language, we take σ to be a pair of the set of traces of the thread, $\mathbb{P}(i)$, together with the sequence of labels produced so far. The second component, V , is the *thread view*, a mapping from memory locations to the maximum timestamp of a message that the thread is aware of.² The third component, P , records the set of outstanding promises made by thread i ; i.e., the set of messages that thread i has added to the shared memory without yet having performed the corresponding writes.

As already mentioned, PS's memory, M , is just a set of messages, representing the writes (and promises) that have been performed. Each message m records the location written ($m.loc$), the value written ($m.wval$), a timestamp ($m.ts$), as well as some other components that we will not use in this section.³ For a location x , we write $\maxmsg(M, x)$ for message in M with location x that has the maximum timestamp.

The initial thread state map, $\mathcal{TS}_{init}(\mathbb{P})$, maps each thread i to the state $\langle \mathbb{P}(i), \epsilon \rangle$, the bottom view (that maps every location to the bottom timestamp), and the empty set of promises. The initial program state, $MS_{init}(\mathbb{P})$, has as components the initial thread state map, $\mathcal{TS}_{init}(\mathbb{P})$, the bottom global SC view, and the initial memory, M_{init} , that contains an initialization message $\langle x : 0 @ \perp \rangle$ for each location x .

PS supports all the operations of our programming language except for sc accesses (reads, writes, and updates). It does support sc fences. Executing an sc fence simply updates the $\mathcal{S}$ component: it computes the elementwise maximal of the global SC view ($\mathcal{S}$) and its thread-local view (V), and updates both components to be that maximal view. In other words, for each location x , it sets both $\mathcal{S}$ and V to the maximum of $\mathcal{S}(x)$ and $V(x)$.

We have already informally discussed in §2.4 how the components of a state are affected by read, write, and promise transitions. For complete descriptions of these transitions and the formal definition of the promise machine, we refer the reader to Kang et al. [2017].

²Actually, each thread has three thread views: the *current* view, the *acquire* view, and the *release* view. The release and acquire views are used to assign appropriate semantics to release and acquire fences. Further details can be found in Kang et al. [2017] and in our technical appendix. Here, for simplicity, we describe only the current view.

³The additional components are another timestamp $m.ts_{from}$, useful for reserving timestamp ranges and giving appropriate semantics to updates, and a message view.

We take the behavior of a machine state in a promise execution to be its final memory contents excluding any unfulfilled promised messages. That is, for each location x , we select the non-promised message in M that has the maximal timestamp and return its value.

$$\text{Behavior}(MS) \triangleq \{(x, v) \mid x \in \text{Locs} \wedge \max_{\text{msg}}(MS.M \setminus \bigcup_i MS.\mathcal{TS}(i).P, x).wval = v\}$$

The behavior of a program $\mathbb{P}$ is the set of behaviors of all *final* machine states that can arise by executing the program (i.e., all machine states reachable from the initial program state that cannot reduce further):

$$\text{Behavior}_{\text{PS}}(\mathbb{P}) \triangleq \{\text{Behavior}(MS) \mid MS_{\text{init}}(\mathbb{P}) \rightarrow^* MS \wedge MS \nrightarrow\}.$$

We note that recording the final memory contents suffices for distinguishing programs that differ only in their final thread-local states. We can place these programs in a context that writes the contents of the local states to main memory, and then use the definition above to distinguish them.

4.2 Relating the WEAKEST Model to Promising Semantics

We now state the main result of this section, which says that WEAKEST admits all the program behaviors that PS allows.

Theorem 1. *For a program $\mathbb{P}$, $\text{Behavior}_{\text{PS}}(\mathbb{P}) \subseteq \text{Behavior}_{\text{WEAKEST}}(\mathbb{P})$.*

To prove this theorem, consider a final promise machine state MS of the program $\mathbb{P}$. We have to show that there exists a WEAKEST event structure G of $\mathbb{P}$ with a consistent execution $X \in \text{ex}_{\text{WEAKEST}}(G)$ such that MS and X have same behavior. Our proof consists of three stages:

- (1) First, we define a simulation relation $G \sim_{\Pi} MS$ between the WEAKEST event structure G and the promise machine state MS . The relation is parameterized by a set of threads, Π , in which the event structure is lagging behind MS . The simulation relation relates various components of WEAKEST event structure to certain components of promise machine state by a number of mapping functions. We also define a set of relations to identify the desired execution in the WEAKEST event structure.
- (2) Next, we prove that steps of the promising machine *preserve* the simulation relation. More precisely, starting from related states, every step of PS can be matched by zero or more steps of the WEAKEST event structure construction yielding related states. To handle promise steps, we split this proof obligation in two lemmas. Lemma 1 considers the transitions by non-promise operations of a particular thread, while Lemma 2 takes care of state transition on entire promise machine state.
- (3) Finally, the simulation relation is defined so as to also identify an execution X of G such that the behavior of X and MS coincide. So, when PS reaches a final state, we extract that execution from G and complete the proof.

We move on to the definition of the simulation relation, $G \sim_{\Pi} MS$, which says that G simulates MS modulo the threads Π . When $\Pi = \emptyset$, we write $G \sim MS$ and say that G simulates MS . Our definition uses the following mapping functions, predicates, and relations.

- $\mathbb{W} : (G.E \cap \mathcal{W}) \rightarrow M$ is a partial function that maps certain write events of G to the corresponding messages in the promise machine memory.
- $\mathbb{S} \subseteq G.E$ records the set of *covered events* in G , namely the events that are fully executed by the promising machine. In particular, any promised write in $\mathbb{S}$ must have been fulfilled in MS . We write $\mathbb{S}_i$ for the covered events of thread i (i.e., $\mathbb{S} \cap G.E_i$).
- $\text{sc} \subseteq (\mathbb{S} \cap \mathcal{F}_{\text{sc}}) \times (\mathbb{S} \cap \mathcal{F}_{\text{sc}})$ is a relation that enforces a total order on the sc fences following the S view of the promise machine. Note that PS disallows sc fence steps to appear in certification runs, which is why $\text{sc} \subseteq \mathbb{S} \times \mathbb{S}$.

$$\begin{aligned}
\text{spo} &\triangleq G.\text{po} \cap (\mathbb{S} \times \mathbb{S}) & \text{srf} &\triangleq G.\text{rf} \cap (\mathbb{S} \times \mathbb{S}) & \text{smo} &\triangleq \text{mo} \cap (\mathbb{S} \times \mathbb{S}) \\
\text{sfr} &\triangleq (\text{srf}^{-1}; \text{smo}) \setminus [G.E] & \text{seco} &\triangleq (\text{srf} \cup \text{smo} \cup \text{sfr})^+ \\
\text{ssw} &\triangleq [\mathcal{E}_{\sqsubseteq_{\text{REL}}}] ; ([\mathcal{F}] ; \text{spo})^? ; [\mathcal{W}] ; \text{spo}_{\text{loc}}^? ; [\mathcal{W}_{\sqsubseteq_{\text{RLX}}}] ; \text{srf}^+ ; [\mathcal{R}_{\sqsubseteq_{\text{RLX}}}] ; (\text{spo} ; [\mathcal{F}])^? ; [\mathcal{E}_{\sqsubseteq_{\text{ACQ}}}] \\
\text{shb} &\triangleq (\text{spo} \cup \text{ssw})^+
\end{aligned}$$

Fig. 11. Auxiliary definitions for the simulation relation.

Using these relations, we define $\text{mo} \subseteq (G.\mathcal{W} \times G.\mathcal{W})|_{\text{loc}}$ to order the writes to each location x according to the timestamps of the respective messages in the promise machine.

$$\text{mo} \triangleq \{(e_1, e_2) \mid e_1.\text{loc} = e_2.\text{loc} \wedge \mathbb{W}(e_1).\text{ts} < \mathbb{W}(e_2).\text{ts}\} \setminus G.\text{cf}$$

In Fig. 11, we further define restrictions of the program order, the reads-from relation, the modification order, the reads-before relation, the extended coherence order, the synchronization relation and the happens-before order to events in $\mathbb{S}$. Intuitively, the execution $\langle \mathbb{S}, \text{spo}, \text{srf}, \text{smo} \rangle$ corresponds to the promise machine state.

We also define the behavior of the event structure with respect to $\mathbb{W}$ and $\mathbb{S}$ as follows.

$$\text{Behavior}(G, \mathbb{W}, \mathbb{S}) \triangleq \{(x, v) \mid \exists e \in \mathcal{W}_x \cap \mathbb{S}. e.\text{wval} = v \wedge \nexists e_1 \in \mathbb{S}. \text{mo}(e, e_1)\}$$

Using these auxiliary definitions, we now define the simulation relation as follows.

Definition 6. Let $\mathbb{P}$ be a program with T threads, $\Pi \subseteq T$ be a subset of threads, G be a **WEAKEST** event structure, and $\text{MS} = \langle \mathcal{TS}, \mathbb{S}, \text{M} \rangle$ be a promise machine state. We say that $G \sim_{\Pi} \text{MS}$ holds iff there exist $\mathbb{W}, \mathbb{S}$, and sc such that the following conditions hold:

- (1) G is consistent according to the **WEAKEST** model: $\text{isCons}_{\text{WEAKEST}}(G)$.
- (2) The local state of each thread in MS contains the program of that thread along with the sequence of covered events of that thread: $\forall i. \mathcal{TS}(i).\sigma = \langle \mathbb{P}(i), \text{labels}(\text{sequence}_{\text{spo}}(\mathbb{S}_i)) \rangle$.
- (3) Whenever $\mathbb{W}$ maps an event of G to a message in MS , then the location accessed and the written values match: $\forall e \in \text{dom}(\mathbb{W}). e.\text{loc} = \mathbb{W}(e).\text{loc} \wedge e.\text{wval} = \mathbb{W}(e).\text{wval}$.
- (4) All outstanding promises of threads $(T \setminus \Pi)$ have corresponding write events in G that are po-after $\mathbb{S}$: $\forall i \in (T \setminus \Pi). \forall e \in (\mathbb{S}_0 \cup \mathbb{S}_i). \mathcal{TS}(i).P \subseteq \{\mathbb{W}(e') \mid (e, e') \in G.\text{po}\}$.
- (5) For every location x and thread i , the thread view of x in the promise state MS records the timestamp of the maximal write visible to the covered events of thread i .

$$\forall i, x. \mathcal{TS}(i).V(x) = \max\{\mathbb{W}(e).\text{ts} \mid e \in \text{dom}([\mathcal{W}_x]; G.\text{jf}^?; \text{shb}^?; \text{sc}^?; \text{shb}^?; [\mathbb{S}_i])\}$$

- (6) The $\mathbb{S}$ events satisfy coherence: $\text{shb}; \text{seco}^?$ is irreflexive.
- (7) The atomicity condition holds for the $\mathbb{S}$ events: $\text{sfr}; \text{smo}$ is irreflexive.
- (8) The sc fences are appropriately ordered by sc : $[\mathcal{F}_{\text{sc}}]; (\text{shb} \cup \text{shb}; \text{seco}; \text{shb}); [\mathcal{F}_{\text{sc}}] \subseteq \text{sc}$.
- (9) The behavior of MS matches that of the $\mathbb{S}$ events: $\text{Behavior}(\text{MS}) = \text{Behavior}(G, \mathbb{W}, \mathbb{S})$.

Next, we establish Lemmas 1 and 2, which concern the preservation of the simulation relation by program steps. First, we show that non-promising steps of threads i preserve simulation modulo thread i .

Lemma 1. $G \sim_{\{i\}} \text{MS} \wedge \text{MS} \xrightarrow{i}^{np} \text{MS}' \implies \exists G'. G \rightarrow_{\mathbb{P}, \text{WEAKEST}}^* G' \wedge G' \sim_{\{i\}} \text{MS}'$.

We prove this lemma by case analysis over the promise machine state transition. We simply perform the corresponding event structure construction step and define updated $\mathbb{W}, \mathbb{S}$, and sc to show that $G' \sim_{\{i\}} \text{MS}'$ holds.

Using this lemma, we establish the following stronger property:

Lemma 2. $G \sim MS \wedge MS \rightarrow MS' \implies \exists G'. G \rightarrow_{\mathbb{P}, \text{WEAKEST}}^* G' \wedge G' \sim MS'$.

To prove Lemma 2, we first use Lemma 1 to establish $G_1 \sim_{\{i\}} MS'$ for an appropriate G_1 . Then, we consider the PS certification run of thread i , $MS' \xrightarrow{np}_i^* MS''$ with $MS''.\mathcal{TS}(i).P = \emptyset$. We inductively apply Lemma 1 on this certification run to extend G_1 to G' such that $G' \sim_{\{i\}} MS''$. Since, however, MS'' has no outstanding promises for thread i , it follows that $G' \sim MS''$, and consequently also $G' \sim MS'$ (for a smaller $\mathbb{S}$), as required.

Finally, Theorem 1 follows from Lemma 2 by induction.

5 DATA-RACE-FREEDOM GUARANTEES

In this section we define the notion of a data race and then show that our models provide certain guarantees for programs without races. These guarantees are stated as DRF (“data-race-freedom”) theorems, and allow programmers to understand the behavior of a certain class of programs without needing to understand the *WEAKESTMO* model.

In the context of an execution X , we say that two non-conflicting events are *concurrent* if they are not related by the happens-before relation. Two concurrent events are *racy* if they access the same location and at least one of them is a write. Let $X.\text{Race}$ be the set of all racy events of execution X :

$$X.\text{Race} \triangleq \text{dom}(((X.E \times X.E) \setminus (X.\text{hb}_{\text{C11}}^\equiv \cup X.\text{cf}))|_{\text{loc}} \cap \text{one}(\mathcal{W}))$$

where $\text{one}(A)$ is the relation saying that at least one of its components belongs to the set A ; that is, $\text{one}(A)(x, y) \triangleq (x \in A \vee y \in A)$. An execution X is *RA-race-free* if its races are confined to SC accesses (i.e., $X.\text{Race} \subseteq \mathcal{E}_{\text{SC}}$). Execution X is *RLX-race-free* if $X.\text{Race} \subseteq \text{Ld}_{\sqsupseteq \text{ACQ}} \cup \text{St}_{\sqsupseteq \text{REL}} \cup \text{U}_{\sqsupseteq \text{ACQ-REL}}$.

We recall the notion of RC11-consistency. An execution, X is RC11-consistent if it is consistent according to Definition 5 and $X.\text{po} \cup X.\text{rf}$ is acyclic. Then, we can establish the following result.

Theorem 2 (DRF-RLX). *Given a program $\mathbb{P}$, suppose its RC11-consistent executions are RLX-race-free. Then, $\text{Behavior}_{\text{WEAKESTMO}}(\mathbb{P}) = \text{Behavior}_{\text{RC11}}(\mathbb{P})$.*

To prove this theorem, we first show that $G.\text{jf} \subseteq G.\text{hb}$ holds for every event structure G constructed from $\mathbb{P}$. It then follows that at most one full execution X can be extracted from G (with $X.\text{rf} = G.\text{jf}$), and so that execution has $X.\text{po} \cup X.\text{rf}$ be acyclic. A more detailed proof can be found in our technical appendix [Chakraborty and Vafeiadis 2018].

Composing our DRF-RLX theorem with the DRF-SC theorem of Lahav et al. [2017, Theorem 4], we can derive a standard DRF-SC theorem, which says that programs whose races under SC are restricted to SC accesses exhibit only SC behavior.

Theorem 3 (DRF-SC). *Given a program $\mathbb{P}$, suppose its SC-consistent executions are RA-race-free. Then, $\text{Behavior}_{\text{WEAKESTMO}}(\mathbb{P}) = \text{Behavior}_{\text{SC}}(\mathbb{P})$.*

We can further combine our DRF-RLX theorem with Lahav et al. [2017, Theorem 5] to get another simple criterion for ensuring the absence of weak behaviors. Namely, if all accesses of a program $\mathbb{P}$ to shared locations are atomic (i.e., at least RLX) and shared-location accesses in the same thread are separated by an SC-fence, then that program has only SC behaviors.

Finally, composing our DRF-RLX theorem with the DRF-RA theorem of Kang et al. [2017, Theorem 2], we can derive a DRF-RA theorem for *WEAKESTMO*. The DRF-RA theorem states that a program whose races under release-acquire consistency (i.e., by treating all accesses as having release/acquire semantics) are confined to release/acquire accesses (or stronger) exhibits only release-acquire consistent behavior.

Theorem 4 (DRF-RA). *Given a program $\mathbb{P}$, suppose its RA-consistent executions are RLX-race-free. Then, $\text{Behavior}_{\text{WEAKESTMO}}(\mathbb{P}) = \text{Behavior}_{\text{RA}}(\mathbb{P})$.*

6 COMPILATION RESULTS

WEAKESTMO can serve as an intermediate memory model for an optimizing compiler from C/C++ to target architectures such as x86, PowerPC, and ARM. Typically, compilation takes place in multiple steps. First, the C/C++ program is mapped to the compiler’s intermediate representation (IR), which will follow the WEAKESTMO semantics. Then, a number of optimizing transformations are performed in the IR, and finally the IR is mapped to the code for the respective target architecture. We first define the transformation correctness and then study the correctness of these transformation steps. The correctness proofs of these transformations are in our technical appendix [Chakraborty and Vafeiadis 2018].

Definition 7. A transformation of program $\mathbb{P}_{\text{src}}$ in memory model M_{src} to program $\mathbb{P}_{\text{tgt}}$ in model M_{tgt} is *correct* if it does not introduce new behaviors: i.e., $\text{Behavior}_{M_{\text{tgt}}}(\mathbb{P}_{\text{tgt}}) \subseteq \text{Behavior}_{M_{\text{src}}}(\mathbb{P}_{\text{src}})$.

Variants of WEAKESTMO. Concerning the semantics of data races, we define two variants of the WEAKESTMO model.

WEAKESTMO-C11 According to the C and C++ standards, if a program has a consistent execution with NA-race, then the program has *undefined* behavior: i.e., the program may generate any arbitrary outcome [ISO/IEC 14882 2011; ISO/IEC 9899 2011]. To model this semantics of races, we say that a program has arbitrary behavior if it contains an extracted execution with a NA-race.

WEAKESTMO-LLVM The LLVM concurrency semantics differs from C11 in handling racy programs [Chakraborty and Vafeiadis 2017]. LLVM distinguishes between *write-write* races and *read-write* races. A write-write race is one happening between a pair of write events, whereas a read-write race is between a load event and a concurrent write event.

According to LLVM, only write-write NA-races result in undefined behavior. In contrast, read-write NA-races have defined behavior: the racy read may return an **undef(u)** expression which can be materialized to an arbitrary constant value of the same data type.

The WEAKESTMO-C11 has the same set of rules as WEAKESTMO; the WEAKESTMO-LLVM construction rules differ and are shown in our technical appendix [Chakraborty and Vafeiadis 2018]. In the following discussion, ‘WEAKESTMO’ refers to both variants of the model.

6.1 Mapping from C/C++ to WEAKESTMO

WEAKESTMO supports all the features of the C11 model except for `memory_order_consume` loads. Thus, after strengthening the consume loads into acquire loads (a transformation that is meant to be sound in C11), the mapping from C11 to WEAKESTMO is simply the identity mapping.

We take the C11 concurrency model to be its revised formal definition, namely the weakRC11 model by Lahav et al. [2017]. Since we use the same weakRC11 model for constraining inconsistent executions (§3.3), the correctness of the mapping is straightforward.

Theorem 5. *The identity mapping from weakRC11 to WEAKESTMO is correct (for both models).*

6.2 Optimizations as WEAKESTMO Source-to-Source Transformations

We move on to consider the correctness of standard compiler optimizations. Following the related work [Vafeiadis et al. 2015; Ševčík 2011], we consider compiler optimizations as being composed of a number of simple thread-local source-to-source transformations—introductions, reorderings, and eliminations of memory accesses—and we restrict attention to the correctness of those basic transformations.

To describe these transformations on a thread i , we use the informal syntax $\alpha \rightsquigarrow \beta$, where α and β are sequences of memory accesses possibly with metavariables ranging over both α and β .

Table 1. Allowed reorderings $a \cdot b \rightsquigarrow b \cdot a$ where $\ell \neq \ell'$.

$\downarrow a \setminus b \rightarrow$	$\text{Ld}_{\sqsubseteq \text{ACQ}}(\ell)$	$\text{Ld}_{\text{SC}}(\ell)$	$\text{St}_{\text{NA}}(\ell)$	$\text{St}_{\text{RLX}}(\ell)$	$\text{St}_{\sqsupset \text{REL}}(\ell)$	$\text{U}_{\sqsubseteq \text{ACQ}}(\ell)$	$\text{U}_{\sqsupset \text{REL}}(\ell)$	F_{ACQ}	F_{REL}	F_{SC}
$\text{Ld}_{\text{NA}}(\ell')$	✓	✓	✓	✓	✗	✓	✗	✓	✗	✗
$\text{Ld}_{\text{RLX}}(\ell')$	✓	✓	✓	✓	✗	✓	✗	✗	✗	✗
$\text{Ld}_{\sqsupset \text{ACQ}}(\ell')$	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗
$\text{St}_{\sqsubseteq \text{REL}}(\ell')$	✓	✓	✓	✓	✗	✓	✗	✓	✗	✗
$\text{St}_{\text{SC}}(\ell')$	✓	✗	✓	✓	✗	✓	✗	✓	✗	✗
$\text{U}_{\sqsubseteq \text{REL}}(\ell')$	✓	✓	✓	✓	✗	✓	✗	✗	✗	✗
$\text{U}_{\sqsupset \text{ACQ}}(\ell')$	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗
F_{ACQ}	✗	✗	✗	✗	✗	✗	✗	=	✗	✗
F_{REL}	✓	✓	✓	✗	✓	✗	✓	✓	=	✗
F_{SC}	✗	✗	✗	✗	✗	✗	✗	✗	✗	=

$\text{St}_{o'}(x, v') \cdot \text{St}_o(x, v) \rightsquigarrow \text{St}_o(x, v)$ (Overwritten write)

$\text{St}_o(x, v) \cdot \text{Ld}_{o'}(x, v) \rightsquigarrow \text{St}_o(x, v)$ (Read after write)

$\text{U}_o(x, v', v) \cdot \text{Ld}_{o'}(x, v) \rightsquigarrow \text{U}_o(x, v', v)$ (Read after update)

$\text{Ld}_o(x, v) \cdot \text{Ld}_{o'}(x, v) \rightsquigarrow \text{Ld}_o(x, v)$ (Read after read)

$\text{Ld}_{\text{NA}}(x, v) \cdot \text{St}_{\text{NA}}(x, v) \rightsquigarrow \epsilon$ (Non-atomic read write)

$\text{St}_{\text{NA}}(x, v') \cdot \tau \cdot \text{St}_{\text{NA}}(x, v) \rightsquigarrow \tau \cdot \text{St}_{\text{NA}}(x, v)$ (Non-adjacent overwritten write)

$\text{St}_{\text{NA}}(x, v) \cdot \tau \cdot \text{Ld}_{\text{NA}}(x, v) \rightsquigarrow \text{St}_{\text{NA}}(x, v) \cdot \tau$ (Non-adjacent read after write)

Fig. 12. Safe eliminations where $o' \sqsubseteq o$ and τ does not contain any x -accesses nor any release-acquire pairs. That is, for all τ_1, τ_2 such that $\tau = \tau_1 \cdot \tau_2$, either τ_1 does not contain a release label or τ_2 does not contain an acquire label.

By this syntax, we mean that every trace of memory accesses by the target code $\mathbb{P}_{\text{tgt}}(i)$ is also a trace of the source code of $\mathbb{P}_{\text{src}}(i)$ where at most one subsequence α of the trace has been replaced with β . More formally, $\mathbb{P}_{\text{tgt}}(i) \subseteq \mathbb{P}_{\text{src}}(i) \cup \{\tau \cdot \beta \cdot \tau' \mid \tau \cdot \alpha \cdot \tau' \in \mathbb{P}_{\text{src}}(i)\} \wedge \forall j \neq i. \mathbb{P}_{\text{tgt}}(j) = \mathbb{P}_{\text{src}}(j)$.

Reordering of Independent Accesses. Compilers often reorder a pair of independent instructions as part of instruction scheduling or in order to enable further optimizations. The safe (✓) and unsafe (✗) reorderings are listed in Table 1. We prove the correctness of the safe transformations by constructing for each target execution a source execution with the same behavior.

Theorem 6. *The safe reorderings in Table 1 are correct in both WEAKESTMO models.*

Counterexamples for the unsafe reorderings can be found in Vafeiadis et al. [2015]. As a side remark, GCC and LLVM do not exploit the full range of correct reorderings: they typically reorder only non-atomic accesses with respect to other accesses, but do not reorder pairs of atomic accesses.

Redundant Access Elimination. We enlist a number of correct elimination transformations in Fig. 12. The first four transformations remove an access because of an adjacent access to the same location. Compilers like GCC and LLVM perform such transformations only when the eliminated access is non-atomic (i.e., $o' = \text{NA}$); the eliminations are, however, also correct even for atomic accesses as long as the memory order of the eliminated access is not stronger than that of the justifying access (i.e., $o' \sqsubseteq o$). Combining the adjacent access eliminations with a number of safe reordering steps enables us to also eliminate certain non-adjacent accesses. Next, the fifth

WEAKESTM0	via RC11 [Lahav et al. 2017]			SPower
	x86	PowerPC	ARMv7	
store _{RLX}	mov	st	st	st
store _{REL}	mov	lwsync; st	dmb; st	lwsync; st
store _{SC}	mov; mfence	hwsync; st	dmb; st	hwsync; st
load _{NA}	mov	ld	ld	ld
load _{RLX}	mov	ld; cmp; bc	ld; cmp; bc	ld
load _{ACQ}	mov	ld; cbisync	ld; cbisb	ld; lwsync
load _{SC}	mov	hwsync; ld; cbisync	dmb; ld; cbisb	hwsync; ld; lwsync
CAS _{RLX}	lock cmpxchg	U	U	U
CAS _{REL}	lock cmpxchg	lwsync; U	dmb; U	lwsync; U
CAS _{ACQ}	lock cmpxchg	U ; cbisync	U ; cbisb	U ; lwsync
CAS _{ACQ-REL}	lock cmpxchg	lwsync; U ; cbisync	dmb; U ; cbisb	lwsync; U ; lwsync
CAS _{SC}	lock cmpxchg	hwsync; U ; cbisync	dmb; U ; cbisb	hwsync; U ; lwsync
F _{RLX}	mfence	lwsync	dmb	lwsync
F _{SC}	mfence	hwsync	dmb	hwsync

PowerPC $U \triangleq L$: lwarx; cmpw; bne L' ; stwcx.; bne L ; L' : cbisync $\triangleq$ cmp; bc; isync

ARMv7 $U \triangleq L$: ldrex; mov; teq L' ; strexeq; teq L ; L' : cbisb $\triangleq$ cmp; bc; isb

Fig. 13. Compilation schemes to x86, PowerPC, and ARM.

transformation removes a non-atomic load-store pair corresponding to the assignment $x = x$ in a programming language. The last two transformations eliminate a non-atomic access that is redundant because of a non-atomic store to the same location.

Theorem 7. *The eliminations in Fig. 12 are correct in both WEAKESTM0 models.*

Speculative Load Introduction. One subtle transformation that the LLVM compiler performs is speculative load introduction. This happens in cases such as the following

$$\text{if}(cond) \ a = X_o; \quad \rightsquigarrow \quad t = X_o; \text{if}(cond) \ a = t;$$

where a conditional shared memory load is hoisted outside the conditional by introducing a fresh temporary variable, which is used only when the condition holds. As the conditional move between registers can be performed without a conditional branch, LLVM's transformation avoids introducing a branch instruction at the cost of possibly performing an unnecessary load (if the condition is false).

This transformation is incorrect in the WEAKESTM0-C11 model because it may introduce a data race in the program when *cond* is false. It is, however, correct in WEAKESTM0-LLVM model.

Theorem 8. *The transformation $\epsilon \rightsquigarrow \text{Ld}_o(x, _)$ is correct in WEAKESTM0-LLVM.*

Strengthening. Finally, it is sound to strengthen the memory order of an access or fence (e.g., $F_o \rightsquigarrow F_{o'}$ for $o \sqsubset o'$) as well as to introduce a new fence instruction in a program (i.e., $\epsilon \rightsquigarrow F_o$). These transformations are correct because WEAKESTM0 is monotone.

6.3 Mapping from WEAKESTM0 to x86, PowerPC, and ARMv7

Finally, we present some results about the lowering transformations from WEAKESTM0 to the x86, PowerPC, and ARMv7 architectures.

First, by observing that *WEAKESTM0* is weaker than RC11, we use the RC11 mapping correctness results of [Lahav et al. \[2017\]](#) to get correct mappings to the aforementioned architectures (see Fig. 13). While these mappings are correct, they are suboptimal for PowerPC and ARMv7 in that they insert a fake conditional branch after each relaxed load.

To resolve this problem, we also prove the correctness of the intended mapping to the SPower model, a stronger version of the Power model due to [Lahav and Vafeiadis \[2016\]](#). In addition to the Power model's constraints, SPower requires $(po \cup rf)$ to be acyclic, which makes the fake control dependencies after relaxed loads unnecessary. The interest of the compilation scheme to SPower is twofold. First, even though load-store reordering is architecturally allowed by the PowerPC model, as far as we know, no existing implementations actually perform such reorderings, and hence SPower is a correct description of existing implementations (at the time of writing). Second, [Lahav and Vafeiadis \[2016\]](#) show that excluding CAS and SC accesses, PowerPC is equivalent to a model that first transforms the program by arbitrarily reordering independent memory accesses and then returns examines the behaviors of the transformed program according the SPower model. Thus, following the proof strategy of [Kang et al. \[2017\]](#), since *WEAKESTM0* supports reorderings of independent memory accesses, to prove correctness of compilation to PowerPC, it suffices to prove correctness of compilation to SPower.

Proving the correctness of the intended optimal mappings from *WEAKESTM0* to the full PowerPC model and to ARMv7, as well as to the ARMv8-Flat model [\[Pulte et al. 2018\]](#), is left for future work. To carry out these proofs, we intend to compile via the recent IMM model of [Podkopaev et al. \[2019\]](#), from which compilation results to multiple architectures have been developed.

7 RELATED WORK

As explained in the introduction, defining an adequate concurrency model that does not allow *out-of-thin-air* behaviors has been a longstanding research challenge, and has led to a number of proposals. Among them, we are aware of five that correctly differentiate *CYC*, *LB*, and *LBfd*.

In an early attempt, [Manson et al. \[2005\]](#) defined a fairly complicated operational model for Java. Their model, however, turned out to be unsound: [Ševčík and Aspinall \[2008\]](#) showed that it disallows some outcomes that were actually observable on Java programs. Nevertheless, they provided a set of 20 causality benchmarks, which became a standard benchmark for comparing memory models. (We discuss the outcomes of our models on those benchmarks in §3.5.)

[Jeffrey and Riely \[2016\]](#) later defined a model for a fragment of Java based on event structures. Similar to the Java memory model, executions are constructed by adding one event at a time. Each newly added read event has to be AE-justified: for all moves of the opponent (choosing how to continue executing the program), there must exist an extension containing a write whence the read gets its value. While AE-justification assigns the right semantics to the *RNG* example, its additional complexity over our more basic existential justification leads to some problems. First, enumerating all possible executions of a program is computationally infeasible, because one has to prove an assertion with N quantifier alternations (where N is the length of the execution). Second, their model fails to validate the reordering of independent read events, and therefore cannot be compiled to Power and ARM without additional fences. The paper, however, sketches an extension of the model that seems to fix this problem. It remains unclear whether the extension satisfies DRF-SC and can be compiled to hardware with the intended mappings.

[Pichon-Pharabod and Sewell \[2016\]](#) proposed a model for a fragment of C11 based on plain event structures (with only program order and conflict edges). They define an operational semantics over such event structures that iteratively constructs an execution. It does so by either committing a minimal event from the structure or by performing one of set of pre-determined compiler optimizations that rewrites the remaining event structure. The resulting model is quite complicated

and, as such, does not have much metatheory developed about it. It also fails to explain some weak program behaviors (cf. the ARM-weak program of [Kang et al. 2017]).

Chakraborty and Vafeiadis [2017] developed another model based on event structures for a fragment of LLVM concurrency without relaxed accesses (a.k.a. monotonic in LLVM’s terminology). As such, their event structures are substantially simpler than the ones in this paper as, e.g., the distinction between `WEAKEST` and `WEAKESTM0` does not arise in their restricted setting. This paper can be seen as an extension of that work both in terms of C11 feature coverage and also of the results proved.

Kang et al. [2017] defined the promising semantics, which comes with a number of results, which we are able to exploit for our `WEAKEST` model. It has also led to some follow up work, which proved correctness of compilation to ARM [Podkopaev et al. 2017, 2019] and the correctness of a program logic over it [Svendsen et al. 2018]. Nevertheless, as discussed, the promising semantics also has a number of shortcomings, which are difficult to resolve because of the model’s complexity and brittleness. We believe there are three main reasons for this complexity: (1) the use of timestamps and messages to represent executions as opposed to execution graphs; (2) the ability to promise a write at any point during execution in conjunction with the certification of promises, which can be of arbitrary length, which make it difficult to use the promising semantics as an execution oracle; and (3) the quantification over all ‘future’ memories in the certification of promises [Kang et al. 2017, §3], which is needed for handling RMWs. Its shortcomings include (a) admitting some questionable behaviors (e.g., that of `Coh-CYC`); (b) not supporting global optimizations; (c) providing overly strong semantics of RMWs, whose compilation to ARMv8 incurs an unintended performance cost (cf. `FADD`); and (d) not supporting SC accesses. Our work provides solutions to problems (a), (c), and (d), albeit our solution for (c) is not supported by a compilation correctness proof.

Besides these proposals, some papers have argued for disallowing the weak outcome of `LB` at the expense of requiring a more expensive compilation scheme to Power and ARM [Batty et al. 2013; Boehm and Demsky 2014; Lahav et al. 2017; Ou and Demsky 2018; Vafeiadis and Narayan 2013]. This seems adequate for the C/C++ setting, which forbids races on non-atomic accesses. It is, however, considered to be too expensive for Java, where for type-safety reasons all races must have defined semantics. Our work actually shows that one can still provide a sound and useful model that does allow the weak outcomes of `LB` and `LBfd`.

8 FUTURE WORK AND CONCLUSION

In this paper, we proposed a memory model based on event structures. Our model has two variants: the weaker version, `WEAKEST`, which simulates the promising semantics [Kang et al. 2017] and the stronger version, `WEAKESTM0`, which removes certain questionable behaviors allowed in promising semantics. We showed that the `WEAKESTM0` resolves out-of-thin-air problem, provides standard DRF guarantees, allows the desirable optimizations, and can be mapped to the architectures like X86, PowerPC, ARMv7. Additionally our models are flexible enough to adopt existing results from other papers [Chakraborty and Vafeiadis 2017; Lahav and Vafeiadis 2016; Lahav et al. 2017].

Even though `WEAKESTM0` and promising semantics are not comparable, we observe the same outcomes for the examples in Kang et al. [2017]. Specifically, similar to promising semantics, we allow the weak outcome of the program in Kang et al. [2017, §6]. Also, in the example in Kang et al. [2017, §7], `WEAKESTM0` disallows $x = 1$ in the source but allows the $x = 1$ outcome in the target program after sequentialization. Thus, similar to promising semantics, `WEAKESTM0` cannot handle sequentialization.

In the future, we intend to prove the correctness of the mappings from `WEAKESTM0` to the new ARMv8 model [Pulte et al. 2018] as well as the optimal mappings to PowerPC model [Alglave et al.

2014] via IMM [Podkopaev et al. 2019]. Additionally, we would like to mechanize our results in an interactive proof assistant, so as to provide further confidence about their correctness.

Further ahead, for WEAKESTMO to be established as a good programming language consistency model, we would have to develop techniques to reason about concurrent programs running in the WEAKESTMO model.

In terms of *program logics*, there are several works handling different fragments of RC11 (e.g., [Doko and Vafeiadis 2016, 2017; Kaiser et al. 2017; Turon et al. 2014; Vafeiadis and Narayan 2013]). Their soundness proofs, however, rely heavily on the acyclicity of $\text{po} \cup \text{rf}$ and it is unclear whether they can be adapted to the weaker setting of WEAKESTMO. Among these logics, RSL [Vafeiadis and Narayan 2013] should be sound under our model, because Svendsen et al. [2018] proved the soundness of a variant of it over the promising semantics. For the more advanced logics, however, the counterexample of Doko and Vafeiadis [2017] shows that certain adaptations will be needed.

In terms of *model checking* for weak memory concurrency, again there are a number of works that assume $\text{po} \cup \text{rf}$ acyclicity (e.g., [Abdulla et al. 2017; Kokologianakis et al. 2018]). There are also tools that do not require $\text{po} \cup \text{rf}$ acyclicity [Abdulla et al. 2016; Alglave et al. 2014; Norris and Demsky 2016]. These tools, however, use more expensive state enumeration techniques, and are significantly slower than the tools that assume $\text{po} \cup \text{rf}$ acyclicity. It remains to be seen whether similar model checking approaches can work for models based on event structures, and if so, how much slower they would be in comparison to the existing tools.

ACKNOWLEDGMENTS

We would like to thank Marko Doko, Evgenii Moiseenko, Anton Podkopaev, Azalea Raad, and the anonymous POPL reviewers for their feedback on earlier drafts of this paper.

REFERENCES

- Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonardsson, and Konstantinos Sagonas. 2017. Stateless model checking for TSO and PSO. *Acta Inf.* 54, 8 (2017), 789–818. <https://doi.org/10.1007/s00236-016-0275-0>
- Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Carl Leonardsson. 2016. Stateless model checking for POWER. In *CAV 2016 (LNCS)*, Vol. 9780. Springer, Heidelberg, 134–156. https://doi.org/10.1007/978-3-319-41540-6_8
- Sarita V. Adve and Mark D. Hill. 1990. Weak ordering—A new definition. In *ISCA 1990*. ACM, New York, NY, USA, 2–14. <https://doi.org/10.1145/325164.325100>
- Jade Alglave, Luc Maranget, Paul E. McKenney, Andrea Parri, and Alan Stern. 2018. Frightening small children and disconcerting grown-ups: Concurrency in the Linux kernel. In *ASPLOS 2018*. ACM, New York, NY, USA, 405–418. <https://doi.org/10.1145/3173162.3177156>
- Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding cats: Modelling, simulation, testing, and data-mining for weak memory. *ACM Trans. Program. Lang. Syst.* 36, 2 (2014), 7:1–7:74. <https://doi.org/10.1145/2627752>
- Mark Batty, Mike Dodds, and Alexey Gotsman. 2013. Library abstraction for C/C++ concurrency. In *POPL’13*. ACM, New York, NY, USA, 235–248. <https://doi.org/10.1145/2429069.2429099>
- Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod, and Peter Sewell. 2015. The Problem of Programming Language Concurrency Semantics. In *ESOP 2015 (LNCS)*, Vol. 9032. Springer, Heidelberg, 283–307. https://doi.org/10.1007/978-3-662-46669-8_12
- Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ concurrency. In *POPL’11*. ACM, New York, NY, USA, 55–66. <https://doi.org/10.1145/1926385.1926394>
- Hans-J. Boehm and Brian Demsky. 2014. Outlawing ghosts: Avoiding out-of-thin-air results. In *MSPC 2014*. ACM, New York, NY, USA, Article 7, 6 pages. <https://doi.org/10.1145/2618128.2618134>
- Soham Chakraborty and Viktor Vafeiadis. 2017. Formalizing the concurrency semantics of an LLVM fragment. In *CGO 2017*. IEEE, Piscataway, NJ, USA, 100–110. <https://doi.org/10.1109/CGO.2017.7863732>
- Soham Chakraborty and Viktor Vafeiadis. 2018. Technical appendix. Available at <http://plv.mpi-sws.org/weakest/>.
- Marko Doko and Viktor Vafeiadis. 2016. A program logic for C11 memory fences. In *VMCAI 2016 (LNCS)*, Vol. 9583. Springer, Heidelberg, 413–430. https://doi.org/10.1007/978-3-662-49122-5_20
- Marko Doko and Viktor Vafeiadis. 2017. Tackling real-life relaxed concurrency with FSL++. In *ESOP 2017 (LNCS)*, Vol. 10201. Springer, Heidelberg, 448–475. https://doi.org/10.1007/978-3-662-54434-1_17

- ISO/IEC 14882. 2011. Programming Language C++.
- ISO/IEC 9899. 2011. Programming Language C.
- Alan Jeffrey and James Riely. 2016. On thin air reads: Towards an event structures model of relaxed memory. In *LICS 2016*. ACM, New York, NY, USA, 759–767. <https://doi.org/10.1145/2933575.2934536>
- Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong logic for weak memory: Reasoning about release-acquire consistency in Iris. In *ECOOP 2017 (LIPIcs)*, Vol. 74. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 17:1–17:29. <https://doi.org/10.4230/LIPIcs.ECOOP.2017.17>
- Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A promising semantics for relaxed-memory concurrency. In *POPL 2017*. ACM, New York, NY, USA, 175–189. <https://doi.org/10.1145/3009837.3009850>
- Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. 2018. Effective stateless model checking for C/C++ concurrency. *Proc. ACM Program. Lang.* 2, POPL (2018), 17:1–17:32. <https://doi.org/10.1145/3158105>
- Ori Lahav and Viktor Vafeiadis. 2016. Explaining relaxed memory models with program transformations. In *FM 2016 (LNCS)*, Vol. 9995. Springer, Heidelberg, 479–495. https://doi.org/10.1007/978-3-319-48989-6_29
- Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing sequential consistency in C/C++11. In *PLDI 2017*. ACM, New York, NY, USA, 618–632. <https://doi.org/10.1145/3062341.3062352>
- Leslie Lamport. 1979. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers* 28, 9 (1979), 690–691. <https://doi.org/10.1109/TC.1979.1675439>
- Jeremy Manson, William Pugh, and Sarita V. Adve. 2004. Causality Test Cases. <http://www.cs.umd.edu/~pugh/java/memoryModel/unifiedProposal/testcases.html>.
- Jeremy Manson, William Pugh, and Sarita V. Adve. 2005. The Java memory model. In *POPL 2005*. ACM, New York, NY, USA, 378–391. <https://doi.org/10.1145/1040305.1040336>
- Brian Norris and Brian Demsky. 2016. A practical approach for model checking C/C++11 code. *ACM Trans. Program. Lang. Syst.* 38, 3 (2016), 10:1–10:51. <https://doi.org/10.1145/2806886>
- Peizhao Ou and Brian Demsky. 2018. Towards understanding the costs of avoiding out-of-thin-air results. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 136:1–136:29. <https://doi.org/10.1145/3276506>
- Jean Pichon-Pharabod and Peter Sewell. 2016. A concurrency semantics for relaxed atomics that permits optimisation and avoids thin-air executions. In *POPL’16*. ACM, New York, NY, USA, 622–633. <https://doi.org/10.1145/2676726.2676995>
- Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis. 2017. Promising compilation to ARMv8 POP. In *ECOOP 2017 (LIPIcs)*, Vol. 74. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 22:1–22:28. <https://doi.org/10.4230/LIPIcs.ECOOP.2017.22>
- Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis. 2019. Bridging the gap between programming languages and hardware weak memory models. *Proc. ACM Program. Lang.* 3, POPL (2019), 69:1–69:31. <https://doi.org/10.1145/3290382>
- Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2018. Simplifying ARM concurrency: Multicopy-atomic axiomatic and operational models for ARMv8. *Proc. ACM Program. Lang.* 2, POPL (2018), 19:1–19:29. <https://doi.org/10.1145/3158107>
- Kasper Svendsen, Jean Pichon-Pharabod, Marko Doko, Ori Lahav, and Viktor Vafeiadis. 2018. A separation logic for a promising semantics. In *ESOP 2018 (LNCS)*, Vol. 10801. Springer, Heidelberg, 357–384. https://doi.org/10.1007/978-3-319-89884-1_13
- Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. 2014. GPS: Navigating weak-memory with ghosts, protocols, and separation. In *OOPSLA 2014*. ACM, New York, NY, USA, 691–707. <https://doi.org/10.1145/2660193.2660243>
- Viktor Vafeiadis, Thibaut Balabonski, Soham Chakraborty, Robin Morisset, and Francesco Zappa Nardelli. 2015. Common compiler optimisations are invalid in the C11 memory model and what we can do about it. In *POPL 2015*. ACM, New York, NY, USA, 209–220. <https://doi.org/10.1145/2676726.2676995>
- Viktor Vafeiadis and Chinmay Narayan. 2013. Relaxed separation logic: A program logic for C11 concurrency. In *OOPSLA 2013*. ACM, New York, NY, USA, 867–884. <https://doi.org/10.1145/2509136.2509532>
- Jaroslav Ševčík. 2011. Safe optimisations for shared-memory concurrent programs. In *PLDI’11*. ACM, New York, NY, USA, 306–316. <https://doi.org/10.1145/1993498.1993534>
- Jaroslav Ševčík and David Aspinall. 2008. On validity of program transformations in the Java memory model. In *ECOOP’08*. Springer, Heidelberg, 27–51. https://doi.org/10.1007/978-3-540-70592-5_3
- Glynn Winskel. 1986. Event Structures. In *Advances in Petri Nets*. Springer, Heidelberg, 325–392. https://doi.org/10.1007/3-540-17906-2_31