

Model Checking on Multi-execution Memory Models

EVGENII MOISEENKO, JetBrains Research

MICHALIS KOKOLOGIANNAKIS, MPI-SWS, Germany

VIKTOR VAFEIADIS, MPI-SWS, Germany

Multi-execution memory models, such as Promising and Weakestmo, are an advanced class of weak memory consistency models that justify certain outcomes of a concurrent program by considering multiple candidate executions collectively. While this key characteristic allows them to support effective compilation to hardware models and a wide range of compiler optimizations, it makes reasoning about them substantially more difficult. In particular, we observe that Promising and Weakestmo inhibit effective model checking because they allow some surprisingly weak behaviors that cannot be generated by examining one execution at a time.

We therefore introduce Weakestmo2, a strengthening of Weakestmo by constraining its multi-execution nature, while preserving the important properties of Weakestmo: DRF theorems, compilation to hardware models, and correctness of local program transformations. Our strengthening rules out a class of surprisingly weak program behaviors, which we attempt to characterize with the help of two novel properties: *load buffering race freedom* and *certification locality*. In addition, we develop WMC, a model checker for Weakestmo2 with performance close to that of the best tools for per-execution models.

CCS Concepts: • **Theory of computation** → **Verification by model checking**; • **Software and its engineering** → **Semantics**; **Concurrent programming languages**.

Additional Key Words and Phrases: Weak memory models, model checking

ACM Reference Format:

Evgenii Moiseenko, Michalis Kokologiannakis, and Viktor Vafeiadis. 2022. Model Checking on Multi-execution Memory Models. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 152 (October 2022), 70 pages. <https://doi.org/10.1145/3563315>

1 INTRODUCTION

A *weak memory model* is a formal definition of the semantics of shared-memory concurrent programs, which allows more program outcomes (i.e., reachable thread configurations) than can be explained by a straightforward interleaving of the threads of a program. Consider, for instance, the load-buffering (LB) program below. (In our examples, we assume all variables are initialized to 0.)

$$\begin{array}{l} r_1 := x \text{ // reads } 1 \\ y := 1 \end{array} \parallel \begin{array}{l} r_2 := y \text{ // reads } 1 \\ x := r_2 \end{array} \quad (\text{LB})$$

The annotated outcome, where both threads read the value 1, cannot be explained by simply interleaving the instructions of the two threads, since the first instruction to execute can only read 0 (the initial value). Multicore Arm processors, however, do exhibit this outcome because they often execute independent instructions out of order. For example, Thread 1 may first perform the $y := 1$ write, then Thread 2 can read $y = 1$ and write $x := 1$, and then Thread 1 can finish by reading $x = 1$.

Authors' addresses: Evgenii Moiseenko, JetBrains Research, evgeniy.moiseenko@jetbrains.com; Michalis Kokologiannakis, MPI-SWS, Germany, michalis@mpi-sws.org; Viktor Vafeiadis, MPI-SWS, Germany, viktor@mpi-sws.org.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/10-ART152

<https://doi.org/10.1145/3563315>

Most memory models—including all those for hardware architectures—are defined in a *per-execution* style, where each possible program outcome is explained by a single program execution witnessing that outcome. As Batty et al. [2015] have shown, however, this per-execution style is inadequate for defining the semantics of programming languages like C/C++ that strive to provide features with optimal efficiency.

For this reason, advanced language-level memory models (e.g., [Chakraborty et al. 2019; Jagadeesan et al. 2020; Jeffrey et al. 2016; 2022; Kang et al. 2017; Manson et al. 2005; Paviotti et al. 2020; Pichon-Pharabod et al. 2016]) instead adopt a *multi-execution* style, where multiple candidate program executions are used together to justify a given program outcome. For example, Promising [Kang et al. 2017] achieves this by augmenting the regular in-order program execution with an additional step: a promise to execute a future write that is backed up with an additional execution justifying that it is possible to fulfill the promise. Promising explains the **LB** outcome as follows. Thread 1 first promises to write $y := 1$, which is fulfillable by running the thread to completion. Then Thread 2 reads 1 and writes $x := 1$, and then Thread 1 reads 1, arriving at the desired outcome. The fact that promises have to be justified by other promise-free executions is what makes Promising a multi-execution model. And this aspect of Promising is crucial: removing promise certification would lead to an overly weak model that would allow some clearly undesirable outcomes known as ‘out-of-thin-air’ outcomes in the literature. Other multi-execution models, such as Weakestmo [Chakraborty et al. 2019], explain the **LB** outcome in a similar way, but explicitly represent the multiple program executions as an event structure.

While multi-execution models, such as Promising and Weakestmo, can be implemented efficiently on a wide range of hardware architectures [Moiseenko et al. 2020; Podkopaev et al. 2019], they have a significant drawback: they are very difficult to reason about, especially in an automated fashion. Although there are numerous effective automated techniques for verifying programs under per-execution weak memory models (e.g., [Abdulla et al. 2015a; 2018; 2015b; Barnat et al. 2013; Bouajjani et al. 2013; Demsky et al. 2015; Huang et al. 2016; Kokologiannakis et al. 2017; 2019]), there are no automated techniques for reasoning about multi-execution models. Verification of finite-state programs (with loops) is undecidable [Abdulla et al. 2021], and even model checking (i.e., enumerating all possible outcomes) of small loop-free programs is typically intractable.

We believe that the difficulty in automated verification on multi-execution models is largely due to the unconstrained nature of the models’ out-of-order execution mechanisms. To make such a model amenable to model checking, one has to constrain the use of multiple executions in two distinct ways: (1) on *when* multiple executions are introduced to explain a certain behavior (i.e., in terms of Promising, when a promise can be made); and (2) on *how much* these multiple executions interact with one another. Naturally, we desire to restrict multi-execution nature as much as possible: additional executions should only be allowed when there is a good reason to do so, and they should not be allowed to diverge too much from one another.

Our first contribution is to propose two properties that constrain multiple executions in the aforementioned ways, namely *load buffering race freedom* and *certification locality* (§2). Besides their use for model checking, these properties rule out certain program outcomes that cannot be observed by any combination of reasonable compiler optimizations on existing hardware platforms, and so may be of independent interest.

Subsequently, we introduce Weakestmo2, a strengthening of Weakestmo [Chakraborty et al. 2019] that satisfies load buffering race freedom and certification locality (§3). Further, we show that Weakestmo2 preserves the soundness of Weakestmo’s efficient compilation schemes to hardware-level models shown by Moiseenko et al. [2020] (§3.4) as well as the soundness of local program transformations (§3.5).

Finally, we develop an effective model checking algorithm, WMC, for verifying programs running on Weakestmo2 (§4), and implement it as an extension of the GENMC model checker [Kokologianakis et al. 2019; 2021]. Our experiments (§5) demonstrate that WMC’s performance is superior to that of the few other tools for (per-execution) weak memory models admitting the LB behavior (i.e., the annotated outcome of the `LB` program), and comparable to the best tools for models that forbid the LB behavior.

2 OVERVIEW

In this section, we recall the basic terminology of weak memory models (§2.1) and motivate our Weakestmo2 model. For the latter, we introduce two properties of multi-execution models that are needed for effective model checking but are not satisfied by Promising and Weakestmo.

Load buffering race freedom (LBRF, §2.2) restricts the multi-execution nature of a model to affect only programs with *load buffering races* (LB races). LBRF is defined analogously to the well-known *data race freedom* (DRF) guarantee [Adve et al. 1996; Manson et al. 2005] replacing the notion of a data race with a stronger novel notion of an LB race.

Certification locality (CL, §2.3) concerns multi-execution models with a certification mechanism and restricts *how much* certification executions can differ from the main execution. CL allows one to determine locally whether a certain load-store reordering is allowed, which disallows certain ‘bait-and-switch’ behaviors [Jagadeesan et al. 2020].

We conclude this section by explaining how these properties enable effective model checking (§2.4).

2.1 Execution Graphs and Data Race Freedom

In the literature of axiomatic (a.k.a. declarative) per-execution memory models, the possible executions of a program P under a model M are represented as a set of *execution graphs* that correspond to the instructions of P and satisfy M ’s consistency predicate. Execution graphs consist of:

- a set of nodes, called *events*, which represent the individual operations performed by the program that are relevant for concurrency (e.g., reads R , writes W , fences F), and
- various kinds of directed edges between events, such as:
 - the *program order* (po), relating events in the same thread in their control-flow order as well as initialization events before other events, and depicted as a solid black edge;
 - the *reads-from* (rf) relation, connecting each read to the write it is reading from, and depicted as a dashed green edge from the write to the read;
 - the *happens-before* (hb) order, a subset of $\text{porf} \triangleq (\text{po} \cup \text{rf})^+$ that includes po , capturing ordering due to intra-thread control-flow and inter-thread synchronization. (For simplicity, the examples of this paper do not contain inter-thread synchronization, and so $\text{hb} = \text{po}$.)

The strongest useful model in this framework is *sequential consistency* (SC) [Lamport 1979], which requires that there be a total order $<_{\text{SC}}$ among all events of an execution graph extending porf such that each read reads from the most recent same-location write that precedes it in $<_{\text{SC}}$. Other models place weaker constraints, with models such as x86-TSO [Owens et al. 2009] and RC11 [Lahav et al. 2017] requiring (among other things) that porf be acyclic.

Figure 1 shows four execution graphs corresponding to the `LB` program from §1. These graphs are generated by picking every possible value for each read event that matches the value written by the write event to the same location whence the read is reading from. The first three graphs are SC-consistent (and therefore also RC11-consistent). By contrast, the fourth graph, witnessing the ‘load buffering’ behavior, is not RC11-consistent because it contains a porf cycle.

A standard property that is expected of memory models is the DRF_{SC} guarantee [Adve et al. 1996; Manson et al. 2005], which constrains non-SC behaviors to occur only on programs with data

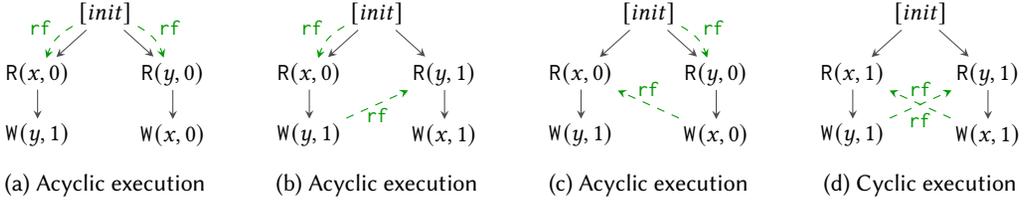


Fig. 1. Execution graphs of LB.

aces. We say that two events (of different threads) are *concurrent* if they are not ordered by **hb**. A pair of events form a *data race* in an execution graph if they are concurrent memory accesses to the same location, at least one of which is a write event. A program P is *data-race-free* under a memory model M if no consistent execution graph of P under M contains a data race.

Definition 2.1 (DRF). A memory model M provides the *data race freedom guarantee* with respect to a stronger memory model M' (written $\text{DRF}_{M'}$) if, for any data-race-free program P under M' , its consistent executions under M are exactly the same as under M' .

A memory model providing the DRF_{SC} guarantee allows programmers to adopt a defensive programming strategy of avoiding data races (e.g., by using locks), which incurs some performance degradation but relieves them from the need to learn and understand the memory model definition.

2.2 Load Buffering Race Freedom

Load buffering race freedom is analogous to DRF. We say that a **po** edge between a read and a write is *reorderable* (**rpo**) if the two accesses are *relaxed* following C11 terminology (i.e., weaker than release/acquire) and there is no fence between them. A reorderable edge signifies that the two events may be executed out of order (e.g., under Promising [Kang et al. 2017], the write may be promised), and thus contribute to a load buffering behavior.

Definition 2.2 (Load buffering race). A pair of events r and w form a *load buffering race* (LB race) in an execution graph if r is a read, w is a concurrent write to the same location, and there is a **rpo**; (**rf** \ **po**); **porf** path from r to w (i.e., a **porf**-path starting with a reorderable edge).

For instance, execution graph (b) of Fig. 1 has a load buffering race between the $R(x, 0)$ and $W(x, 1)$ events. Similarly, graph (c) has an LB race between the $R(y, 0)$ and $W(y, 1)$ events.

Existence of a **porf**-path between w and r indicates that the write might depend on the read. In models like RC11, it means that r cannot read from w because that would create a **porf**-cycle, such as the one in Fig. 1(d). We insist that the first edge along this path is reorderable to rule out cases where an explicit fence has been added to prevent the load-buffering behavior.

Definition 2.3 (LB-race-free program). A program P is *LB-race-free* under a memory model M if no consistent execution graph of P under M contains a load buffering race.

Definition 2.4 (LBRF). A memory model M provides the *load buffering race freedom* guarantee with respect to a stronger memory model M' (written $\text{LBRF}_{M'}$) if, for any LB-race-free program P under M' , its consistent executions under M are exactly the same as under M' .

Normally, we take M' to be some standard model that forbids **porf**-cycles, such as RC11. Similar to DRF, $\text{LBRF}_{\text{RC11}}$ allows one to program defensively against a model M without even knowing its definition by avoiding LB races. Since absence of LB races is to be checked with respect to RC11, one can use any of the existing tools and methodologies that reason about program correctness

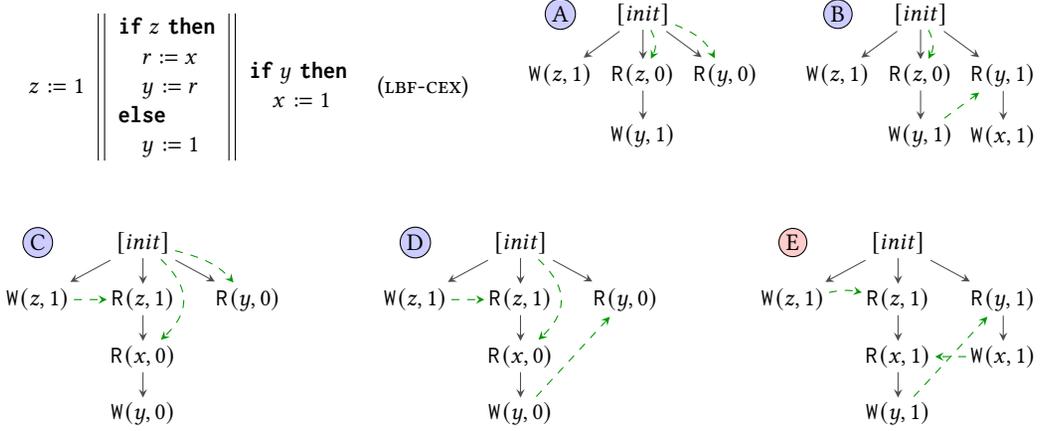


Fig. 2. A program with its RC11-consistent execution graphs and a problematic cyclic execution.

under RC11 (e.g., iGPS, HERD, RCMC, GENMC, etc). Moreover, LB races can easily be removed by making the `porf`-paths from the racy read to the write start with non-reorderable edges. This can be achieved, for example, by strengthening the access mode of the read to be an acquire or by adding an acquire or release fence right after it. In §5.4, we report on a script that does so automatically. Based on the results of Ou et al. [2018] and our experience, the run-time overhead incurred by these extra fences is extremely low, if at all perceptible.

2.2.1 LBRF and Existing Models. Among the axiomatic per-execution models that allow LB behaviors, the original C11 model [Batty et al. 2011] does not satisfy $\text{LBRF}_{\text{RC11}}$ because it allows out-of-thin-air outcomes and does not even satisfy DRF_{SC} . Lower-level models that track syntactic dependencies between instructions, such as IMM [Podkopaev et al. 2019], Power [Alglave et al. 2014], and ARM-8 [Pulte et al. 2018], satisfy LBRF wrt. their strengthenings with `porf`-acyclicity. The additional behaviors they allow over their strengthenings are `porf`-cycles with at least one `po` edge from a load to a store being reorderable. By changing the `rf` edge of that load to read from a prior store, and relying on “receptiveness” of the mapping from programs to executions, we can construct an LB race. As an example of this reasoning, we prove the following theorem:

THEOREM 2.5. IMM satisfies $\text{LBRF}_{\text{RC11}}$. (See Appendix A for the proof.)

Among the multi-execution models, Promising [Kang et al. 2017; Lee et al. 2020] and Weakestmo [Chakraborty et al. 2019] do not satisfy $\text{LBRF}_{\text{RC11}}$. To see this, consider the LBF-CEX program along with its executions shown in Fig. 2. Although none of its RC11-consistent executions (A, B, C, D) contains an LB race, both Promising and Weakestmo allow the additional execution E where r gets the value 1. This execution can arise in the following manner. First, thread 2 promises the $W(y, 1)$ store; the promise is allowed because thread 2 can read $z = 0$ and fulfill it. Then, thread 3 executes: it reads $y = 1$ and writes 1 to x . Finally, threads 1 and 2 execute: thread 2 reads $z = 1$ and $x = 1$, and subsequently writes 1 to y thereby fulfilling its promise.

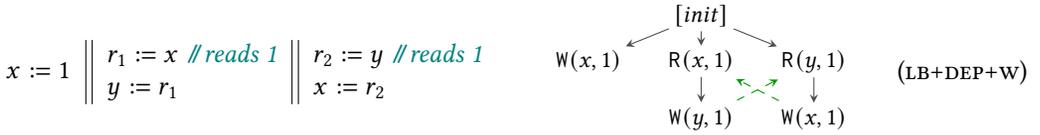
That said, it is fairly straightforward to strengthen Promising and Weakestmo to satisfy LBRF by restricting when promises can be made. The existing condition of Promising requires there to be a thread-local execution that certifies the promise. In addition to that, we can require there to be an execution witnessing an LB race. As we shall shortly see, however, satisfying LBRF alone is not

sufficient for effective model checking. We need another locality property that completely forbids non-local certifications that depend on external writes like the $z = 1$ write in the example above.

2.3 Certification Locality

LBRF restricts *when* writes may be promised (i.e., only upon LB races), but not *how*. In models like Promising and Weakestmo (even if artificially strengthened to satisfy LBRF), certifications of outstanding promises can differ a lot from the actual executions whose promises they are certifying. As we will shortly see, this leads to some rather weak outcomes.

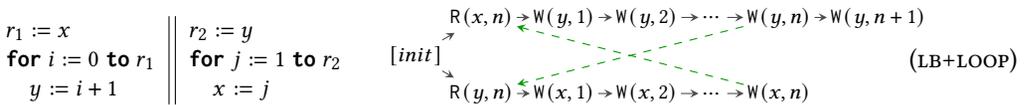
We start with an example that does not directly constitute an odd behavior but that is indicative of what can go wrong with non-local certifications.



The annotated outcome is perfectly valid if Thread 2 gets the value 1 from Thread 1 and propagates it to Thread 3—this outcome is even allowed under SC. What is odd, however, is to justify the outcome by the depicted execution graph, where Threads 2 and 3 read from each other without any flow of information from Thread 1 to either of these threads. The problem is that both Promising and Weakestmo essentially allow this justification. After executing Thread 1, Thread 2 can promise to write 1 to y (by reading from Thread 1). Thread 3 then executes to completion, reading from Thread 2, and writing 1 to x . Finally, Thread 2 continues, reads from Thread 3 and fulfills its promise, thus resulting in the graph above. What went wrong in this execution is that, while the write of Thread 1 was needed to enable the early execution of the write in Thread 2, this dependency is then forgotten in the final execution.

Certification locality forbids exactly this pattern. We say that a write is *non-local* at a certain point in a thread if it neither happens-before that point nor is read by an event happening-before that point. CL requires that every non-local write that is read in the process of certifying a promise also be read by the same thread while issuing the promise and vice versa. In other words, the two executions of the thread issuing a promise should read exactly the same non-local writes between the point the promise is issued (i.e., when the executions start diverging because they read from different writes) and the points where the promises are fulfilled.

To further justify CL, we next show an example of a really weak behavior allowed by Promising and Weakestmo. Consider the **LB+LOOP** program below and the associated execution graph that reads an arbitrary natural number n into r_1 and r_2 .



The displayed outcome is allowed by Promising (and similarly by Weakestmo) because Thread 1 can initially promise $y := 1$ (which it can trivially fulfill), then Thread 2 can promise $x := 1$ (which it can fulfill by reading $y = 1$ from Thread 1), then Thread 1 can promise $y := 2$ (by reading $x = 1$), then Thread 2 can promise $x := 2$ and so on. Again, the problem in this unbounded execution is the lack of certification locality. Once a thread depends on an external write to justify a promise, it should not “change its mind” and ignore that write in favor of a different one.

2.3.1 Enforcing CL. Repairing Promising to satisfy CL is not as straightforward as it is for LBRF because it does not track the exact set of writes read by a thread. We therefore instrument the Promising state by attaching to each promise a set of external writes that must be read before fulfilling the promise. We then check that as long as a thread contains outstanding promises, it can only read from external writes that are recorded in its promise set and, moreover, that when a promise is fulfilled all its attached external writes have been read. We present a suitable definition in Appendix B. While our adapted definition achieves certification locality, we have not investigated whether it also satisfies the remaining properties of the original Promising model, namely correctness of compilation and source-to-source transformations.

By contrast, repairing Weakestmo is much easier because the certification runs are available as part of the event structure. To enforce CL, one can simply augment the Weakestmo definition with an axiom that rules out ‘bait-and-switch’ behaviors. We defer the formal definition of our resulting model, Weakestmo2, to §3. There we also establish three important results about Weakestmo2: it provides the LBRF_{RC11} and DRF_{SC} guarantees (§3.3), its expected compilation schemes to hardware models are sound (§3.4), and it supports the expected local reorderings and eliminations (§3.5).

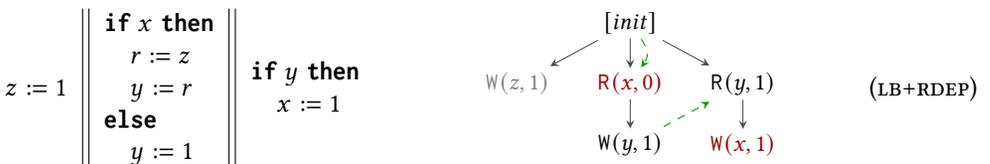
2.4 WMC: Effective Model Checking for Multi-Execution Memory Models

Apart from serving as criteria to rule out certain weak behaviors of multi-execution memory models, LBRF and CL are instrumental in making model checking of programs on such models feasible. We briefly describe how WMC, our model checking algorithm for Weakestmo2, exploits these properties. (We defer the full presentation of WMC to §4.)

Stateless model checkers, such as GENMC [Kokologiannakis et al. 2019], generate all the consistent executions of a given program in an incremental fashion. Starting with the empty execution graph, they add one event at a time in all possible consistent ways. With multi-execution memory models, there are two ways in which events can be added: either by adding events in order or by promising a write out of order. Considering all possible promises, however, is infeasible because of the huge state-space that would need to be explored.

This is where LBRF helps: Promises only need to be considered when an LB race is encountered. WMC therefore first generates the RC11-consistent executions of a program using techniques from the literature. Whenever it discovers an LB race, it uses an additional promising/certification mechanism to explore executions with *porf* cycles. As an example of how WMC works at a high level, consider the **LB** program from §1. Under RC11, **LB** has 3 consistent executions (Figures 1a to 1c), and WMC starts by enumerating those. While doing so, however, it notices that the execution of Fig. 1b contains a load buffering race between events $R(x, 0)$ and $W(x, 1)$. Thus, WMC creates the execution of Fig. 1d by *promising* the write $W(y, 1)$ in Thread 1 (so that $R(y, 1)$ can read from it), and then by making $R(x, 0)$ read from $W(x, 1)$. Subsequently, it *certifies* the promised write in Thread 1, and generates the cyclic execution.

This is exactly the point where CL is useful: It ensures that a promise can quickly be certified or withdrawn by executing only the thread containing the racy read. We illustrate this point with the program below, which is a slight variant of **LBF-CEX** with the reads of x and z swapped in Thread 2. This program has an acyclic execution with an LB race on the x accesses (shown to the side), and so the model checker will have to consider promising $W(y, 1)$ in Thread 2.



The model checker will then try to certify the promise by re-executing Thread 2 when it reads 1 for x . The only way to certify the promise is by reading from $W(z, 1)$. However, depending on the scheduling strategy, this write might not be available at the point when the LB race was detected. Therefore, to avoid missing any executions, the model checker would have to postpone the certification of $W(y, 1)$ and first explore other parts of the program, an exploration which may itself involve more promises and certifications. This approach is not only complicated to implement correctly, but also rather inefficient because in the fairly common case that the promise cannot be certified, one would wastefully explore all the possible executions of other parts of the program, leading to many blocked explorations.

CL enables a crucial optimization: Since (under CL) certification runs can read only from writes in the `porf`-prefix of the promise (see §3), the promise can be certified *locally* and *immediately* by re-executing the thread in question. In our example, regardless of whether $W(z, 1)$ has been added to the graph or not, WMC will only consider the read $R(z, 0)$, and therefore conclude that the promise $W(y, 1)$ cannot be certified.

As we show in §5, our technique for promising writes *only* upon detecting an LB race and certifying promised writes *immediately* and *locally* scales much better than the existing techniques that tackle similar memory models. Indeed, the few existing model checkers that handle such models (e.g., [Norris et al. 2013; Pulte et al. 2019]) explore all possible certifiable promises, irrespective of whether they result in additional behaviors. Such wasteful blind exploration of promises is the key factor contributing to their poor performance.

Further, the optimization due to LBRF can also be used to improve by a moderate amount model checking on per-execution models that allow LB behaviors, such as IMM. To demonstrate this, we took HMC [Kokologiannakis et al. 2020], a model checker that operates under IMM, and implemented HMC_{LBRF} , a version of HMC that leverages LBRF. Similarly to WMC, HMC_{LBRF} starts by enumerating RC11-consistent executions, and falls back to dependency tracking only upon detecting an LB race. As we show in §5, HMC_{LBRF} outperforms HMC in LB-race-free programs, thereby further showcasing LBRF’s usefulness in verification.

3 REPAIRING WEAKESTMO

In this section, we describe how Weakestmo2, our strengthening of Weakestmo, supports LBRF and CL. In what follows, we assume a simplified version of the model, containing only relaxed accesses and fences. For the full model, we refer the reader to Chakraborty et al. [2019] and Appendix C.

Weakestmo is an axiomatic multi-execution memory consistency model. Unlike conventional axiomatic models, which determine the validity of particular outcome based on a consistency predicate on a single execution graph, Weakestmo considers the execution graph consistent if it can be extracted from some consistent *event structure*. Event structures encompass multiple runs of a program in a single graph. That is, event structures can contain several execution branches of the same thread, which are used to model the speculative out-of-order execution of instructions.

Definition 3.1. An event is a tuple $\langle id, tid, lab \rangle$ where $id \in \mathbb{N}$ is a unique identifier for the event, $tid \in \mathbb{N}_\perp$ identifies the thread to which the event belongs (\perp for initialization events), and lab is a *label* of the form: (1) $R(x, v)$ for a read of $v \in \text{Val}$ from $x \in \text{Loc}$; (2) $W(x, v)$ for a write of $v \in \text{Val}$ to $x \in \text{Loc}$; (3) F for a fence.

We write Event for the set of all events. The set of all reads is $R \triangleq \{\langle i, t, l \rangle \mid l = R(\cdot)\}$. The sets of all writes and fences are defined analogously. Given an event e , we write $id(e)$, $tid(e)$, and $lab(e)$ to project its components, and $loc(e)$ and $val(e)$ to project its location and value respectively.

Further, given a relation r , we use $dom(r)$ and $rng(r)$ to denote its domain and codomain, while $r^?$, r^+ , and r^* denote its reflexive, transitive, and reflexive-transitive closures, respectively. We write

from the structure in Fig. 3 does not form an execution graph corresponding to the program **LB**. In addition, we must take all **po**-predecessors of the event and the writes that justify them. There is, however, one subtle point. Event e_1^1 is justified by e_{21}^1 , which is in conflict with e_{12}^1 ; thus, instead of e_{21}^1 , we pick the equivalent write e_{22}^1 . To achieve that, we define the derived *reads-from* relation $\mathbf{rf} \triangleq (\mathbf{ew}^*; \mathbf{jf}) \setminus \mathbf{cf}$ by extending the **jf** relation to \mathbf{ew}^* equivalence classes.

Definition 3.5. A justified configuration C of the event structure S is a subset of its events $C \subseteq S.E$, s.t.

- C is conflict free: $\mathbf{cf} \cap C \times C = \emptyset$;
- C is closed w.r.t. **po**-prefixes: $\text{dom}(\mathbf{po}; [C]) \subseteq C$;
- C is **rf**-complete: $C \cap R \subseteq \text{rng}([C]; \mathbf{rf})$.

Definition 3.6. An execution graph G is extracted from an event structure S , denoted as $S \triangleright G$, if $G.E$ is a justified configuration of S s.t. $G.x = S.x|_{G.E}$ for $x \in \{\mathbf{po}, \mathbf{ew}, \mathbf{co}\}$ and $G.\mathbf{jf} = G.\mathbf{rf} = S.\mathbf{rf}|_{G.E}$.

Definition 3.7. The behavior of an execution graph G , denoted as $\mathcal{B}(G)$, is a mapping assigning to each location x its final value, that is, the value written by the **co**-maximal write to x in G .

For example, in Fig. 3 the set of events marked by  forms a justified configuration and thus induces an execution graph with the behavior $\{x \mapsto 1, y \mapsto 1\}$.

3.1 Weakestmo Consistency

To filter out nonsensical event structures, Weakestmo defines a number of *consistency constraints*, which depend on the following auxiliary definitions.

- $\mathbf{hb} \triangleq (\mathbf{po} \cup \mathbf{sw})^+$ — *Happens-before* is the transitive closure of the program order and *synchronizes-with* relations. The latter connects synchronized events. For example, it connects two fences if there exists a $\mathbf{po}; \mathbf{rf}; \mathbf{po}$ path between them.
- $\mathbf{ecf} \triangleq (\mathbf{hb}^{-1})^?; \mathbf{cf}; \mathbf{hb}^?$ — *Extended conflict* propagates the conflict relation along **hb**.
- $\mathbf{eco} \triangleq (\mathbf{co} \cup \mathbf{rf} \cup \mathbf{rf}^{-1}; \mathbf{co})^+$ — *Extended coherence* is almost a total order on accesses to a given location; it orders every pair of such accesses except for equal writes, and reads reading from the same write.
- $\mathbf{jo} \triangleq (\mathbf{jf} \setminus \mathbf{po}); (\mathbf{po} \cup \mathbf{jf})^*$ — *Justification order*. The **jo** predecessors of an event e are all the writes that cause the event e indirectly through some inter-thread communication. We call these writes the *justification set* of the event e .
- A write event w is a *promise* w.r.t. an event e if $\langle w, e \rangle \in \mathbf{cf} \cap \mathbf{jo}$. Additionally:
 - if $\langle w, e \rangle \in \mathbf{ew}^+$; $\mathbf{po}^?$ then w is a *certified promise*;
 - if $\langle w, e \rangle \notin \mathbf{ew}^+$; $\mathbf{po}^?$ then w is a *pending promise*.

Consider Fig. 3 again. Write event e_{21}^1 is a pending promise w.r.t. e_{12}^1 and a certified promise w.r.t. e_{22}^1 . The later event is exactly the equal write that certifies the promise.

Definition 3.8. Event structure S is Weakestmo-consistent if the following conditions hold.

- | | |
|---|---------------------|
| • \mathbf{ecf} is irreflexive. | (NON-CONTRADICTORY) |
| • $\mathbf{jf} \cap \mathbf{ecf} = \emptyset$ | (WELL-JUSTIFIED) |
| • $\mathbf{po} \cup \mathbf{jf}$ is acyclic | (NO-THIN-AIR) |
| • $\mathbf{hb}; \mathbf{eco}^?$ is irreflexive. | (COHERENT) |
| • $[F]; \mathbf{po}; \mathbf{ew}^+ \subseteq \mathbf{po}$ | (WELL-FENCED) |
| • $\mathbf{cf} \cap \mathbf{jo} \subseteq \mathbf{ew}^+; (\mathbf{po} \cup \mathbf{po}^{-1})^?$ | (CERTIFIED) |
| • $\mathbf{ew} \subseteq (\mathbf{cf} \cap (\mathbf{jo} \cup \mathbf{jo}^{-1}))^+$ | (GROUNDED) |

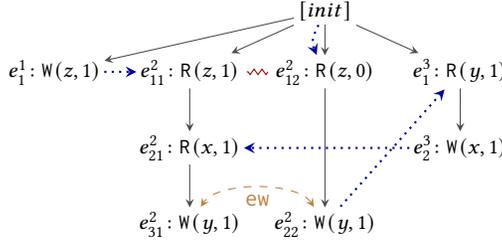


Fig. 4. Weakestmo2 **inconsistent** event structure of **LBF-CEX**.

The first two constraints forbid nonsensical event structures where some event either is in conflict with itself, or justifies a conflicting event. **NO-THIN-AIR** prevents the values of reads to appear out-of-thin-air. **COHERENT** enforces the *coherence* property at the level of the whole event structure¹. The last three axioms are related to promises and certification: **WELL-FENCED** prevents promises to be issued across fences; **CERTIFIED** ensures that all promises are eventually certified; and **GROUND** guarantees that the **ew** relation is only used to certify promises, and not to link arbitrary writes.

Definition 3.9. Execution graph G is Weakestmo-consistent if there exists Weakestmo-consistent event structure S s.t. G can be extracted from S .²

3.2 Weakestmo2 Consistency

As already mentioned in §2.2.1, Weakestmo-consistency guarantees neither $\text{LBRF}_{\text{RC11}}$ nor CL. For example, Figure 4 depicts a Weakestmo-consistent event structure of **LBF-CEX** that justifies the weak outcome $r_1 = 1$, which is forbidden by $\text{LBRF}_{\text{RC11}}$.

In order to get the cyclic execution of **LBF-CEX**, Thread 2 first issues the promise e_{22}^2 . Through Thread 3, this promise justifies another branch of Thread 2, namely $e_{11}^2 \rightarrow e_{21}^2 \rightarrow e_{31}^2$, and can be certified thanks to the equivalent write e_{31}^2 . The problem is that in the two branches of Thread 2, the read of z is justified by two different writes. In other words, Thread 2 *bait*s other threads with the promise, assuming that read of z gets value \emptyset , but then *switches* by picking value 1 for this read.

To forbid this kind of behavior, we place an additional consistency constraint enforcing CL. We say that a write w *externally justifies* a read r if it justifies r and does not happen-before r nor justifies some other read that happens before r . For example, in Fig. 4, e_1^1 externally justifies e_{11}^2 and e_2^3 externally justifies e_{21}^2 . We require that whenever an event structure contains two conflicting branches due to some promise, the external justifications of the two branches agree modulo the external justification that created the conflict between the branches.

Definition 3.10. An event structure S is Weakestmo2-consistent if it is Weakestmo-consistent and also:

- $(\text{jf} \setminus (\text{jf}^? ; \text{hb})) ; \text{po} ; \text{ew} \subseteq \text{jf} ; (\text{po} \cup \text{lpmat})$ (NO-BAIT-AND-SWITCH)

where $\text{lpmat} \triangleq \text{cf}_{\text{imm}} ; [\text{rng}(\text{jf} \cap (\text{jf}^? ; \text{hb}))]$; po denotes the *load buffering pattern*.

An execution graph is Weakestmo2-consistent if it can be extracted from some Weakestmo2-consistent event structure.

¹To see why coherence is enforced for the whole event structure rather than on per-execution basis consult Chakraborty et al. [2019, §2.3].

²The full version of the model (see §C) filters out executions that violate atomicity of read-modify-write operations or sequential consistency of sc accesses.

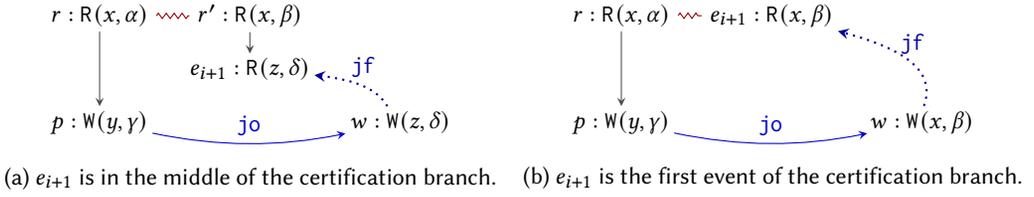


Fig. 5. Illustration to the proof of the LBRF theorem.

The **NO-BAIT-AND-SWITCH** axiom requires that external justifications of one branch either also justify some read in the conflicting branch ($jf; po$) or be the very reason why the branch was created ($jf; lbp$): the externally justified read should be in immediate conflict with a non-externally justified read heading the other branch.

Due to this axiom, the event structure of Fig. 4 is not Weakestmo2-consistent because event e_2^3 externally justifies e_{21}^2 but does not justify any event po -before e_{22}^2 . In contrast, the event structure of Fig. 3 is Weakestmo2-consistent. The only relevant external justification is that of e_{12}^1 by e_2^2 , which is allowed because e_{12}^1 is in immediate conflict with e_{11}^1 , which is po -before e_{21}^1 .

3.3 Load Buffering and Data Race Freedom

We recall the definition of RC11-consistency and show that Weakestmo2 satisfies $LBRF_{RC11}$.

Definition 3.11. An execution graph G is *RC11-consistent* if the following hold:

- $po \cup rf$ is acyclic (NO-THIN-AIR)
- $hb; eco^?$ is irreflexive. (COHERENT)

THEOREM 3.12. Weakestmo2 satisfies $LBRF_{RC11}$. (See Appendix D for the full proof.)

PROOF SKETCH. We introduce a subclass of *promise-free* event structures for which $cf \cap jo = \emptyset$ holds. It is easy to show that for Weakestmo2-consistent promise-free event structure S justified-from relation coincides with reads-from relation: $S.jf = S.rf$. Therefore every extracted execution of a consistent promise-free event structure is RC11-consistent, since $po \cup rf$ acyclicity follows immediately from $po \cup jf$ acyclicity. Thus it suffices to show that every Weakestmo2-consistent event structure of an LB-race-free program is promise-free.

The latter can be shown by induction on the construction of an event structure. That is, given a Weakestmo2-consistent event structure S one can consider a sequence of events $\{e_1, \dots, e_n\}$ of S ordered according to some total order extending $(S.po \cup S.jf)^*$. It is possible to construct S step-by-step by adding single event at each step. Then by induction we can show that each event structure S_i obtained on i -th step is promise-free.

Trivially, empty event structure S_0 is promise free. Also it can be shown that if the event added on $i + 1$ step e_{i+1} is a write or a fence, then the relation $cf \cap jo$ cannot increase and thus the event structure remains promise-free. The only non-trivial case is when e_{i+1} is a read event.

This situation is depicted in Fig. 5a. We have a newly added read event e_{i+1} and a promise p . Read events r and r' are the first events at which two branches of the event structure diverge and become conflicting. The constraint **NO-BAIT-AND-SWITCH** guarantees that the read e_{i+1} cannot observe promise p in the middle of the certification branch, thus it has to be that $e_{i+1} = r'$ (see Fig. 5b). Then it is easy to see that events r and w form a load-buffering race. Both of these events belong to the event structure S_i obtained on the previous step. Because of our inductive assumption, S_i is promise-free. Thus we can extract an RC11-consistent execution containing a load-buffering race,

contradicting our assumption that P is LB-race-free under RC11. Therefore, it has to be that S_{i+1} remains promise-free. \square

Composing $\text{LBRF}_{\text{RC11}}$ with RC11's DRF_{SC} theorem, we get DRF_{SC} for Weakestmo2.

COROLLARY 3.13. *Weakestmo2 satisfies DRF_{SC} .*

3.4 Soundness of Compilation Mappings

One of the main objectives of the advanced multi-execution weak memory models is to enable efficient compilation mappings to the hardware architectures. In this section, we show that our modification of the Weakestmo preserves this property. We prove soundness of the optimal compilation schemes, i.e., those that do not require to insert fences or fake dependencies when compiling relaxed accesses, from Weakestmo2 to memory models of x86, ARMv7, ARMv8, and POWER.

To achieve this goal, we adjust the proof of the compilation correctness for Weakestmo by Moiseenko et al. [2020]. The proof uses the IMM model as a mediator between Weakestmo and the hardware models. Since the correctness of compilation mappings from IMM to hardware models is already established by Podkopaev et al. [2019], it suffices to show correctness of compilation from Weakestmo2 to IMM, which boils down to the following statement [Moiseenko et al. 2020, § 2.3].

THEOREM 3.14. *Let P be a program, and G be an IMM-consistent execution graph of P . Then, there exists an Weakestmo2-consistent event structure S of P such that $S \triangleright G$.*

(See Appendix F for the full proof.)

PROOF SKETCH. To prove the theorem, following the original proof, we construct the required event structure S step by step following a *traversal* of the IMM graph G [Podkopaev et al. 2019, § 6.2]. Traversal of the graph G induces operational small-step semantics $G \vdash TC \xrightarrow{e} TC'$ where TC and TC' are *traversal configurations* and e is an event being traversed. Traversal configuration is a tuple $\langle C, I \rangle$, where $C \subseteq G.E$ is a set of *covered events* and $I \subseteq G.W$ is a set of *issued writes*.

Intuitively, covering an event corresponds to in-order execution of an instruction of the program, while issuing a write corresponds to our-of-order speculative execution of some store. An event can be covered whenever (i) all of its po predecessors are covered and (ii) it is already issued or it reads from an issued write. In order to issue a write, one must first issue all the writes of other threads on which it depends via the preserved program order ppo . These constraints can be manifested as the following invariants of the traversal configuration $\langle C, I \rangle$:

$$\text{dom}(\text{po}; [C]) \subseteq C \quad C \cap W \subseteq I \quad \text{dom}(\text{rf}; [C]) \subseteq I \quad \text{dom}((\text{rf} \setminus \text{po}); \text{ppo}; [I]) \subseteq I$$

Giving the operational semantics of traversal $G \vdash TC \xrightarrow{e} TC'$, and the operational semantics of event structure construction $S \xrightarrow{e} S'$ the proof then proceeds using the standard *simulation* argument. As such, the main challenge of our modification of the proof was to show that the new axiom **NO-BAIT-AND-SWITCH** is preserved during the simulation. It turned out that in order to ensure that we need to slightly modify the construction from Moiseenko et al. [2020].

We demonstrate the problem with the original construction and our key idea on how to repair it with an example. Consider program **LB-IMM** in Fig. 6. Its annotated outcome is allowed by IMM and can be seen as a result of reordering the syntactically independent instructions of Thread 1.

To generate this outcome, the first step of the traversal issues $W(a, 1)$ in Thread 1. To simulate this action, we create the branch $e_{11}^1 \rightarrow e_{21}^1 \rightarrow e_{31}^1 \rightarrow e_{41}^1 \rightarrow e_{51}^1$ (see Fig. 6) using the *receptiveness* property [Podkopaev et al. 2019, § 6.4]. Receptiveness allows us to pick arbitrary values for intermediate read events in the branch if there is no dependency from these reads to the issued write. For each such read, we choose some “stable” justification write [Moiseenko et al. 2020, § 4.3.1]. In this case, the stable justification writes happen to be the initialization writes (omitted on Fig. 6 for brevity).

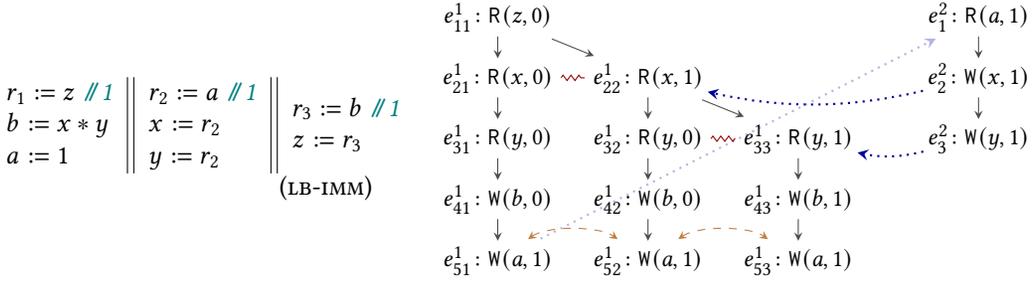


Fig. 6. **LB-IMM** and its partial Weakestmo2-consistent event structure.

The next steps in the traversal cover $R(a, 1)$ and then issue and immediately cover $W(x, 1)$ and $W(y, 1)$ in Thread 2. To match these steps, we add events $e_1^2 \rightarrow e_2^2 \rightarrow e_3^2$ to the event structure.

Subsequently, the traversal issues $W(b, 1)$. At this point, the construction of Moiseenko et al. [2020] adds a branch $e_{22}^1 \rightarrow e_{33}^1 \rightarrow e_{43}^1 \rightarrow e_{53}^1$. It does so by picking suitable justification writes for all reads on which the issued write $W(b, 1)$ depends (via **ppo**). However, changing justification for multiple reads at once in the new branch violates **NO-BAIT-AND-SWITCH**.

We repair the construction by showing that we can replace justification writes for G .**ppo**-preceding reads *incrementally* one-by-one, constructing a series of certification branches, as shown in Fig. 6. That is, we first construct the intermediate branch $e_{22}^1 \rightarrow e_{32}^1 \rightarrow e_{42}^1 \rightarrow e_{52}^1$. Doing so satisfies **NO-BAIT-AND-SWITCH** because the new branch differs from the previous branch only at the point of immediate conflict: $e_{21}^1 \leftrightarrow e_{22}^1$. Starting from this intermediate branch, it becomes possible to add the required branch $e_{33}^1 \rightarrow e_{43}^1 \rightarrow e_{53}^1$ which now has the event e_{43}^1 matching the issued write $W(b, 1)$.

The remaining part of the simulation process is not shown in Fig. 6 because it proceeds unchanged compared to Moiseenko et al. [2020]: first $R(b, 1)$ is covered, then $W(z, 1)$ is issued and covered, and finally all the events of Thread 1 are covered. The corresponding events are added to the event structure in a straightforward way to match these traversal steps, arriving at the final execution graph justifying the annotated outcome. \square

3.5 Soundness of Program Transformations

Another important objective of the advanced weak memory models is to justify source-to-source program transformations applied by optimizing compilers. We next discuss the implications of strengthening Weakestmo for their soundness. A transformation tr from a source program P_{src} to a target program P_{tgt} is *sound* if does not add any new behaviors. That is, for every Weakestmo2-consistent execution G_{tgt} of P_{tgt} , there exists some Weakestmo2-consistent execution G_{src} of P_{src} with the same behavior: $\mathcal{B}(G_{src}) = \mathcal{B}(G_{tgt})$.

In Appendix G, we took the results of Chakraborty et al. [2019, §6.2] for the original version of Weakestmo and showed that all sound **reorderings** and **eliminations** of relaxed loads and stores **remain sound** for Weakestmo2. As an example of our reasoning we give a proof sketch for the soundness of the load/store reordering.

THEOREM 3.15. *The reordering of two adjacent independent instructions $a = (r_1 := x)$ and $b = (y := r_2)$ is a sound source-to-source transformation.*

PROOF SKETCH. The proof proceeds by induction on the construction of the target event structure using the simulation argument of Chakraborty et al. [2019, §F]. Given a Weakestmo2-consistent execution graph G_{tgt} of P_{tgt} , we consider event structure S_{tgt} , s.t. $S_{tgt} \triangleright G_{tgt}$, and we built it

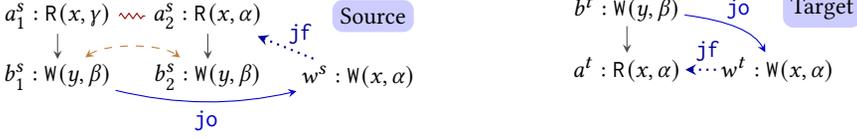


Fig. 7. A fragment of the event structure construction that justifies load/store reordering.

incrementally from the initial event structure: $S_{\text{init}}(P_{\text{tgt}}) \rightarrow^* S_{\text{tgt}}$. Following these steps, we will construct S_{src} and the required graph G_{src} .

To simulate the target event structure construction step $S_{\text{tgt}} \xrightarrow{e} S'_{\text{tgt}}$, if the event added, e , is **not** a result of executing instruction b or a , the source event structure can be augmented by adding the same event $S_{\text{src}} \xrightarrow{e} S'_{\text{src}}$.

Otherwise, let us consult Fig. 7. It depicts a fragment of the source (on the left) and target (on the right) event structures. The target event structure construction adds the event $b^t : W(y, \beta)$ corresponding to the instruction b . In the source, however, one has to execute instruction a first. In doing so, it might be not possible to add an event with the label $R(x, \alpha)$ that will be added later in the target event structure. The corresponding justifying write event might have yet been added to the source event structure because it may, in turn, depend on the write event added as a result of executing b itself (see events b^t and w^t in the target event structure).

The construction of Chakraborty et al. [2019, §F] therefore creates an auxiliary execution branch in the source event structure consisting of events a_1^s and b_1^s . The construction justifies the $a_1^s : R(x, \gamma)$ event by choosing the S_{src} .**co**-maximal non-conflicting write from S_{src} .**jf**? ; S_{src} .**hb** prefix of a_1^s . Since instructions a and b are assumed to be independent, the choice of the value γ for the read a_1^s cannot affect the value of the write $b_1^s : W(y, \beta)$.

The construction then proceeds by adding events to the target and source event structures until the target reaches the event a^t . At this point, the construction adds another branch to the source consisting of events a_2^s and b_2^s . Now, the event a_2^s can have the required label $R(x, \alpha)$ because there is already a justifying event w^s in the source event structure matching the target's justification event w^t . Finally, the write b_2^s is announced to be equal to b_1^s .

What remains to be shown is that the two conflicting branches $a_1^s \rightarrow b_1^s$ and $a_2^s \rightarrow b_2^s$ satisfy the axiom **NO-BAIT-AND-SWITCH**. Indeed, they only differ at the point of the immediate conflict, and moreover, the read a_1^s is justified from a $(S_{\text{src}}$.**jf**? ; S_{src} .**hb**)-preceding write. \square

4 WMC: WEAKESTMO2 MODEL CHECKING

In this section, we present WMC, our model checking algorithm for Weakestmo2, which we build on top of GENMC [Kokologiannakis et al. 2019], an existing state-of-the-art, open-source stateless model checker for RC11 programs. Our algorithm is largely parametric in the underlying memory model, and so can in principle be adapted to other multi-execution memory models satisfying LBRF and CL. We proceed with a brief description of how GENMC operates under RC11 (§4.1), and then describe our extensions for Weakestmo2 (§4.2).

4.1 GENMC: Model Checking under RC11

GENMC, like other *dynamic partial order reduction* (DPOR) algorithms [Abdulla et al. 2014; Flanagan et al. 2005], verifies a program by enumerating its executions one at a time, while recording alternative exploration options along the way. This high-level procedure is depicted in Algorithm 1. (The **highlighted** code represents our extensions for Weakestmo2 and can be ignored for now.)

Algorithm 1 Main exploration algorithm

```

1: procedure VISIT( $P, G, \Pi$ )
2: if  $\neg \text{cons}_{\text{RC11}} \setminus \text{porf}(G)$  then return
3: switch  $a \leftarrow \text{next}_{P, \Pi}(G)$  do
4:   case  $a = \perp$ 
5:     if  $\Pi = \emptyset$  then output “Exec OK”
6:   case  $a \in \text{error}$ 
7:     exit (“Erroneous execution”)
8:   case  $a \in R$ 
9:     for  $w \in \text{GETRFs}(G, \Pi, a)$  do
10:      VISIT( $P, \text{SetRF}(G, w, a), \Pi$ )
11:   case  $a \in W$ 
12:      $\Pi' \leftarrow \{ \langle w, G_w \rangle \in \Pi \mid w \neq a \}$ 
13:     VISIT( $P, G, \Pi'$ )
14:     VISITREVISITS( $P, G, \Pi, \Pi', a$ )
15:   otherwise VISIT( $P, G, \Pi$ )

1: function GETRFs( $G, \Pi, r$ )
2:  $W \leftarrow G.W_{\text{loc}}(r)$ 
3: if  $\Pi \neq \emptyset$  then
4:    $L \leftarrow \{ w' \mid \exists \langle w, G' \rangle \in \Pi. \langle w', w \rangle \in G'.\text{rf}^?; G.\text{hb} \}$ 
5:    $W \leftarrow W \cap L$ 
6: return  $W$ 

```

Algorithm 2 Exploration of revisits

```

1: procedure VISITREVISITS( $P, G, \Pi, \Pi', a$ )
2: if  $\Pi' \neq \emptyset$  then  $\langle Rs, \Pi_c \rangle \leftarrow \langle \emptyset, \emptyset \rangle$  ▷ In cert
3: else if  $\Pi \neq \emptyset$  then ▷ Cert OK
4:    $\langle Rs, \Pi_c \rangle \leftarrow \langle \text{CERTREVS}(G, a), \{a\} \rangle$ 
5: else ▷ Normal exec
6:    $\langle Rs, \Pi_c \rangle \leftarrow \langle \text{GETREVS}(G, a), \emptyset \rangle$ 
7:   for  $\langle w, r \rangle \in Rs$  do
8:      $G' \leftarrow \text{RESTRICT}(G, r, w) \setminus \text{rng}([r]; G.\text{po})$ 
9:      $\Pi' \leftarrow \text{PROMISES}(G, r, \{w\} \cup \Pi_c \cup \text{InCyc}(G, r))$ 
10:    VISIT( $P, G', \Pi'$ )

1: function GETREVS( $G, a$ )
2: return  $\{ \langle a, r \rangle \mid r \in (G.T_{\text{loc}}(a) \setminus \text{dom}(G.\text{porf}; [a])) \cup \text{dom}(G.\text{lbrace}; [a]) \}$ 

1: function PROMISES( $G, r, S$ )
2: return  $\{ \langle w', G \rangle \mid \langle r, w' \rangle \in G.\text{po}; [\text{dom}(G.\text{rfe}) \wedge w' \in \text{dom}(G.\text{porf}; [S])] \}$ 

1: function CERTREVS( $G, e$ )
2: return  $\{ \langle r, w \rangle \mid \langle r, e \rangle \in \text{po} \wedge G.\text{rf}; [r] \subseteq G.\text{rf}^?; G.\text{hb} \wedge w \in G.W_{\text{loc}}(r) \} \cup \bigcup_{w \in \text{dom}([w]; G.\text{po}; [e])} \text{GETREVS}(G, T, w)$ 

```

GENMC’s VISIT procedure explores all consistent executions of a program P under RC11³ recursively. During the exploration, VISIT maintains the current exploration graph G which is augmented with a *revisit set* $G.T$ that records all reads in G whose reads-from edge can be changed. Initially, VISIT is called with an empty execution graph.

At each step, as long as the current graph remains consistent (Line 2), VISIT picks the next event to add, and adds it to the graph using the next function (Line 3).

The role of next is twofold: it schedules a thread t and also adds the next event of t in the graph. If no thread can be scheduled (e.g., if all threads are finished), then the execution is complete, and next returns \perp (Line 4). If an error is encountered (e.g., if an assertion in the program is violated), next returns the error token error (Line 7). Otherwise, it returns the event it added a .

If a is a read, VISIT has to consider all possible *rf* edges for it. To that end, for each possible *rf* option w (given by GETRFs), VISIT recursively calls itself with the graph recording that a read from w (Lines 9 and 10).

If a is a write, apart from simply recursing further (Line 13), VISIT has to also check whether w can *revisit* any of the existing reads in G . This is necessary because, when a read r is added to the graph, it may well be the case that some write from which r could also read from has not yet been added to the graph. Thus, whenever VISIT adds a write a , it also checks whether any of the existing reads can be revisited to read from a , and explores these options via VISITREVISITS (Line 14).

Otherwise, VISIT simply recurses further.

³Adapting the algorithm for a different model m merely requires changing the consistency check in Line 2 of Algorithm 1.

The VISITREVISITS Procedure. The calculation of revisitable reads is performed by VISITREVISITS, by means of GETREVS. As can be seen in Algorithm 2, GETREVS returns a set where a is paired with all revisitable reads to the same location that are not in its `porf` prefix.

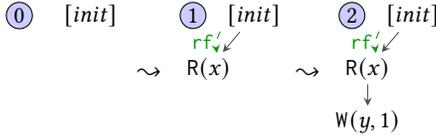
Subsequently (and in accordance to standard DPOR approaches), for each such revisit pair $\langle r, w \rangle$, VISITREVISITS first restricts G (Line 8) so that it only contains the events that were added before r , as well as the events that are `porf`-before w (as w was added after r). Analogously, it restricts $G.T$ so that it only contains reads added before (and including) r . Finally, VISITREVISITS simply calls VISIT recursively to generate the corresponding executions.

GENMC: An Example. Let us now illustrate how GENMC works with an example. Consider a variant of the load buffering program from §1 with no dependencies between instructions.

$$\begin{array}{l} r_1 := x \\ y := 1 \end{array} \parallel \begin{array}{l} r_2 := y \\ x := 1 \end{array} \quad (\text{LB-NODEP})$$

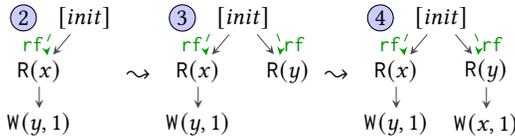
We choose this example because it has 3 consistent executions under RC11 (Figures 1a to 1c) and one additional execution under Weakestmo2. Running GENMC on LB-NODEP highlights the difficulties arising when trying to enumerate the executions of programs with `porf` cycles. In the presentation below, we omit the values read by reads in the graphs, as they can be deduced from the `rf` edges.

GENMC starts with an empty graph and then adds two events corresponding to $r_1 := x$ and $y := 1$, respectively, as can be seen below.



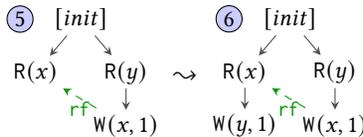
Note that, because there are no other reads-from options for $R(x)$ except for 0, the loop in Line 9 in VISIT will only add one child node to the recursion tree. Similarly, because there are no revisit options for $W(y, 1)$, VISITREVISITS is not executed, and GENMC only calls VISIT in Line 13 for the newly added write.

However, when the read event corresponding to $r_2 := y$ is added in the next step, VISIT starts two recursive explorations: one where $R(y)$ reads 0, and an alternative one where it reads from $W(y, 1)$. Let us assume that VISIT first proceeds with the one where $R(y)$ reads 0.



In a similar manner, when the write corresponding to $x := 1$ is added to the graph, VISIT will initiate two recursive explorations: one where $W(x, 1)$ does not revisit any reads, and one where it revisits $R(x)$. Indeed, because $R(x)$ is *not* `porf`-before $W(x, 1)$, it will be considered by GETREVS, and VISITREVISITS will initiate a recursive exploration.

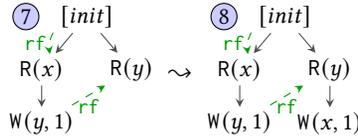
In the non-revisiting case, the graph is complete (corresponding to Fig. 1a), so let us focus on the revisiting case. The graph created for the revisiting exploration is graph 5 below.



In effect, this graph models a scenario where the $W(x, 1)$ was added just before the $R(x)$: the only events that are present in the graph are the ones added before $R(x)$ (i.e., the read itself), as well as those that are absolutely necessary in order to trigger $W(x, 1)$ (here, its po -predecessor). Had we only kept the events that were added before $R(x)$ when revisiting (as VISIT does in the read case), $W(x, 1)$ would not exist upon restriction, and $R(x)$ would not have been able to read from it.

Continuing with the revisiting case, $W(y, 1)$ is added once again to the graph. This time, however, it does not revisit $R(y)$, as the latter is both $porf$ -before $W(y, 1)$, and also not revisitable (recall that T is restricted to only contain events that were added before $R(x)$). The graph is now complete again (corresponding to Fig. 1c), and thus VISIT backtracks and explores the alternative rf for $R(y)$.

The final exploration can be seen below. Since the recursive call corresponding to ⑦ was initiated by the addition of $R(y)$, graph ⑦ is identical to ③ with the only exception being $R(y)$'s rf edge.



In the final step of the algorithm, VISIT adds $W(x, 1)$ again in the graph. Since $R(x)$ is $porf$ before $W(x, 1)$, no reads are considered by GETREVS, and the final exploration is concluded (cf. Fig. 1b).

4.2 WMC: Model Checking under Weakestmo2

We just saw how GENMC explores all three RC11-consistent executions of the LB program. When it comes to generating the execution of Fig. 1d, however, GENMC fails, for two major reasons.

First, GETREVS does not return reads that are $porf$ -before the revisiting write. As such, even though $W(y, 1)$ (resp. $W(x, 1)$) could revisit $R(y)$ (resp. $R(x)$) in executions ⑥ and ⑧ to obtain the cyclic execution, that was impossible due to a $porf$ -path between the read and the write.

Second, revisiting $porf$ -earlier reads (which would seemingly solve the first issue above) is not enough on its own, as such reads may be non-revisitable (as e.g., $R(y)$ in ⑥).

We next show how WMC overcomes these difficulties and generates the execution of Fig. 1d, using the ideas of §2.4. Our WMC extensions are highlighted in Algorithms 1 and 2.

4.2.1 WMC: Overview. In order to generate LB behaviors, the first thing that needs to be changed is the function $GETREVS(G, a)$. Besides the revisitable reads that are not $porf$ -before a (as in GENMC's case), WMC also returns any reads that are $porf$ -before a , as long as they are in an LB race with a :

$$lbrace \triangleq ([R] ; =_{loc} ; [W]) \cap (G.rpo ; (G.rf \setminus G.po) ; G.porf) \setminus (hb \cup hb^{-1})$$

Put differently, GETREVS tries to create executions with LB cycles only when an LB race is detected.

That said, simply revisiting $porf$ -prior events is not a sound way of generating executions with LB cycles. The problem is that when a write w revisits a $porf$ -earlier read r , the graph that RESTRICT would create also includes r 's po -suffix. This po -suffix *has* to be removed from the graph, since its very existence may depend on the value read by r (e.g., due to control flow). Yet, some events of the po -suffix do need to be re-added in the graph if $porf$ -cyclic executions are to be generated.

WMC resolves this problem with a two-step approach. As a first step, when a write w revisits a $porf$ -earlier read r , WMC removes the $porf$ -prefix of w that is po -after r (Line 8). The writes that are po -after r and are read externally are kept in a *promise set* calculated by PROMISES (Line 9), which is in turn used during the recursive exploration (Line 10). Promise sets, Π , are sets of pairs consisting of the write that is promised and the execution graph when the promise was issued (the latter is used to constrain the reads within promise certifications).

As a second step, whenever VISIT encounters a non-empty promise set, WMC initiates a *certification phase*. During this phase, WMC operates in a restricted mode: all promises of Π have

to be fulfilled (i.e., the corresponding writes have to be re-added to the graph), and all non-local explorations (see §2.3) are postponed until the certification phase succeeds.

To that end, WMC modifies GENMC's original algorithm in the following ways.

First, the next function is changed so that when WMC is in a certification phase (i.e., $\Pi \neq \emptyset$), next returns the events of the thread under certification until either all promises have been fulfilled, or all the events of this thread have been added to the graph (in which case next returns \perp).

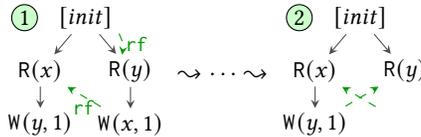
Second, the GETRFs function is modified to constrain the possible reads-from options during the certification phase (Line 3). The modified version of the function calculates the set of local writes that were used to issue promises (Line 4). Reading from these writes is safe, since it cannot lead to bait-and-switch behaviors. All other (non-local) reads-from options will be considered only after the certification procedure has succeeded (see below).

Third, VISITREVISITS is changed to behave differently under certification. After VISIT removes any promise in Π fulfilled by a (Line 12), VISITREVISITS checks whether the certification is over. If the certification is not over yet, no revisits are performed (Line 2). (Analogously to GETRFs, a 's revisits will be calculated later once the certification procedure completes successfully.) If the certification is over (i.e., all promises have been fulfilled), then all *rf* options that were discarded by GETRFs and all revisits that were skipped by VISITREVISITS will now be considered (Line 4). To that end, VISITREVISITS calculates all non-local options for the reads and writes of the thread that completed the certification (by means of CERTREVS; Line 4), and then recursively explores them.

Arguably, the most intricate parts in the changes described above are (1) the calculation of non-local options at the end of a certification, and (2) the calculation of the promise set for a recursive exploration. But, before diving into these parts, let us see an example of WMC in action.

4.2.2 WMC: An Example. Using the modifications above, WMC explores all 4 executions of Fig. 1.

Initially, the exploration remains the same as with GENMC. In graph ⑥, however (and in contrast to GENMC), WMC also considers $R(y)$ to be revisited by $W(y, 1)$, eventually leading to graph ② below, where $W(x, 1)$ needs to be certified:



Let us now see what happens when VISIT proceeds with exploration corresponding to graph ②. Since Π is non-empty, WMC enters a certification phase. The next event to be added is $W(x, 1)$, which fulfills the promised write of the first thread. Since no revisits were skipped (and no *rf* edges were restricted) during the certification phase, CERTREVS returns the empty set, and the execution (corresponding to Fig. 1d) is complete.

The rest of the exploration proceeds in a similar manner. When WMC encounters graph ⑧, it will notice that graph ⑧ contains an LB race, and therefore generate another (duplicate) execution with the weak LB behavior⁴. In general, however, this is not something WMC could have predicted, and thus has no way of avoiding it without doing some extra bookkeeping (see below).

We conclude this example with two remarks.

Blocked Executions. While in the **LB-NODEP** program above all promises could be fulfilled, this is not the case in general. In the **LB+CTRL** program below, whenever a read gets revisited in hope that a cyclic execution will be generated (as in execution ② above), its promise cannot be fulfilled because of the control dependency between the read and its subsequent write. Although there can

⁴The exploration is similar to the one presented for ② above, and is thus omitted for brevity.

be a fair number of blocked executions in a test case, the overhead that is induced on WMC by such explorations is modest because WMC does not fully explore these blocked executions: it discards them as soon as the certification phase fails.

$$\begin{array}{l} r_1 := x \\ \text{if } r_1 = 0 \text{ then} \\ \quad y := 1 \end{array} \parallel \begin{array}{l} r_2 := y \\ \text{if } r_2 = 0 \text{ then} \\ \quad x := 1 \end{array} \quad (\text{LB+CTRL}) \qquad \begin{array}{l} r_1 := x \\ y := 1 \end{array} \parallel \begin{array}{l} r_2 := z \\ r_3 := y \\ x := 1 \end{array} \parallel z := 2 \qquad (\text{RLB+W})$$

Duplicate Executions. One may wonder what does WMC do about duplicate executions in general. For LB-race-free programs, WMC essentially follows GENMC: instead of exploring the executions of a program recursively, WMC uses a stack and some auxiliary structures that allow it to not encounter any duplication at the cost of some extra memory (see [Kokologiannakis et al. 2019]).

For LB-racy programs like **LB-NODEP**, however, it is not possible to avoid exploring duplicate executions simply by recording the LB cycles that have already been encountered in the exploration. Indeed, it turns out this naive approach of recording LB cycles is not sound, as it may lead to the disposal of valid executions. To see this, consider the **RLB+W** program above and assume that we record the encountered cycles. Note that **RLB+W** has two executions with LB cycles: one where $r_2 = 0$, and one where $r_2 = 2$. Assuming that events are added from left to right and that we first encounter the cyclic execution where $r_2 = 0$, when we encounter the second execution where $r_2 = 2$, it will be erroneously disposed, as it involves the exact same cycle as the one with $r_2 = 0$.

A sound way to avoid duplication is to record the whole **porf**-prefix of such LB cycles. Although this can amount to recording complete execution graphs, it does not induce a significant space overhead in practice because executions with LB cycles are rare.

4.2.3 WMC: Promises and Non-Local Revisits. Finally, let us now return to the last remaining parts of VISITREVISITS: the calculation of non-local options at the end of a certification, and the calculation of the promise set Π .

Starting with the calculation of non-local options, when the certification is over, we have to a) calculate alternative **rf** edges for the reads that read locally during certification, and b) calculate revisits for writes the revisits of which were skipped during certification. As can be seen in Algorithm 2, CERTREVS performs exactly these two actions.

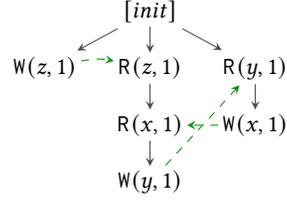
Continuing with the calculation of the promise set, given a revisit $\langle w, r \rangle$, the definition of PROMISES simply returns the writes after **po**-after r that are read externally, and are **porf**-before its last parameter S .

The only question that remains to be answered is *what* should be passed as the S argument of PROMISES (cf. VISITREVISITS, Line 9). Clearly, one thing that should be passed as part of S is the revisiting write w indeed, this is the first component passed to PROMISES. Passing just w , however, is not enough. In fact, there are two more cases we have to take into account.

As an example of the first case, consider the end of a successful certification phase triggered by a revisit $\langle w, r \rangle$, as well as a revisit $\langle w', r' \rangle$ returned by CERTREVS at the certification end. Let us further assume that the revisit $\langle w', r' \rangle$ initiates a new certification phase. In such cases, it is inadequate to consider as promises only the writes that are **porf**-before w' . Since w' is itself **porf**-before w (and part of the cycle that the revisit $\langle w, r \rangle$ created), we also have to include w in S in order to ensure that w 's cycle will be preserved; that is the role of S 's second component, Π_c . (Note that revisiting r' from w' without preserving the cycle of $\langle w, r \rangle$ is also possible, but that execution will be obtained in another graph without the $\langle w, r \rangle$ cycle.)

As an example of the second case, consider the program below along with its Weakestmo2-consistent execution where all reads read 1:

$$z := 1 \left\| \begin{array}{l} r := z \\ s := x \\ \text{if } r = 0 \vee s = 1 \text{ then} \\ \quad y := 1 \end{array} \right\| \left\| \begin{array}{l} a := y \\ x := a \end{array} \right.$$



Perhaps surprisingly, it is impossible to get the above graph if we do not revisit the read of z after the LB-cycle between x and y has been created. Indeed, let us assume that WMC executes the program in a left to right manner. When the second thread is executed, $r := z$ can read either 0 or 1, while $s := x$ can only read 0. Given these options, the write $y := 1$ (and, by extension, the LB cycle) will only appear in the exploration where $r := z$ reads 0. In other words, while it is consistent to have both reads of the second thread reading 1, it is impossible to arrive at this scenario after the first read reads 1.

To account for this problem, whenever a certification is complete and an LB cycle is created, we have to (1) revisit reads in the certification thread that are po -before the cycle, and (2) ensure that the cycle will continue to exist after these revisits take place. Luckily, our definition of `CERTREVS` already alleviates the first issue: instead of considering revisits just for reads that participate in the cycle, it also considers revisits of reads that are po -before the cycle. To solve the second issue (i.e., guarantee that the cycle will exist after these revisits take place), we use the `INCyc` function, which returns all writes in the certification thread that are po -after its argument, and also participate in the cycle⁵.

5 EVALUATION

We evaluate WMC by answering the following questions:

- §5.1 How often do LB races appear in practice? What overhead is there for detecting them?
- §5.2 How well does WMC perform against other tools that handle similar memory models?
- §5.3 How does WMC scale in synthetic benchmarks containing many load buffering races?
- §5.4 Can we use WMC to automatically remove load buffering races?

To do so, we compare WMC against the following tools.

- GENMC [Kokologiannakis et al. 2019; 2021] is the stateless model checker that we build upon. It supports the RC11 memory model, which does not permit LB behaviors.
- HMC [Kokologiannakis et al. 2020] is an extension of GENMC that verifies programs under IMM [Podkopaev et al. 2019], a weak memory model that allows certain LB behaviors by keeping track of dependencies between instructions and forbidding cycles consisting solely of dependencies and `rf` edges.
- HMC_{LBRF} is a variant of HMC that we implemented in order to further demonstrate the benefits of LBRF. HMC_{LBRF} leverages LBRF and starts calculating dependencies only when it finds a load buffering race.
- NIDHUGG [Abdulla et al. 2015a; 2016] is a stateless model checker that, among others, supports the POWER memory model by tracking dependencies (similarly to HMC).
- CDSCHECKER [Norris et al. 2013] verifies programs under an informally defined strengthening of the original C11 model, which forbids thin-air behaviors using a notion of promises.
- RMEM [rmem 2009] is a memory model simulator that supports a number of hardware memory models. In our benchmarks, we employ the Promising-Arm model [Pulte et al. 2019] because

⁵As a further optimization, `INCyc` returns the empty set if its argument is reading non-locally. We do not have to consider non-local reads, as such reads will be local in another exploration.

rMEM usually runs much faster with Promising-Arm than with any of its other supported models that allow LB. As the name suggests, Promising-Arm also uses a notion of promises to induce load-buffering behaviors while forbidding thin-air behaviors.

In summary, we observe that LB races are fairly rare in our set of non-synthetic benchmarks and that exploiting LBRF and CL makes WMC a very effective model checker that has a very small average overhead over GENMC and outperforms the other tools.

Experimental Setup. We conducted all experiments on a Dell PowerEdge M620 blade system, with two Intel Xeon E5-2667 v2 CPUs (8 cores @ 3.3 GHz) and 256GB of RAM, running a custom Debian-based distribution. We used LLVM 7 for HMC (v0.5), GENMC (v0.5), and NIDHUGG (v0.3), commit #da671f7 for CDSCHECKER, and commit #85c8130 for rMEM (v0.1). All reported times are in seconds, unless explicitly noted otherwise. We set the timeout limit to 30 minutes.

5.1 Load Buffering Races and WMC's Overhead

To measure the LB race detection overhead, we conducted two case studies.

In the first case study, we took the lock implementations used by Oberhauser et al. [2021] (13 in total), 6 queue implementations used by GENMC, and 10 data-structure benchmarks used by Ou et al. [2018]⁶. Running WMC on these benchmarks confirmed our expectation that realistic implementations rarely contain LB races: out of 29 implementations, we found LB races only in 2. One of them was due to the porting of a non-C11-compliant queue to C11, while the other was an intentional race part of a lock implementation (`musl_lock`), that could not lead to LB behaviors.

In the second case study, we took the entire GENMC benchmark suite (241 tests), which is a combination of small litmus tests and larger concurrent programs. We split these programs into two categories: those that have LB races (28 tests, mostly litmus tests) and those that do not (213 tests). The results are shown in Table 1.

As far as WMC is concerned, detecting LB races imposes roughly 25% overhead over GENMC, which is substantially lower than the overhead of calculating dependencies in HMC, especially when no LB races are present.

As far as HMC_{LBRF} is concerned, observe that HMC_{LBRF} significantly improves the running time of HMC on benchmarks without LB races, as it effectively runs WMC, and imposes a negligible overhead with respect to HMC on benchmarks with LB races. The reason for the latter that it typically detects LB races very quickly, in a fraction of the time needed for verifying the program.

In addition, also observe that HMC_{LBRF} performs really close to GENMC and WMC in benchmarks without LB races. The reason for the additional slowdown compared to GENMC and WMC is that HMC_{LBRF} has to maintain some extra data structures which will be necessary if it has to fall back to dependency tracking.

5.2 Comparison with Other Model Checkers

In order to compare WMC and HMC_{LBRF} with other model checkers, we use both synthetic benchmarks and more realistic data structure benchmarks from the literature.

Synthetic Benchmarks. We extracted synthetic benchmarks from SV-COMP [2019] (pthread and pthread-atomic categories) and rewrote them to utilize weakly-ordered atomic accesses so as to contain LB races (see Table 2).

⁶We could not port all 43 of their benchmarks as they are written in C++, only a subset of which is supported by GENMC.

Table 2. Synthetic benchmarks adapted from SV-COMP [2019]

	GENMC	HMC	HMC _{LBRF}	NIDHUGG	RMEM	CDSCHECKER	WMC
reorder(2)	0.06	0.11	0.07	1.37	77.10	0.04	0.06
singleton	0.01	0.01	0.01	0.12	96.21	0.00	0.01
fib_bench	13.09	25.98	26.68	237.73	⊙	⊙	37.66
szymanski(1)	0.05	0.07	0.08	0.13	5.56	0.28	0.06
szymanski(2)	246.23	399.79	402.34	3.74	529.85	1688.60	321.23
szymanski(3)	⊙	⊙	⊙	170.88	⊙	⊙	⊙
peterson(10)	0.03	0.06	0.07	0.65	⊙	0.25	0.14
peterson(20)	0.12	0.31	0.30	3.81	⊙	1.36	1.09
peterson(30)	0.32	0.84	0.85	13.54	⊙	3.86	4.43
dekker	0.01	0.02	0.02	0.14	30.09	0.23	0.02
sigma	141.56	352.86	168.94		⊙	⊙	157.91
indexer(12)	0.02	0.02	0.02		⊙	0.90	0.02
indexer(13)	0.05	0.07	0.07		⊙	116.71	0.05
indexer(14)	0.30	0.48	0.50		⊙	⊙	0.34
indexer(15)	2.45	4.03	4.07		⊙	⊙	2.85

In a nutshell, throughout Table 2, RMEM and NIDHUGG are generally much slower than all other tools: RMEM maintains a total order of stores across all different memory locations, while NIDHUGG requires expensive consistency checks at each program step. We also note that NIDHUGG does not support RMW instructions and thus some entries in Table 2 are blank.

In addition, as also noted in § 5.1, HMC is approximately 2x slower than GENMC; HMC_{LBRF}, however, does not always perform that much worse compared to GENMC. More specifically, when LB races are present (e.g., for fib_bench, szymanski and peterson), HMC_{LBRF} performs similarly to HMC. When LB races are not present though (e.g., for sigma), HMC_{LBRF} outperforms HMC and performs similarly to GENMC, providing us a first testament to LBRF’s usefulness.

Let us now move to a more detailed comparison between the different tools.

Starting from the upper part of Table 2, we observe that CDSCHECKER is faster than all other tools for the first two benchmarks. This is attributed to two factors: (1) CDSCHECKER operates on binaries thus avoiding the interpretation overhead that GENMC, HMC, and NIDHUGG have to face, and (2) CDSCHECKER utilizes a coarser equivalence partitioning in its DPOR that does not totally order the stores of each memory location; that way, CDSCHECKER explores fewer executions than the other tools. Although GENMC, HMC, and WMC can also operate under a similar partitioning, we chose to not use it, so that WMC better reflects the model presented in § 3.

For the next three benchmarks, however, the situation is reversed: CDSCHECKER performs much worse compared to GENMC, and HMC, despite its coarser equivalence partitioning. In fact, for szymanski, both GENMC and HMC are faster than CDSCHECKER, even though they explore two orders of magnitude more executions than CDSCHECKER. This big disparity in running times is due to CDSCHECKER exploring a large number of infeasible executions, caused by unfulfilled promises.

Nevertheless, WMC’s performance is not on par with that of GENMC and HMC for the same three benchmarks, as one might have hoped: WMC is slower than GENMC in all three benchmarks, while it manages to outperform HMC only for szymanski. For peterson and szymanski WMC’s performance is not attributed to scalability limitations, but rather to Weakestmo2: for these benchmarks the model allows for more executions compared to the models of the other tools. Still, for szymanski, WMC is faster than HMC/HMC_{LBRF} despite the fact that it explores approximately 18% more executions, as calculating dependencies proves to be more expensive. For fib_bench, the slowdown WMC experiences is due to the many promises that remain unfulfilled: CDSCHECKER suffers from the same problem in fib_bench, but does not even manage to terminate within the time limit, as it explores all possible promises, and not just those dictated by LBRF.

Table 3. Benchmarks adapted from Norris et al. [2013]

	GENMC	HMC	HMC _{LB} RF	CDS _{CHECKER}	WMC
barrier-bnd	1.29	2.67	1.60	4.13	1.34
mpmc-queue-bnd	1.87	4.18	2.14	19.78	2.09
treiber-stack-bnd	0.56	1.28	0.65	⊙	0.58
linuxrwlocks-bnd	0.29	0.57	0.58	⊙	0.46
seqlock-bnd	45.40	109.24	111.21	⊙	47.78

Finally, we mention in passing that NIDHUGG is faster than GENMC, HMC and WMC in `szymanski` due to the way the latter tools handle SC fences. These tools treat SC accesses and fences as release/acquire, as part of an optimization. Thus, they explore 5 orders of magnitude more executions for `szymanski` compared to NIDHUGG, resulting in worse performance for this benchmark.

Moving on to the lower part of Table 2, WMC outperforms CDS_{CHECKER} by a large margin. Take `indexer`, for instance: CDS_{CHECKER} explores 4 orders of magnitude more infeasible executions than consistent executions, which makes it much slower compared to GENMC, HMC and WMC.

We conclude the discussion for this table with an observation. While WMC was outperformed by HMC and HMC_{LB}RF in a few cases, benchmarks where that happens employ little to no synchronization: this is extremely uncommon for realistic benchmarks, since using relaxed accesses with no other synchronization gives no guarantees. On the other hand, in cases where synchronization is purposefully employed (e.g., `dekker`), or where there are no LB races (e.g., `indexer`), WMC greatly outperforms all other tools that handle similar memory models.

Data Structure Benchmarks. We next consider some more realistic benchmarks (cf. Table 3) from Norris et al. [2013], whose loops have been manually unrolled to ensure fairness across tools. We exclude NIDHUGG from the comparison because it does not support RMW instructions under POWER, as well as RMEM because it is difficult to encode these benchmarks in its input language.

The observations here are similar to the ones we made for Table 2. CDS_{CHECKER} performs worse than all other tools due to exploring too many infeasible executions. For the first three benchmarks where there are no LB races, WMC outperforms all other tools operating under similar memory models. In addition, HMC_{LB}RF outperforms HMC, as it operates under RC11. However, even in the last two benchmarks where there are LB races, WMC outperforms all other tools, as it avoids the overhead of tracking dependencies. In the case of `linuxrwlocks` specifically, WMC does have some overhead due to LB races and promises that remain unfulfilled, but it still remains very competitive. Again note that the overhead of detecting LB races by HMC_{LB}RF over HMC is negligible in these examples because an LB race is found in the first few executions.

5.3 Load-Buffering Benchmarks

We next evaluate WMC's performance on synthetic test cases with many LB races, and hence potentially many duplicate executions (cf. Table 4). Although such patterns appear rarely in non-adversarial programs, it is pedagogical to examine how WMC performs in such cases.

Test cases `LB+ctrl(N)` and `LB+data(N)` are similar to program `LB` from §1, except that in these tests the `porf` cycle spans N threads. Also, `LB+ctrl(N)` has control dependencies between instructions in place of data dependencies as in `LB+data(N)` and `LB`. Test case `LB-nodep(N)` resembles litmus test `LB-NODEP`, but, similarly, its `porf` cycle spans N threads. Finally, `LB-pairs(N)` contains $N/2$ independent pairs of threads with each pair constituting a load buffering pattern with no dependencies.

We used two versions of WMC (cf. first two distinct columns of Table 4): one where duplicate executions are fully explored, and one where executions that will lead to duplicates are blocked as soon as they are detected, using the technique of §4.

WMC explores the same number of unique executions for all benchmarks of Table 4 except for LB+ctrl; the latter has control dependencies that preclude LB executions. Except for LB-pairs, the percentage of duplicate/blocked executions remains very low or zero. For LB-pairs, memorizing cyclic executions that have been explored and pruning unnecessary exploration prevents the blocked executions outgrowing the number of consistent executions and thus dominating verification time.

5.4 Automatically Fixing Load-Buffering Races

Finally, we investigate how easily LB races can be eliminated. For that, we constructed a script that tries to automatically repair a test case by strengthening the access mode of some instructions in the race’s `porf` path. The script does that incrementally: when an LB race is detected by WMC, the script modifies the test’s source code by strengthening the access mode of two instructions constituting an `rf` edge in the race’s `porf` path, and then runs WMC again until no race exists.

We ran our script on all 30 benchmarks with LB races from §5.1. In most cases (21/30), our script eliminated the LB races with a single iteration. Four tests needed two iterations to be repaired, while one test (`szymanski`) required 10 iterations. The reason is that `szymanski` employs very little synchronization in the form of SC fences. As such, it contains many LB races that need to be repaired. Finally, there were 4 benchmarks which the script was unable to repair: since our script performs syntactic transformations to the test’s source code, it cannot repair tests that use custom primitives for shared memory accesses.

6 RELATED WORK AND CONCLUSION

There has been a lot of work on developing weak memory models for Java/C/C++ that resolve the “out-of-thin-air” (OOTA) problem [Chakraborty et al. 2019; Jagadeesan et al. 2020; Jeffrey et al. 2016; Kang et al. 2017; Lee et al. 2020; Manson et al. 2005; Paviotti et al. 2020; Pichon-Pharabod et al. 2016] with more recent solutions typically criticizing the earlier ones because of their outcomes on certain litmus tests or the authors’ stylistic preferences. Given the lack of an agreed formal criterion for classifying OOTA behaviors, load buffering race freedom and certification locality can be seen as more rigorous appraisals of weak memory models. That said, although LBRF is defined in a model-independent fashion, CL is much more tied to the certification mechanism that is present in the Promising [Kang et al. 2017; Lee et al. 2020] and Weakestmo [Chakraborty et al. 2019] models, and may not be easy to apply to multi-execution memory models with a very different definitional structure. Somewhat surprisingly, both Promising and Weakestmo violate both LBRF and CL; yet, as we have shown, Weakestmo can be adapted to satisfy these properties, while maintaining correctness of the compilation schemes to hardware architectures.

Recently, Jagadeesan et al. [2020] developed a very interesting multi-execution memory model that allows LB but forbids bait-and-switch behaviors, as well as a temporal logic that can be used to dismiss the dubious outcome of `LBF-CEX`. We expect that their model satisfies LBRF and probably also CL, but have not yet investigated a proof of these conjectures.

Table 4. Load buffering benchmarks

	# Execs	Dupls	Blocked
LB+ctrl(10)	11	0	0
LB+ctrl(12)	13	0	0
LB+ctrl(14)	15	0	0
LB+data(10)	1024	0	0
LB+data(12)	4096	0	0
LB+data(14)	16 384	0	0
LB-noddep(10)	1024	0.9%	0.9%
LB-noddep(12)	4096	0.3%	0.3%
LB-noddep(14)	16 384	0.1%	0.1%
LB-pairs(10)	1024	205.2%	33.3%
LB-pairs(12)	4096	281.5%	33.3%
LB-pairs(14)	16 384	376.8%	33.3%

In terms of program verification under weak memory models, there is a significant body of work for `porf`-acyclic models, such as TSO and RC11. These range from program logics for manual verification (e.g., [Doko et al. 2016; 2017; Kaiser et al. 2017; Ridge 2010; Sieczkowski et al. 2015; Turon et al. 2014]) to model checking tools for safety verification (e.g., [Abdulla et al. 2015a; 2018; Barnat et al. 2013; Demsky et al. 2015; Huang et al. 2016; Kokologiannakis et al. 2017; 2019]) and fence insertion tools for enforcing robustness (e.g., [Abdulla et al. 2015b; Bouajjani et al. 2013]). As discussed in §2.2, LBRF provides a simple mechanism for programmers to use these results on weaker memory models at the cost of a few extra fences. Similarly, there are a number of papers that can handle dependency-tracking models, such as ARM and POWER (e.g., [Abdulla et al. 2016; Alglave et al. 2017; 2013; Kokologiannakis et al. 2020; Pulte et al. 2019]).

By contrast, there is hardly any work on verifying concurrent programs on multi-execution memory models: Svendsen et al. [2018] develop a program logic for Promising, which is considerably simpler than corresponding logics for TSO and RC11 [Doko et al. 2016; 2017; Kaiser et al. 2017; Sieczkowski et al. 2015; Turon et al. 2014], and CDSHECKER [Norris et al. 2013] is a model checker for an earlier adaptation of the C11 model with promises similar to those in Promising. As discussed, the reason for this lack of work is that multi-execution models are substantially more complicated and not very amenable to automated formal verification.

The notions we introduced in this work—load buffering race freedom and certification locality—are first steps towards enabling formal verification for such multi-execution models, and indeed have been instrumental in the design of WMC. We hope that they will lead to further exploration in this space, and that they will also prove useful for program logics.

ACKNOWLEDGMENTS

We would like to thank Soham Chakraborty and Anton Podkopaev for their insights on Weakestmo and IMM memory models, and the anonymous reviewers for their valuable feedback. This work was supported by a European Research Council (ERC) Consolidator Grant for the project “PERSIST” under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101003349).

REFERENCES

- Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonardsson, and Konstantinos Sagonas (2015a). “Stateless model checking for TSO and PSO.” In: *TACAS 2015*. Vol. 9035. LNCS. Berlin, Heidelberg: Springer, pp. 353–367. doi: [10.1007/978-3-662-46681-0_28](https://doi.org/10.1007/978-3-662-46681-0_28).
- Parosh Aziz Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas (2014). “Optimal dynamic partial order reduction.” In: *POPL 2014*. New York, NY, USA: ACM, pp. 373–384. doi: [10.1145/2535838.2535845](https://doi.org/10.1145/2535838.2535845).
- Parosh Aziz Abdulla, Mohamed Faouzi Atig, Adwait Godbole, S. Krishna, and Viktor Vafeiadis (2021). “The Decidability of Verification under PS 2.0.” In: *ESOP 2021*. Ed. by Nobuko Yoshida. Cham: Springer International Publishing, pp. 1–29.
- Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Carl Leonardsson (2016). “Stateless model checking for POWER.” In: *CAV 2016*. Vol. 9780. LNCS. Berlin, Heidelberg: Springer, pp. 134–156. doi: [10.1007/978-3-319-41540-6_8](https://doi.org/10.1007/978-3-319-41540-6_8).
- Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Tuan Phong Ngo (Oct. 2018). “Optimal stateless model checking under the release-acquire semantics.” In: *Proc. ACM Program. Lang.* 2.OOPSLA, 135:1–135:29. doi: [10.1145/3276505](https://doi.org/10.1145/3276505).
- Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Ngo Tuan Phong (2015b). “The best of both worlds: Trading efficiency and optimality in fence insertion for TSO.” In: *ESOP 2015*. Vol. 9032. LNCS. Springer, pp. 308–332. doi: [10.1007/978-3-662-46669-8_13](https://doi.org/10.1007/978-3-662-46669-8_13).
- Sarita V. Adve and Kourosh Gharachorloo (Dec. 1996). “Shared memory consistency models: A tutorial.” In: *IEEE Comput.* 29.12, pp. 66–76.
- Jade Alglave and Patrick Cousot (2017). “Ogre and Pythia: an invariance proof method for weak consistency models.” In: *POPL 2017*. Ed. by Giuseppe Castagna and Andrew D. Gordon. ACM, pp. 3–18. url: <http://dl.acm.org/citation.cfm?id=3009883>.

- Jade Alglave, Daniel Kroening, and Michael Tautschnig (2013). “Partial orders for efficient bounded model checking of concurrent software.” In: *CAV 2013*. Vol. 8044. LNCS. Berlin, Heidelberg: Springer, pp. 141–157. doi: [10.1007/978-3-642-39799-8_9](https://doi.org/10.1007/978-3-642-39799-8_9).
- Jade Alglave, Luc Maranget, and Michael Tautschnig (July 2014). “Herding cats: Modelling, simulation, testing, and data mining for weak memory.” In: *ACM Trans. Program. Lang. Syst.* 36.2, 7:1–7:74. doi: [10.1145/2627752](https://doi.org/10.1145/2627752).
- J. Barnat, L. Brim, and V. Havel (July 2013). “LTL model checking of parallel programs with under-approximated TSO memory model.” In: *ACSD 2013*, pp. 51–59. doi: [10.1109/ACSD.2013.8](https://doi.org/10.1109/ACSD.2013.8).
- Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod, and Peter Sewell (2015). “The problem of programming language concurrency semantics.” In: *ESOP 2015*. Vol. 9032. LNCS. Berlin, Heidelberg: Springer, pp. 283–307. URL: http://dx.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 2011*. Austin, Texas, USA: ACM, pp. 55–66. doi: [10.1145/1926385.1926394](https://doi.org/10.1145/1926385.1926394).
- Ahmed Bouajjani, Egor Derevenet, and Roland Meyer (2013). “Checking and enforcing robustness against TSO.” In: *ESOP 2013*. Vol. 7792. LNCS. Springer, pp. 533–553.
- Soham Chakraborty and Viktor Vafeiadis (Jan. 2019). “Grounding thin-air reads with event structures.” In: *Proc. ACM Program. Lang.* 3.POPL, 70:1–70:28. doi: [10.1145/3290383](https://doi.org/10.1145/3290383).
- Brian Demsky and Patrick Lam (2015). “SATCheck: SAT-directed stateless model checking for SC and TSO.” In: *OOPSLA 2015*. Pittsburgh, PA, USA: ACM, pp. 20–36. doi: [10.1145/2814270.2814297](https://doi.org/10.1145/2814270.2814297).
- Marko Doko and Viktor Vafeiadis (2016). “A Program Logic for C11 Memory Fences.” In: *VMCAI 2016*. Ed. by Barbara Jobstmann and K. Rustan M. Leino. Vol. 9583. LNCS. Springer, pp. 413–430. doi: [10.1007/978-3-662-49122-5_20](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*. Ed. by Hongseok Yang. Vol. 10201. LNCS. Springer, pp. 448–475. doi: [10.1007/978-3-662-54434-1_17](https://doi.org/10.1007/978-3-662-54434-1_17).
- Cormac Flanagan and Patrice Godefroid (2005). “Dynamic partial-order reduction for model checking software.” In: *POPL 2005*. New York, NY, USA: ACM, pp. 110–121. doi: [10.1145/1040305.1040315](https://doi.org/10.1145/1040305.1040315).
- Shiyong Huang and Jeff Huang (2016). “Maximal Causality Reduction for TSO and PSO.” In: *CONF_OOPSLA 2016*. New York, NY, USA: ACM, pp. 447–461. doi: [10.1145/2983990.2984025](https://doi.org/10.1145/2983990.2984025).
- Radha Jagadeesan, Alan Jeffrey, and James Riely (Nov. 2020). “Pomsets with preconditions: A simple model of relaxed memory.” In: *JNL_PACMPL* 4.OOPSLA. doi: [10.1145/3428262](https://doi.org/10.1145/3428262).
- Alan Jeffrey and James Riely (2016). “On thin air reads: Towards an event structures model of relaxed memory.” In: *CONF_LICS 2016*. New York, NY, USA: ACM, pp. 759–767. doi: [10.1145/2933575.2934536](https://doi.org/10.1145/2933575.2934536).
- Alan Jeffrey, James Riely, Mark Batty, Simon Cooksey, Ilya Kaysin, and Anton Podkopaev (2022). “The leaky semicolon: Compositional semantic dependencies for relaxed-memory concurrency.” In: *JNL_PACMPL* 6.POPL, pp. 1–30. doi: [10.1145/3498716](https://doi.org/10.1145/3498716).
- 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: *CONF_ECOOP 2017*. Vol. 74. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 17:1–17:29. doi: [10.4230/LIPIcs.ECOOP.2017.17](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: *CONF_POPL 2017*. Paris, France: ACM, pp. 175–189. doi: [10.1145/3009837.3009850](https://doi.org/10.1145/3009837.3009850).
- Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis (Dec. 2017). “Effective stateless model checking for C/C++ concurrency.” In: *Proc. ACM Program. Lang.* 2.POPL, 17:1–17:32. doi: [10.1145/3158105](https://doi.org/10.1145/3158105).
- Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis (2019). “Model checking for weakly consistent libraries.” In: *PLDI 2019*. New York, NY, USA: ACM. doi: [10.1145/3314221.3314609](https://doi.org/10.1145/3314221.3314609).
- Michalis Kokologiannakis and Viktor Vafeiadis (2020). “HMC: Model checking for hardware memory models.” In: *ASPLOS 2020*. ASPLOS ’20. Lausanne, Switzerland: ACM, pp. 1157–1171. doi: [10.1145/3373376.3378480](https://doi.org/10.1145/3373376.3378480).
- Michalis Kokologiannakis and Viktor Vafeiadis (2021). “GenMC: A model checker for weak memory models.” In: *CAV 2021*. Ed. by Alexandra Silva and K. Rustan M. Leino. Vol. 12759. LNCS. Springer, pp. 427–440. doi: [10.1007/978-3-030-81685-8_20](https://doi.org/10.1007/978-3-030-81685-8_20).
- Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer (2017). “Repairing sequential consistency in C/C++11.” In: *PLDI 2017*. Barcelona, Spain: ACM, pp. 618–632. doi: [10.1145/3062341.3062352](https://doi.org/10.1145/3062341.3062352).
- Leslie Lamport (Sept. 1979). “How to Make a Multiprocessor Computer that Correctly Executes Multiprocess Programs.” In: *IEEE Trans. Computers* 28.9, pp. 690–691. doi: [10.1109/TC.1979.1675439](https://doi.org/10.1109/TC.1979.1675439).
- Sung-Hwan Lee, Minki Cho, Anton Podkopaev, Soham Chakraborty, Chung-Kil Hur, Ori Lahav, and Viktor Vafeiadis (2020). “Promising 2.0: Global optimizations in relaxed memory concurrency.” In: *PLDI 2020*. Ed. by Alastair F. Donaldson and Emina Torlak. ACM, pp. 362–376. doi: [10.1145/3385412.3386010](https://doi.org/10.1145/3385412.3386010).
- Jeremy Manson, William Pugh, and Sarita V. Adve (2005). “The Java memory model.” In: *POPL 2005*. ACM, pp. 378–391. doi: [10.1145/1040305.1040336](https://doi.org/10.1145/1040305.1040336).
- Evgenii Moiseenko, Anton Podkopaev, Ori Lahav, Orestis Melkonian, and Viktor Vafeiadis (2020). “Reconciling Event Structures with Modern Multiprocessors.” In: *ECOOP 2020*. Vol. 166. Leibniz International Proceedings in Informatics

- (LIPICs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 5:1–5:26. doi: [10.4230/LIPICs.ECOOP.2020.5](https://doi.org/10.4230/LIPICs.ECOOP.2020.5).
- Brian Norris and Brian Demsky (2013). “CDSChecker: Checking concurrent data structures written with C/C++ atomics.” In: *OOPSLA 2013*. ACM, pp. 131–150. doi: [10.1145/2509136.2509514](https://doi.org/10.1145/2509136.2509514).
- Jonas Oberhauser et al. (2021). “VSync: Push-Button Verification and Optimization for Synchronization Primitives on Weak Memory Models.” In: *ASPLOS 2021*. Virtual, USA: ACM, pp. 530–545. doi: [10.1145/3445814.3446748](https://doi.org/10.1145/3445814.3446748).
- Peizhao Ou and Brian Demsky (Oct. 2018). “Towards Understanding the Costs of Avoiding Out-of-Thin-Air Results.” In: *Proc. ACM Program. Lang.* 2.OOPSLA. doi: [10.1145/3276506](https://doi.org/10.1145/3276506).
- Scott Owens, Susmit Sarkar, and Peter Sewell (2009). “A better x86 memory model: x86-TSO.” In: *TPHOLs 2009*. Munich, Germany: Springer, pp. 391–407. doi: [10.1007/978-3-642-03359-9_27](https://doi.org/10.1007/978-3-642-03359-9_27).
- Marco Paviotti, Simon Cooksey, Anouk Paradis, Daniel Wright, Scott Owens, and Mark Batty (2020). “Modular relaxed dependencies in weak memory concurrency.” In: *ESOP 2020*. Ed. by Peter Müller. Vol. 12075. LNCS. Springer, pp. 599–625. doi: [10.1007/978-3-030-44914-8_22](https://doi.org/10.1007/978-3-030-44914-8_22).
- Jean Pichon-Pharabod and Peter Sewell (2016). “A concurrency semantics for relaxed atomics that permits optimisation and avoids thin-air executions.” In: *POPL 2016*. St. Petersburg, FL, USA: ACM, pp. 622–633. doi: [10.1145/2837614.2837616](https://doi.org/10.1145/2837614.2837616).
- Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis (Jan. 2019). “Bridging the gap between programming languages and hardware weak memory models.” In: *Proc. ACM Program. Lang.* 3.POPL, 69:1–69:31. doi: [10.1145/3290382](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.” In: *Proc. ACM Program. Lang.* 2.POPL, 19:1–19:29. doi: [10.1145/3158107](https://doi.org/10.1145/3158107).
- Christopher Pulte, Jean Pichon-Pharabod, Jeehoon Kang, Sung-Hwan Lee, and Chung-Kil Hur (2019). “Promising-ARM/RISC-V: A simpler and faster operational concurrency model.” In: *PLDI 2019*. Phoenix, AZ, USA: ACM, pp. 1–15. doi: [10.1145/3314221.3314624](https://doi.org/10.1145/3314221.3314624).
- Tom Ridge (2010). “A rely-guarantee proof system for x86-TSO.” In: *VSTTE 2010*. Vol. 6217. LNCS. Springer, pp. 55–70.
- rmem (2009). *rmem: Executable concurrency models for ARMv8, RISC-V, Power, and x86*. URL: <https://github.com/rem-project/rmem> (visited on Aug. 24, 2019).
- Filip Sieczkowski, Kasper Svendsen, Lars Birkedal, and Jean Pichon-Pharabod (2015). “A separation logic for fictional sequential consistency.” In: *ESOP 2015*. Vol. 9032. LNCS. Berlin, Heidelberg: Springer, pp. 736–761.
- SV-COMP (2019). *Competition on Software Verification (SV-COMP)*. URL: <https://sv-comp.sosy-lab.org/2019/> (visited on Mar. 27, 2019).
- Kasper Svendsen, Jean Pichon-Pharabod, Marko Doko, Ori Lahav, and Viktor Vafeiadis (2018). “A separation logic for a promising semantics.” In: *ESOP 2018*. Ed. by Amal Ahmed. Vol. 10801. LNCS. Springer, pp. 357–384. doi: [10.1007/978-3-319-89884-1_13](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, pp. 691–707. doi: [10.1145/2660193.2660243](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*. Mumbai, India: ACM, pp. 209–220. doi: [10.1145/2676726.2676995](https://doi.org/10.1145/2676726.2676995).

A PROOF OF THE LBRF THEOREM FOR IMM

In this section, we prove that IMM [Podkopaev et al. 2019] provides the $\text{LBRF}_{\text{RC11}}$ guarantee. IMM is an intermediate memory consistency model that abstracts over implementation details of hardware memory models, in particular models of x86, ARMv7, ARMv8, and POWER.

The crucial property of IMM required for our proof is the fact that the model tracks syntactic dependencies between the events and defines a subset of the program order that is preserved (ppo). All typical hardware memory models also satisfy this property, which suggests that they can also be shown to satisfy LBRF with a similar argument.

Before diving deep into the details of the proof, we provide a quick recap of the IMM model. We refer readers to Podkopaev et al. [2019] for a more detailed presentation of the model.

Definition A.1. An IMM execution graph G is an execution graph augmented with *data*, *control*, *address* and *CAS* dependency relations $\langle \text{data}, \text{ctrl}, \text{addr}, \text{casdep} \rangle$. The dependency relation is the union of the above:

$$\text{deps} \triangleq \text{data} \cup \text{ctrl} \cup \text{addr} ; \text{po}^? \cup \text{casdep}$$

Definition A.2. The relation *barrier-ordered-before* is defined as follows:

$$\text{bob} \triangleq \text{po} ; [\text{W}^{\text{rel}}] \cup [\text{R}^{\text{acq}}] ; \text{po} \cup \text{po} ; [\text{F}] \cup [\text{F}] ; \text{po} \cup [\text{W}^{\text{rel}}] ; \text{po} \cap \text{=}_{\text{loc}} ; [\text{W}]$$

Definition A.3. *Preserved program order* is defined as follows:

$$\text{ppo} \triangleq [\text{R}] ; (\text{deps} \cup \text{rfi})^+ ; [\text{W}]$$

Definition A.4. *Detour* relation is defined as follows:

$$\text{detour} \triangleq (\text{coe} ; \text{rfe}) \cap \text{po}$$

Definition A.5. The *global acyclic ordering* is defined as follows:

$$\text{ar} \triangleq \text{rfe} \cup \text{bob} \cup \text{ppo} \cup \text{detour} \cup \text{psc}_F$$

Similarly to Weakestmo and RC11, the IMM model also defines the relations sw , hb , scb , psc_F , psc_{base} , and psc . However, the IMM has a slightly different definition of sw , which also affects other relations depending on it. For the purpose of our proof this discrepancy does not matter. What does matter is that the IMM definitions are stronger than RC11 ones, that is $\text{sw}_{\text{RC11}} \subseteq \text{sw}_{\text{IMM}}$, $\text{hb}_{\text{RC11}} \subseteq \text{hb}_{\text{IMM}}$. In the rest of this section, we omit the subscripts whenever the version of the definition used (i.e., IMM or RC11 one) is clear from the context. The exact definitions of IMM versions of these relations and the motivation behind them can be found in [Podkopaev et al. 2019].

Definition A.6. Execution graph G is IMM-consistent if the following hold:

- ar is acyclic (IMM-NO-THIN-AIR)
- $\text{hb}_{\text{IMM}} ; \text{eco}^?$ is irreflexive. (IMM-COHERENT)
- $\text{rmw} \cap (\text{fr} ; \text{co}) = \emptyset$. (RMW-ATOMIC)
- psc is acyclic. (IMM-SEQUENTIAL-CONSISTENCY)

Definition A.7 (*porf-sequence*). Given an execution graph G , a set of events $\mathcal{E} \subseteq G.E$ is a *porf-sequence* if there exists a linearization $\{e_1, \dots, e_n\} = \mathcal{E}$, s.t. for all $1 \leq i < n$ either events e_i and e_{i+1} are either immediate po neighbors or e_{i+1} reads externally from e_i :

$$\langle e_i, e_{i+1} \rangle \in \text{po}_{\text{imm}} \cup \text{rfe}$$

Definition A.8 (*porf-cycle*). A set of events $\mathcal{E} \subseteq G.E$ is a *porf-cycle* if it is a closed *porf-sequence*, i.e., $\langle e_n, e_1 \rangle \in \text{po}_{\text{imm}} \cup \text{rfe}$.

Definition A.9 (porf-span). Let \mathcal{E} be a **porf**-cycle of graph G . Its *span* is the subsequence $\mathcal{S} = \mathcal{E} \cap (\text{dom}(G.\text{rfe}; [\mathcal{E}]) \cup \text{rng}([\mathcal{E}]; G.\text{rfe}))$. In other words, the span contains writes read externally by some event in the cycle, and read events reading externally from the cycle.

LEMMA A.10. *G be an IMM-consistent graph of program P and let \mathcal{E} be a **porf**-cycle, s.t. the **porf**-prefix of the cycle does not contain other cycles, i.e., for $X = \text{dom}(G.\text{porf}; [\mathcal{E}]) \setminus \mathcal{E}$ the graph $G|_X$ is **porf**-acyclic. Then, there exists a RC11-consistent execution graph G' of P with an LB race.*

PROOF. First, note that the cycle \mathcal{E} cannot consist only of **po** or only of **rfe** edges. In the former case, it would contradict the fact that **po** is partial order and thus it is acyclic. In the latter case, we have that $\text{rfe}; \text{rfe} = \emptyset$, since for **rfe** its domain (write events) and codomain (read events) are disjoint sets.

Second, the cycle \mathcal{E} can contain several smaller subcycles. Below we will show that these subcycles can be eliminated incrementally one-by-one yielding a sequence of IMM consistent graphs each one having smaller number of cycles, and that upon the disposal of the last cycle we will arrive to **porf** acyclic RC11 consistent graph with LB race.

Consider the span \mathcal{S} of \mathcal{E} . We will show that it is always possible to find an edge $\langle r, w \rangle$ and event w' in the span \mathcal{S} , s.t. the following is true:

- $r \in G.R^{\text{rlx}} \setminus G.R_{\text{ex}}$ and $w \in G.W^{\text{rlx}}$
- $\langle r, w \rangle \in G.\text{po} \setminus (G.\text{ppo} \cup G.\text{bob})$
- $\text{dom}([G.R_{G.\text{loc}(r)} \cap \mathcal{S}]; G.\text{po}; [r])$ is empty
- $\langle w', r \rangle \in G.\text{rfe}$
- $\langle w', r \rangle \notin \text{hb}_{\text{IMM}} \cup G.\text{hb}_{\text{IMM}}^{-1}$

We do this by induction on the length n of the span \mathcal{S} .

Base Case: $P(i = 0)$.

In this case we immediately arrive at contradiction. The fact that the span of the cycle is empty implies that the cycle consists only of **po** edges, which cannot be true as we have mentioned above.

Induction Step: $P(i) \Rightarrow P(i + 1)$.

First, note that it cannot be the case that the span consists only of **ppo**, **bob** and **rfe** edges, that is the following is false:

- $\langle e_i, e_{i+1} \rangle \in \text{ppo} \cup \text{bob} \cup \text{rfe}$ for all $1 \leq i < n$;
- $\langle e_n, e_1 \rangle \in \text{ppo} \cup \text{bob} \cup \text{rfe}$.

Otherwise we would get the **ppo** \cup **bob** \cup **rfe** cycle, which is forbidden by IMM consistency.

Therefore, there should exist at least one edge $\langle a, b \rangle \in \text{po} \setminus (\text{ppo} \cup \text{bob})$. Then consider the cases for a and b .

First, assume $a \in G.R$ and $b \in G.W$. Then note that it should be the case that $a \in G.R^{\text{rlx}} \setminus G.R_{\text{ex}}$ since otherwise we would have $\langle a, b \rangle \in [G.R^{\text{acq}}]; \text{po} \subseteq \text{bob}$. Similarly if $a \in G.R_{\text{ex}}$ then $\langle a, b \rangle \in [G.R_{\text{ex}}]; \text{po} \subseteq \text{ppo}$. By the same argument $w \in G.W^{\text{rlx}}$, because otherwise $\langle a, b \rangle \in \text{po}; [W^{\text{re1}}] \subseteq \text{bob}$.

Second, suppose there exists $c \in G.R_{G.\text{loc}(a)} \cap \mathcal{S}$, s.t. $\langle c, a \rangle \in G.\text{po}$. Then we can construct a span of smaller length with an edge $\langle c, b \rangle$ which, by inductive assumption, should contain edge $\langle r, w \rangle$ satisfying our constraints. Therefore, there exists no such c .

Then let w_a be a write event of the span \mathcal{S} s.t. $\langle w_a, a \rangle \in G.\text{rfe}$. Suppose $\langle w_a, a \rangle \in \text{hb}_{\text{IMM}} \cup \overline{G.\text{hb}_{\text{IMM}}^{-1}}$. If $\langle w_a, a \rangle \in \text{hb}_{\text{IMM}}$ then consider the last edge in hb_{IMM} path. It cannot be sw_{IMM} edge since it would imply $a \in G.R^{\text{acq}}$ which cannot be true as we have shown above. If the last edge is **po** edge, then there should exist $c \in \text{rng}([w_a]; G.\text{rfe})$ and $\langle c, a \rangle \in G.\text{po}$. But we have proven there exists no such c . Next assume $\langle a, w_a \rangle \in \text{hb}_{\text{IMM}}$. Then there should exist $d \in W$, s.t. $\langle a, d \rangle \in \text{po}$, and either d is a release write or there exists release fence between a and d . Then there should exist

write d' belonging to a span \mathcal{S} , s.t. $\langle d, d' \rangle \in (\text{po} \cap =_{\text{loc}})^?$. Then again we can construct a smaller span with an edge $\langle a, d' \rangle$ and by induction get suitable edge $\langle r, w \rangle$. If neither $\langle w_a, a \rangle \in \text{hb}_{\text{IMM}}$ nor $\langle a, w_a \rangle \in \text{hb}_{\text{IMM}}$, we have $r = a$, $w = b$ and $w' \in \text{dom}(G.\text{rfe}; [a])$ which satisfy our constraints.

Otherwise, let, for example $a \in G.R$ and $b \in G.R$. Let w_a is write from which a reads, and w_b is a write from which b reads. By definition of the span, both writes are external writes belonging to the porf cycle. But then for the cycle \mathcal{S} to be closed, there should exist $w', r' \in \mathcal{S}$, s.t. $\langle w', r' \rangle \in G.\text{rfe}$, $\langle b, w' \rangle \in \text{po}$, and $\langle r', w_b \rangle \in G.\text{porf}$. Then we would be able to construct a smaller porf span from w' going through r' to b without visiting a . By induction this span should contain the required edge $\langle r, w \rangle$. Other cases for a and b can be handled similarly.

As a result, we get an edge $\langle r, w \rangle$ and write w' satisfying the constraints given above. Note that by definition of bob we can derive that $\langle r, w \rangle \in \text{rpo}$.

Next, using the *receptiveness* property one can construct an IMM consistent graph \bar{G} , s.t. $G \cong \bar{G}$ in the following sense

- $G.\text{lab}|_X = \bar{G}.\text{lab}|_X$ where $X \triangleq G.E \setminus \text{rng}([r]; G.\text{ppo}^?; G.\text{rfi}^?)$
- $\bar{G}.E = G.E$
- $\bar{G}.\text{po} = G.\text{po}$
- $\bar{G}.\text{rf}|_{G.E \setminus \{r\}} = G.\text{rf}|_{G.E \setminus \{r\}}$
- $\bar{G}.\text{rfe}; [r] \subseteq G.\text{rfe}^?; G.\text{hb}_{\text{IMM}}$
- $\bar{G}.\text{co} = G.\text{co}$

Put simply, it is possible to construct a graph similar to the original one, that would only be different in the choice of reads-from write for the read r , and the labels of events that depend on r via preserved program order. The formal definition of receptiveness property and the details on graph construction can be found in [Podkopaev et al. 2019, §6.4]. Below we give some intuitive explanation. Since the ppo relation includes the dependencies relations and reads-from relation, which together capture the control and data flow in intra-thread computation, the value of the read (and its reads-from source) can be safely replaced with a different value, and this change will affect only the ppo descendants of the read. Since we have that $\langle r, w \rangle \notin G.\text{ppo}$, the label of w would not be affected if we'll redirect the read r to read-from $G.\text{rfe}^?; G.\text{hb}_{\text{IMM}}$ prior write. Redirecting to such write is safe, since it cannot violate axiom IMM-COHERENT . Also, because $r \notin G.R_{\text{ex}}$, a redirection of this read cannot violate RMW-ATOMIC .

Therefore we have an IMM consistent graph \bar{G} , which has one porf cycle less.

It left to note that upon the elimination of the last cycle, we would have edge $\langle r, w \rangle$ and a write event w' , s.t.:

- $\text{loc}(r) = \text{loc}(w')$
- $\langle r, w \rangle \in \bar{G}.\text{rpo}$
- $\langle w, w' \rangle \in \bar{G}.\text{rfe}; \bar{G}.\text{porf};$
- $\langle w', r \rangle \in \bar{G}.\text{hb}_{\text{IMM}} \cup \bar{G}.\text{hb}_{\text{IMM}}^{-1}$

That is, the pair $\langle r, w' \rangle$ forms an LB race. Since the sequence \mathcal{E} does not form a cycle in \bar{G} , and its porf prefix $\mathcal{P} \triangleq \text{dom}(\bar{G}.\text{porf}; [\mathcal{E}])$ also contains no cycles by our assumption, we have that $\bar{G}|_{\mathcal{P}}$ is porf acyclic. Since it is also IMM consistent, we have that it also satisfies axioms COHERENT , RMW-ATOMIC , $\text{IMM-SEQUENTIAL-CONSISTENCY}$ (in particular, it is implied by the fact that IMM has a stronger notion of happens-before $\text{hb}_{\text{RC11}} \subseteq \text{hb}_{\text{IMM}}$). Therefore \bar{G} is a RC11-consistent graph of P and $\langle r, w' \rangle$ is an LB race. \square

THEOREM A.11. *Let P be LB-race-free under RC11. Then the set of its RC11 consistent execution graphs coincides with the set of its IMM consistent graphs.*

PROOF. The most challenging part of the proof has already been done in A.10.

Let G be some IMM consistent graph of P and assume it is not RC11 consistent, i.e., it contains some $G.\text{porf}$ cycles. Then we can order non-overlapping $G.\text{porf}$ cycles: the cycle \mathcal{E}_1 is said to be less than \mathcal{E}_2 if $\mathcal{E}_1 \subseteq \text{dom}(G.\text{porf}; [\mathcal{E}_2])$. It is easy to show that this definition gives us partial order. Next pick an arbitrary minimal cycle \mathcal{E} . Because its minimal w.r.t. to the partial order defined above, its $G.\text{porf}$ prefix cannot contain another $G.\text{porf}$ cycle. Therefore the preconditions of lemma A.10 are met and we can construct a RC11 consistent execution graph with LB race, which contradicts our assumption that P is LB-race-free. \square

B DEFINITION OF Promising WITH CERTIFICATION LOCALITY

In this section, we define a version of the Promising semantics that satisfies certification locality. For conciseness, we present only the *relaxed* fragment of the Promising machine, since its extra machinery to handle release/acquire synchronization, RMW operations, and fences is orthogonal to our adaptations.

Promising represents memory, M , as a set of messages, where each message $m = \langle x : v@t \rangle$ is a tuple containing a location x , a value v , and a timestamp $t \in \mathbb{Q}$ and represents a write event. Besides the memory, the state of the operational semantics also contains \mathcal{TS} , which is a map from thread identifiers to the thread-local state of each thread. This thread-local state, TS , contains three components: (1) the actual thread-local state, σ , containing the values of the program counter and the local registers, which is modeled abstractly; (2) the thread view, V , a map from locations to timestamps, recording the latest write for each location that the thread is aware of, i.e., has read or executed itself; and (3) a set, P , of messages that the thread has promised.

In our adaptation, we change this last component of the thread state, P , to be a partial function from messages to sets of messages. The purpose of this change is to represent the set of external writes that are performed in the certification of the promise so that the actual execution performs the exact same set of external writes before fulfilling the promise.

The top-level rule of Promising remains unchanged:

$$\frac{TS = \mathcal{TS}(i) \quad TS''.\text{prm} = \emptyset \quad \langle TS, M \rangle \rightarrow \langle TS', M' \rangle \quad \langle TS', M' \rangle \xrightarrow{\text{promise-free}^*} \langle TS'', M'' \rangle}{\langle \mathcal{TS}, M \rangle \rightarrow \langle \mathcal{TS}[i \mapsto TS'], M' \rangle} \text{ (Machine Step)}$$

The machine can make a step, if some thread i makes a step to some configuration $\langle TS', M' \rangle$ for which all outstanding promises are certifiable. That is, there exists a promise-free execution of that thread that reaches a configuration $\langle TS'', M'' \rangle$ with no outstanding promises.

We move on to the thread-local transitions. The first rule concerns ‘silent’ transitions that do not interact with memory. Here and in subsequent rules, Promising assumes a labeled transition relation saying how to update the values of the registers, $\sigma \xrightarrow{\ell} \sigma'$.

$$\frac{\sigma \xrightarrow{\epsilon} \sigma'}{\langle \langle \sigma, V, P \rangle, M \rangle \rightarrow \langle \langle \sigma', V, P \rangle, M \rangle} \text{ (Silent)}$$

The next two rules concern reads. While the original Promising semantics has a single rule for handling this case, we split it into two rules to distinguish between local and external reads.

$$\frac{\sigma \xrightarrow{R(x,v)} \sigma' \quad V(x) = t \quad \langle x : v@t \rangle \in M}{\langle \langle \sigma, V, P \rangle, M \rangle \rightarrow \langle \langle \sigma', V, P \rangle, M \rangle} \text{ (Read-Local)}$$

The rule for external reads checks that all outstanding promises of the thread allow reading the not previously observed message m , and removes m from the sets of external messages recorded in the P component.

$$\frac{\sigma \xrightarrow{R(x,v)} \sigma' \quad V(x) < t \quad \forall p \in \text{dom}(P). m \in P(p) \quad m = \langle x : v@t \rangle \in M \quad P' = \{\langle p, \mathcal{W} \setminus \{m\} \rangle \mid \langle p, \mathcal{W} \rangle \in P\}}{\langle \langle \sigma, V, P \rangle, M \rangle \rightarrow \langle \langle \sigma', V[x \mapsto t], P' \rangle, M \rangle} \text{ (Read-External)}$$

Next is the rule for promising a message. Promising is essentially unconditional: all that is required is for the message to not already exist in memory. In practice, of course, promises are constrained by the top-level rule, which requires all promises to be certifiable, and so only promises guided by the subsequent instructions of the thread have a chance of being certified.

$$\frac{M' = M \uplus \{m\} \quad P' = P \uplus \{\langle m, \mathcal{W} \rangle\}}{\langle \langle \sigma, V, P \rangle, M \rangle \rightarrow \langle \langle \sigma, V, P' \rangle, M' \rangle} \text{ (Promise)}$$

The next rule is for a write that fulfills an outstanding promise. This rule is also used for normal (non-promised) writes: we simply precede it with the step that promises the necessary write.

$$\frac{\sigma \xrightarrow{W(x,v)} \sigma' \quad V(x) < t \quad \langle m, \emptyset \rangle \in P \quad m = \langle x : v@t \rangle \quad P' = P \setminus \{\langle m, \emptyset \rangle\}}{\langle \langle \sigma, V, P \rangle, M \rangle \rightarrow \langle \langle \sigma', V[x \mapsto t], P' \rangle, M \rangle} \text{ (Fulfill)}$$

Note that this rule requires the fulfilled promise to have associated an empty set of external messages, that is, all associated external messages need to be already have read.

The final rule allows us to extend the set of external reads attached to outstanding promises by a single additional message.

$$\frac{m_1, \dots, m_k \in \text{dom}(P) \quad P' = P[m_1 \mapsto P(m_1) \cup \{m'\}] \dots [m_k \mapsto P(m_k) \cup \{m'\}]}{\langle \langle \sigma, V, P \rangle, M \rangle \rightarrow \langle \langle \sigma, V, P' \rangle, M \rangle} \text{ (Extend-Promise)}$$

Promise-free thread transitions are all the above, except for the (Extend-Promise) step. It can be shown that if there is a sequence of thread-local steps that reaches a configuration without any promises, that sequence can be rearranged so that every (Promise) step is immediately before the corresponding (Fulfill) step.

C COMPLETE DEFINITION OF WEAKESTMO2

Definition C.1. An *event* is a tuple $\langle \text{id}, \text{tid}, \text{lab} \rangle$ where $\text{id} \in \mathbb{N}$ is a unique identifier for the event, $\text{tid} \in \mathbb{N}$ identifies the thread to which the event belongs, and lab is a *label* of the form:

- $R^o(x, v)$ for a read of $v \in \text{Val}$ from $x \in \text{Loc}$ with mode $o \in \{\text{ex}, \text{non-ex}\} \times \text{Mod}$;
- $W^o(x, v)$ for a write of $v \in \text{Val}$ to $x \in \text{Loc}$ with $o \in \{\text{ex}, \text{non-ex}\} \times \text{Mod}$;
- F^o for a fence with mode $o \in \text{Mod}$,

where $\text{Mod} = \{\text{rlx}, \text{acq}, \text{rel}, \text{acqrel}, \text{sc}\}$. The modes are partially ordered by \sqsubset as follows:

$$\begin{array}{ccccc} & & \text{rel} & & \\ & \swarrow & & \searrow & \\ \text{rlx} & & & & \text{acqrel} \sqsubset \text{sc} \\ & \swarrow & & \searrow & \\ & & \text{acq} & & \end{array}$$

Definition C.2. An *event structure* is S is a tuple $\langle E, \text{po}, \text{jf}, \text{ew}, \text{co} \rangle$ where:

- $E \subseteq \text{Event}$ is a set of events.
- $\text{po} \subseteq E \times E$ is *program order* relation.
- $\text{jf} \subseteq E \times E$ is *justified from* relation.
- $\text{ew} \subseteq E \times E$ is *equal-writes* relation.
- $\text{co} \subseteq E \times E$ is *coherence order* relation.

Definition C.3. *Conflict* relation is defined as follows: $\text{cf} \triangleq =_{\text{tid}} \setminus (\text{po} \cup \text{po}^{-1})^?$

Definition C.4. Immediate program order and immediate conflict relations are defined as follows:

$$\begin{aligned} \text{po}|_{\text{imm}} &\triangleq \text{po} \setminus (\text{po} ; \text{po}) \\ \text{cf}|_{\text{imm}} &\triangleq \text{cf} \setminus (\text{cf} ; \text{po} \cup \text{po}^{-1} ; \text{cf}) \end{aligned}$$

Definition C.5. Weakestmo event structure is well-formed if the following conditions are met:

- (1) There exists the initialization event $e_0 \in E$.

We assume that: $e_0 = \langle 0, \perp, W^{r1x}(\perp, 0) \rangle$ For $f \in \{\text{tid}, \text{lab}, \text{loc}\}$ we define:

$$\begin{aligned} <_f &\triangleq \{e_0\} \times (E \setminus \{e_0\}) \\ \leq_f &\triangleq \{e_0\} \times (E \setminus \{e_0\}) \cup =_f \end{aligned}$$

- (2) po is strict partial order.

- (3) po orders events from the same thread (the initialization event is placed before other events):

$$<_{\text{tid}} \subseteq \text{po} \subseteq \leq_{\text{tid}}$$

- (4) po is downward-total: $\text{po} ; \text{po}^{-1} \subseteq (\text{po} \cup \text{po}^{-1})^?$

- (5) Immediate conflict can only appear between relaxed reads: $\text{cf}|_{\text{imm}} \subseteq R_{\text{ex}} \times R_{\text{ex}}$

- (6) Every write-exclusive has a corresponding read part:

$$W_{\text{ex}} \subseteq \text{rng}([R_{\text{ex}}] ; (\text{po}|_{\text{imm}} \cap =_{\text{loc}}))$$

- (7) jf relates write/read pairs of accesses to same location with the same value.

$$\text{jf} \subseteq [W] ; (\leq_{\text{loc}} \cap =_{\text{val}}) ; [R]$$

- (8) jf^{-1} is functional: $\text{jf} ; \text{jf}^{-1} \subseteq [W]$

- (9) Every read event has a justification event: $R \subseteq \text{rng}(\text{jf})$

- (10) ew is a symmetric relation.

- (11) ew relates conflicting write events with the same location and value.

$$\text{ew} \subseteq [W^{r1x}] ; (\text{cf} \cap =_{\text{loc}} \cap =_{\text{val}}) ; [W^{r1x}]$$

- (12) co is strict partial order.

- (13) co orders write events with the same location

(the initialization event is placed before other events).

$$[W] ; <_{\text{loc}} ; [W] \subseteq \text{co} \subseteq [W] ; \leq_{\text{loc}} ; [W]$$

- (14) co does not link the writes from the same ew^* class:

$$\text{co} \cap \text{ew}^* \subseteq \emptyset$$

- (15) co is closed w.r.t. ew^* equivalence classes:

$$\text{ew}^* ; \text{co} ; \text{ew}^* \subseteq \text{co}$$

- (16) co is total on writes to the same location w.r.t. ew^* as equivalence relation:

$$\forall x \in \text{Loc}. W_x \times W_x \subseteq \text{ew}^* \cup \text{co} \cup \text{co}^{-1}$$

Given an event structure, we define the following auxiliary relations:

$$\text{rf} \triangleq (\text{ew}^* ; \text{jf}) \setminus \text{cf} \quad (\text{reads-from})$$

$$\text{jfi} \triangleq \text{jf} \cap \text{po} \text{ and } \text{rfi} \triangleq \text{rf} \cap \text{po} \quad (\text{internal jf/rf})$$

$$\text{jfe} \triangleq \text{jf} \setminus \text{po} \text{ and } \text{rfe} \triangleq \text{rf} \setminus \text{po} \quad (\text{external jf/rf})$$

$$\text{jo} \triangleq \text{jfe} ; (\text{po} \cup \text{jf})^* \quad (\text{justification order})$$

$$\text{lbp} \triangleq \text{cf}|_{\text{imm}} ; [\text{rng}(\text{jf} \cap (\text{jf}^? ; \text{hb}))] ; \text{po} \quad (\text{load buffering pattern})$$

$rmw \triangleq [R_{ex}] ; po _{imm} ; [W_{ex}]$	(read-modify-write pairs)
$sw \triangleq [E^{\exists rel}] ; ([F] ; po)^? ; (jf ; rmw)^* ; jf ; (po ; [F])^? ; [E^{\exists acq}]$	(synchronizes-with)
$hb \triangleq (po \cup sw)^+$	(happens before)
$ecf \triangleq (hb^{-1})^? ; cf ; hb^?$	(extended conflict)
$fr \triangleq S.rf^{-1} ; S.co$	(from-reads)
$eco \triangleq (co \cup rf \cup fr)^+$	(extended coherence order)
$scb \triangleq po \cup po _{\neq loc} ; hb ; po _{\neq loc} \cup hb _{=loc} \cup co \cup fr$	(SC-before)
$p_{sc}_{base} \triangleq ([E^{sc}] \cup [F^{sc}] ; hb^?) ; scb ; ([E^{sc}] \cup hb^? ; [F^{sc}])$	
$p_{sc}_F \triangleq [F^{sc}] ; (hb \cup hb ; eco ; hb) ; [F^{sc}]$	
$p_{sc} \triangleq p_{sc}_{base} \cup p_{sc}_F$	(partial SC order)

For event $e \in S.E$, we call the set $dom(jo ; [e])$ the *justification set* of e .

Definition C.6. A *justified configuration* C of an event structure S is a conflict-free, rf -complete, po -prefix-closed subset of its events. That is,

$$C \subseteq S.E \text{ and } [C] ; cf ; [C] = \emptyset \text{ and } dom(po ; [C]) \subseteq C \text{ and } C \cap R \subseteq rng([C] ; rf)$$

Definition C.7. An execution graph G is *extracted from* an event structure S , denoted as $S \triangleright G$, if there exists a justified configuration C s.t. $G = S|_C$ (in particular, $G.rf = S.rf|_C$).

Definition C.8. Event structure S is *Weakestmo2-consistent* if the following conditions are met.

- ecf is irreflexive. (NON-CONTRADICTIONARY)
- $jf \cap ecf = \emptyset$ (WELL-JUSTIFIED)
- $po \cup jf$ is acyclic (NO-THIN-AIR)
- $hb ; eco^?$ is irreflexive. (COHERENT)
- $(jfe ; [E^{\exists acq}] \cup [F]) ; po ; ew^+ \subseteq (jfe ; [E^{\exists acq}] \cup [F]) ; po$ (WELL-FENCED)
- $cf \cap jo \subseteq ew^+ ; (po \cup po^{-1})^?$ (CERTIFIED)
- $ew \subseteq (cf \cap (jo \cup jo^{-1}))^+$ (GROUNDED)
- $(jf \setminus (jf^? ; hb)) ; po ; ew \subseteq jf ; (po \cup lbpat)$ (NO-BAIT-AND-SWITCH)

Definition C.9. Execution graph G is *Weakestmo2-consistent* if there exists a Weakestmo2-consistent event structure s.t. $S \triangleright G$ and additionally G satisfies the following constraints:

- $G.rmw \cap (G.fr ; G.co) = \emptyset$. (RMW-ATOMIC)
- $G.p_{sc}$ is acyclic. (SEQUENTIAL-CONSISTENCY)

D PROOF OF THE LBRF THEOREM FOR A CORE FRAGMENT OF WEAKESTM02

In this section, we prove that the fragment of Weakestmo2 without read-modify-writes and SC features (i.e., with only relaxed and release/acquire reads, writes and fences) provides the LBRF_{RC11} guarantee. This corresponds to the version of the model presented in §3 extended with release/acquire accesses.

We postpone the proof of the theorem for the full model to Appendix E because the other features of the model are mostly orthogonal to the key ideas of the proof and yet they bring several technical difficulties, which require a refinement of LB race definition or a strengthening of the model.

Definition D.1. Execution graph G is RC11-consistent if the following hold:

- $po \cup rf$ is acyclic (RF-NO-THIN-AIR)
- $hb ; eco^?$ is irreflexive. (COHERENT)

LEMMA D.2. *Every RC11-consistent execution graph is Weakestmo2-consistent.*

PROOF. Let G be a RC11-consistent graph. Let $S = G$ (i.e., with $S.jf = G.rf$ and $G.cf = \emptyset$). Clearly, $S \triangleright G$. Therefore, we only need to show that S is a Weakestmo2-consistent event structure.

Since $S.cf = \emptyset$, axioms **NON-CONTRADICTIONARY**, **WELL-JUSTIFIED** are trivially satisfied. By well-formedness, we also have that $S.ew = \emptyset$ and therefore **WELL-FENCED**, **CERTIFIED**, **GROUNDING**, and **NO-BAIT-AND-SWITCH** are satisfied as well.

Since G satisfies **COHERENT**, then S also satisfies this constraint.

Finally, $(S.po \cup S.jf) = (G.po \cup G.rf)$ and thus its acyclicity follows directly from **RF-NO-THIN-AIR**. \square

Definition D.3. A po-edge between a read and a write in an event structure is *reorderable*, if it connects two relaxed accesses and there is no release or acquire fence in between.

$$rpo \triangleq [R^{rx}] ; (po \setminus (po ; [F] ; po)) ; [W^{rx}]$$

Definition D.4. Given a well-formed execution graph G , a pair of events $\langle e_1, e_2 \rangle$ is called a *load buffering race* if it satisfies the following:

- $G.loc(e_1) = G.loc(e_2)$
- $\langle e_1, e_2 \rangle \notin G.hb \cup G.hb^{-1}$
- $\langle e_1, e_2 \rangle \in [R^{rx}] ; rpo ; rfe ; (po \cup rf)^* ; [W]$

Definition D.5. A program P is said to be LB-race-free under RC11 if no RC11 consistent execution graph of P contains a load buffering race.

Definition D.6. Event structure S is *prefix-closed* if $\forall e \in S.E$ we have $[e]_{(po \cup jf)^*} \subseteq S.E$ where $[e]_r \triangleq dom(r ; [e])$

Definition D.7. Given event structure S and subset of its events $E \subseteq S.E$ we call $S|_E \triangleq \langle E, S.po|_E, S.jf|_E, S.ew|_E, S.co|_E \rangle$ the *restriction* of S onto E .

LEMMA D.8. *For any well-formed S and any $E \subseteq S.E$ if $S|_E$ is prefix-closed then $S|_E$ is well-formed as well.*

PROOF. Follows directly by observing that all well-formedness properties are preserved under the restriction to prefix-closed subset of events. \square

Definition D.9. Event structure S is *promise-free-consistent* if it satisfies all the consistency constraints (see Def. C.8) except **WELL-FENCED**, **CERTIFIED**, **GROUNDING**, and **NO-BAIT-AND-SWITCH**.

LEMMA D.10. *For any promise-free-consistent S and any $E \subseteq S.E$ if $S|_E$ is prefix-closed then $S|_E$ is promise-free-consistent as well.*

PROOF. Follows directly from the fact that properties **NON-CONTRADICTIONARY**, **WELL-JUSTIFIED**, **NO-THIN-AIR**, and **COHERENT** are monotone, in the sense that if they hold for S then they hold for any restriction of S \square

Definition D.11. Event structure S is *promise-free* if it satisfies the following constraint.

- $cf \cap jo = \emptyset$ (PROMISE-FREE)

LEMMA D.12. *If S is PROMISE-FREE then $cf \cap (po \cup jf)^+ = \emptyset$*

PROOF. First note that for two arbitrary relations r_1, r_2 the following equation holds:

$$(r_1 \cup r_2)^+ = r_1^+ \cup r_1^* ; (r_2 ; r_1^*)^*$$

By substitution $r_1 \mapsto \text{po}$ and $r_2 \mapsto \text{jf} \setminus \text{po}$ we get (note that since po is partial order we have $\text{po}^+ = \text{po}$ and $\text{po}^* = \text{po}^?$):

$$\begin{aligned} \text{cf} \cap (\text{po} \cup \text{jf})^+ &= \text{cf} \cap (\text{po} \cup \text{jf} \setminus \text{po})^+ = \\ &= \text{cf} \cap (\text{po} \cup \text{po}^? ; ((\text{jf} \setminus \text{po}) ; \text{po}^?))^* \end{aligned}$$

$\text{cf} \cap \text{po} = \emptyset$ by the definition of cf . Observing that $(\text{jf} \setminus \text{po}) ; \text{po}^? \subseteq \text{jo}$ and that jo is transitive we left to show that $\text{cf} \cap (\text{po}^? ; \text{jo}) = \emptyset$. Suppose we have three events e_1, e_2 , and e_3 s.t. $\langle e_1, e_3 \rangle \in \text{cf}$, $\langle e_1, e_2 \rangle \in \text{po}^?$ and $\langle e_2, e_3 \rangle \in \text{jo}$. But since conflict propagates along program order, i.e., $\text{cf} ; \text{po} \subseteq \text{cf}$, we also have $\langle e_2, e_3 \rangle \in \text{cf}$, which contradicts our assumption that $\text{cf} \cap \text{jo} = \emptyset$. \square

LEMMA D.13. *If S is consistent and PROMISE-FREE then $S.\text{ew} = \emptyset$*

PROOF. By GROUNDED we have $\text{ew} \subseteq (\text{cf} \cap (\text{jo} \cup \text{jo}^{-1}))^+$. By PROMISE-FREE we also have $\text{cf} \cap (\text{jo} \cup \text{jo}^{-1}) = \emptyset$ which implies $\text{ew} \subseteq \emptyset$. \square

LEMMA D.14. *If S is consistent and PROMISE-FREE then $S.\text{jf} = S.\text{rf}$*

PROOF. Follows trivially from the definition of rf , lemma D.13 and WELL-JUSTIFIED axiom.

$$\text{rf} \triangleq (\text{ew}^* ; \text{jf}) \setminus \text{cf} = \text{jf} \setminus \text{cf} = \text{jf}$$

\square

LEMMA D.15. *If S is consistent and PROMISE-FREE then every extracted execution $S \triangleright G$ is RC11-consistent.*

PROOF. Let us prove that G satisfies all the RC11 consistency constraints:

- **RF-NO-THIN-AIR**

By D.14 we have $S.\text{po} \cup S.\text{rf} = S.\text{po} \cup S.\text{jf}$. From axiom NO-THIN-AIR we know that $S.\text{po} \cup S.\text{jf}$ is acyclic and therefore $G.\text{po} \cup G.\text{rf} \subseteq S.\text{po} \cup S.\text{rf} = S.\text{po} \cup S.\text{jf}$ is acyclic as well.

- **COHERENT**

Follows from the coherence of S and the fact that $G.\text{hb} ; G.\text{eco}^? \subseteq S.\text{hb} ; S.\text{eco}^?$.

\square

LEMMA D.16. *Given consistent event structure S , if additionally it is PROMISE-FREE then for every event $e \in S.E$ its $(\text{po} \cup \text{jf})$ prefix is extractable execution graph, i.e., for $G \triangleq S|_E$, where $E = \lfloor e \rfloor_{(\text{po} \cup \text{jf})^+}$, we have $S \triangleright G$.*

PROOF. By definition of justified configuration C.6 and extracted execution C.7 we need to show that E forms a justified configuration. Since E is $(\text{po} \cup \text{jf})^*$ downward closure of the event e , then clearly it is po downward-closed. rf completeness follows from the fact that $\text{jf} = \text{rf}$ by lemma D.14, jf is complete by well-formedness of S , and E is $(\text{po} \cup \text{jf})^*$ downward closure of the event e . The fact that E is conflict free follows from the lemma D.12. \square

LEMMA D.17. *If S is consistent and PROMISE-FREE then for every event e its $(\text{po} \cup \text{jf})$ prefix forms RC11-consistent execution graph.*

PROOF. Direct consequence of lemmas D.15 and D.16 \square

LEMMA D.18. *For a consistent event structure S , the following is true*

$$\text{cf} \cap \text{jo} \subseteq \text{rpo}^{-1} ; \text{cf}|_{\text{imm}} ; \text{po}^?$$

PROOF. Let $\langle w, e \rangle \in \mathbf{cf} \cap \mathbf{jo}$. Then there should exist $\langle r, r' \rangle \in \mathbf{cf}|_{\text{imm}}$, s.t. $\langle r, w \rangle \in \text{po}^?$ and $\langle r', e \rangle \in \text{po}^?$. By well-formedness $r, r' \in S.R^{\text{rlx}}$, and $w \in S.W$. Therefore $\langle r, w \rangle \in \text{po}$. It is left to show that $\langle r, w \rangle \in \text{rpo}$. By **CERTIFIED** there should exist w' , s.t. $\langle w, w' \rangle \in \text{ew}$ and $\langle w', e \rangle \in (\text{po} \cup \text{po}^{-1})^?$. Then we get that $w, w' \in W^{\text{rlx}}$ by well-formedness of S . Next, suppose there is an acquire/release fence f between r and w . Then it would be the case that either $\langle f, w \rangle \in [F]; \text{po}$. But since $\langle r, f \rangle \in \text{po}$ we have $\langle f, w' \rangle \in \mathbf{cf}$. Therefore $\langle f, w' \rangle \notin \text{po}$ which contradicts **WELL-FENCED**. \square

LEMMA D.19. *In the consistent event structure S promises can only be relaxed:*

$$\text{dom}(\mathbf{cf} \cap \mathbf{jo}) \subseteq W^{\text{rlx}}$$

PROOF. Suppose there exists $w \in \text{dom}(\mathbf{cf} \cap \mathbf{jo})$ s.t. $w \in W^{\text{rel}}$. Then by axiom **CERTIFIED** we have that there exists w' , s.t. $\langle w, w' \rangle \in \text{ew}$. But then from well-formedness we can deduce that $w, w' \in W^{\text{rlx}}$ which contradicts assumption that w is at least release write. \square

LEMMA D.20. *Let S be a consistent event structure and let $e \in S.E$ be one of its events. Suppose e depends on a promise p , i.e., $\langle p, e \rangle \in \mathbf{cf} \cap \mathbf{jo}$. Also, assume that a branch to which e belongs has a certification write c for p , i.e., $\langle e, c \rangle \in \text{po}$ and $\langle p, c \rangle \in \text{ew}$. Finally, let r and r' be read events at which two branches diverge, i.e., $\langle r, p \rangle \in \text{po}$, $\langle r', e \rangle \in \text{po}$, and $\langle r, r' \rangle \in \mathbf{cf}|_{\text{imm}}$. Then it has to be that $\langle p, r' \rangle \in \mathbf{cf} \cap \mathbf{jo}$.*

In other words, each promise has to be observed by the neighboring certification branch up to the point of immediate conflict between this certification branch and the branch that issued the promise.

PROOF. Note that $\langle r, p \rangle \in \text{po}$ and $\langle r, r' \rangle \in \mathbf{cf}|_{\text{imm}}$. Therefore, since conflict relation propagates along program order, we have that $\langle p, r' \rangle \in \mathbf{cf}$. It is left to show that $\langle p, r' \rangle \in \mathbf{jo}$.

Suppose that $\langle p, r' \rangle \notin \mathbf{jo}$. Let r'' be the first event in a branch between r' and e that observes p :

- $\langle r', r'' \rangle \in \text{po}$
- $\langle r'', e \rangle \in \text{po}$
- $\langle p, r'' \rangle \in \mathbf{jo}$
- $\langle p, r'' \rangle \notin \mathbf{jo}; \text{po}$

Also let w'' be the write from which r'' reads, i.e., $\langle p, w'' \rangle \in \mathbf{jo}$ and $\langle w'', r'' \rangle \in \mathbf{jf}$.

Suppose $\langle w'', r'' \rangle \in \mathbf{jf}^?; \mathbf{hb}$. Note that $\mathbf{jf}^?; \mathbf{hb} \subseteq (\text{po} \cup \mathbf{jf})^+$. If the last edge in the $\mathbf{jf}^?; \mathbf{hb}$ path from w'' to r'' is a po edge then we would have $\langle p, r'' \rangle \in \mathbf{jo}; (\text{po} \cup \mathbf{jf})^+; \text{po} \subseteq \mathbf{jo}; \text{po}$. This contradicts our assumption. Therefore the last edge in the $\mathbf{jf}^?; \mathbf{hb}$ path has to be $\mathbf{jfe}; [E^{\text{acq}}]$. But then, by **WELL-FENCED** it should be the case that $\langle w'', p \rangle \in \mathbf{jfe}; [E^{\text{acq}}]; \text{po} \subseteq (\text{po} \cup \mathbf{jf})^*$ and since also $\langle p, w'' \rangle \in \mathbf{jo} \subseteq (\text{po} \cup \mathbf{jf})^*$ we have a $(\text{po} \cup \mathbf{jf})$ cycle, which contradicts axiom **NO-THIN-AIR**. Therefore $\langle w'', r'' \rangle \notin \mathbf{jf}^?; \mathbf{hb}$. In other words r'' is justified externally. Thus we have $\langle w'', p \rangle \in (\mathbf{jf} \setminus \mathbf{jf}^?; \mathbf{hb}); \text{po}; \text{ew}$

Also note that $\langle r'', p \rangle \notin \text{lbpat}$, since $\langle r'', p \rangle \notin \mathbf{cf}|_{\text{imm}}; \text{po}$ (the later is true since we know $\langle r', p \rangle \in \mathbf{cf}|_{\text{imm}}; \text{po}$ and $\langle r', r'' \rangle \in \text{po}$).

Thereby the preconditions of axiom **NO-BAIT-AND-SWITCH** apply to $\langle w'', p \rangle$, and we can conclude that $\langle w'', p \rangle \in \mathbf{jf}; \text{po} \subseteq (\text{po} \cup \mathbf{jf})^*$. But then again, using the fact that $\langle p, w'' \rangle \in \mathbf{jo} \subseteq (\text{po} \cup \mathbf{jf})^*$, we arrive at contradiction, because we have a $(\text{po} \cup \mathbf{jf})^*$ cycle forbidden by axiom **NO-THIN-AIR**. \square

LEMMA D.21. *Let P be LB-race-free under RC11. Then every consistent event structure S of P is **PROMISE-FREE**.*

PROOF. Let R be a total extension of $(\text{po} \cup \mathbf{jf})^+$ partial order. Let $\{e_1, \dots, e_n\}$ be sequence of events of S , ordered w.r.t. R :

- $\forall i \in \{1, \dots, n\}. e_i \in S.E$
- $\forall e \in S.E. \exists i. e = e_i$
- $\forall i, j \in \{1, \dots, n\}. i < j \Rightarrow \langle e_i, e_j \rangle \in R$

Let $S_i \triangleq S|_{\{e_0, e_1, \dots, e_i\}}$. We show by induction on i that conditions stated in the lemma holds for $S_0, S_1, \dots, S_n = S$. Note that thank to the lemmas D.8 and D.10 all S_i are well-formed and promise-free-consistent.

Base Case: $P(i = 0)$.

$S_1.E = \{e_0\}$. In this case $S_1.jo = \emptyset$. Thus PROMISE-FREE is trivially satisfied.

Induction Step: $P(i) \Rightarrow P(i + 1)$.

Note the following holds (below r_i is short for $S.r_i$):

$$\begin{aligned} \rho_{i+1} &= \rho_i \cup \rho_{i+1}; [e_{i+1}] \\ \mathbf{jf}_{i+1} &= \mathbf{jf}_i \cup \mathbf{jf}_{i+1}; [e_{i+1}] \\ \mathbf{jo}_{i+1} &= \mathbf{jo}_i; ((\rho_{i+1} \cup \mathbf{jf}_{i+1}); [e_{i+1}])^? \cup \\ &\quad \cup \mathbf{jf}_{i+1}; [e_{i+1}] \setminus \rho_{i+1} \end{aligned}$$

Using these equalities we prove PROMISE-FREE by case analysis on e_{i+1} .

- $e_{i+1} \in S.F$

In this case $\mathbf{jf}_{i+1}; [e_{i+1}] = \emptyset$ and thus $\mathbf{jo}_{i+1} = \mathbf{jo}_i; (\rho_{i+1}; [e_{i+1}])^?$. If e_{i+1} has no ρ_{i+1} -predecessors, then $\mathbf{jo}_{i+1} = \mathbf{jo}_i$ and thus we have $\mathbf{cf}_{i+1} \cap \mathbf{jo}_{i+1} = \mathbf{cf}_i \cap \mathbf{jo}_i = \emptyset$. Otherwise we have

$$\begin{aligned} \mathbf{cf}_{i+1} \cap \mathbf{jo}_{i+1} &= \mathbf{cf}_i \cap \mathbf{jo}_i \cup \\ &\quad \cup \mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \rho_{i+1}; [e_{i+1}]) \end{aligned}$$

We have $\mathbf{cf}_i \cap \mathbf{jo}_i = \emptyset$ by inductive assumption. We prove $\mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \rho_{i+1}; [e_{i+1}]) = \emptyset$ by contradiction. Let $p \in \text{dom}(\mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \rho_{i+1}; [e_{i+1}]))$. Note that $p \in S.E_i$. Let e_j be immediate ρ_{i+1} -predecessor of e_{i+1} . Then $p \in \text{dom}(\mathbf{jo}_{i+1}; \rho_{i+1}; [e_{i+1}])$ implies $p \in \text{dom}(\mathbf{jo}_i; [e_j])$. Event e_j cannot be in conflict with p , since it would violate inductive assumption $\mathbf{cf}_i \cap \mathbf{jo}_i = \emptyset$. Thus either $\langle p, e_j \rangle \in \rho_i$ or $\langle e_j, p \rangle \in \rho_i$. If $\langle p, e_j \rangle \in \rho_i$ then $\langle p, e_{i+1} \rangle \in \rho_{i+1}$, but it contradicts $\langle p, e_{i+1} \rangle \in \mathbf{cf}_{i+1}$. If $\langle e_j, p \rangle \in \rho_i$, then since $\langle p, e_j \rangle \in \mathbf{jo}_i$ by $\mathbf{jo}; \rho \subseteq \mathbf{jo}$, which follows immediately from the definition of \mathbf{jo} , we have $\langle p, p \rangle \in \mathbf{jo}_i$. But $\mathbf{jo}_i \subseteq (\rho_i \cup \mathbf{jf}_i)^+$ and thus we have contradiction with NO-THIN-AIR. Therefore $\mathbf{cf}_{i+1} \cap \mathbf{jo}_{i+1} = \emptyset$.

- $e_{i+1} \in S.W$

This case is similar to the previous one.

- $e_{i+1} \in S.R$

In this case we have:

$$\begin{aligned} \mathbf{cf}_{i+1} \cap \mathbf{jo}_{i+1} &= \mathbf{cf}_i \cap \mathbf{jo}_i \cup \\ &\quad \cup \mathbf{cf}_{i+1} \cap (\mathbf{jf}_{i+1}; [e_{i+1}] \setminus \rho_{i+1}) \cup \\ &\quad \cup \mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \rho_{i+1}; [e_{i+1}]) \cup \\ &\quad \cup \mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \mathbf{jf}_{i+1}; [e_{i+1}]) \end{aligned}$$

By inductive assumption $\mathbf{cf}_i \cap \mathbf{jo}_i = \emptyset$. By well-formedness $\mathbf{cf}_{i+1} \cap \mathbf{jf}_{i+1} = \emptyset$ and thus $\mathbf{cf}_{i+1} \cap (\mathbf{jf}_{i+1}; [e_{i+1}] \setminus \rho_{i+1}) = \emptyset$. $\mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \rho_{i+1}; [e_{i+1}]) = \emptyset$ by the same argument as in cases $e_{i+1} \in S.F$ and $e_{i+1} \in S.W$.

We only left to show that $\mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \mathbf{jf}_{i+1}; [e_{i+1}]) = \emptyset$. Let e_j be immediate po_{i+1} -predecessor of e_{i+1} and let $p \in \text{dom}(\mathbf{cf}_{i+1} \cap (\mathbf{jo}_i; \mathbf{jf}_{i+1}; [e_{i+1}]))$. By lemma D.20 the branch of e_{i+1} should observe promise p up to the point of immediate conflict. Therefore if $\langle p, e_j \rangle \in \mathbf{cf}_i$ it should also be true that $\langle p, e_j \rangle \in \mathbf{jo}_i$. But it would violate the inductive assumption $\mathbf{cf}_i \cap \mathbf{jo}_i = \emptyset$. Thus $\langle p, e_j \rangle \notin \mathbf{cf}_i$.

Then either $\langle p, e_j \rangle \in \text{po}_i$ or $\langle e_j, p \rangle \in \text{po}_i$. The former cannot be true since then we would have $\langle p, e_{i+1} \rangle \in \text{po}_{i+1}$, which contradicts $\langle p, e_{i+1} \rangle \in \mathbf{cf}_{i+1}$. Therefore $\langle e_j, p \rangle \in \text{po}_i$.

Let r be an immediate po_i -successor of e_j s.t. $\langle r, p \rangle \in \text{po}_i$. By definition of immediate conflict we have $\langle r, e_{i+1} \rangle \in \mathbf{cf}_{i+1}|_{\text{imm}}$. By well-formedness we have that $r \in S.R^{\text{rlx}}$ and $S.\text{loc}(r) = S.\text{loc}(e_{i+1})$. Also, by lemma D.18 we get $\langle r, p \rangle \in \text{rpo}$. Let w be a write that justifies e_{i+1} , i.e., $\langle w, e_{i+1} \rangle \in \mathbf{jf}_{i+1}$. By well-formedness $S.\text{loc}(w) = S.\text{loc}(e_{i+1})$. Then $\langle p, e_{i+1} \rangle \in \mathbf{jo}_{i+1}; \mathbf{jf}_{i+1}; [e_{i+1}]$ implies $\langle p, w \rangle \in \mathbf{jo}_i$ which, in turn, implies $\langle p, w \rangle \in (\text{po}_i \cup \mathbf{jf}_i)^+$. Since $\langle r, p \rangle \in \text{rpo}_i$ we also have $\langle r, w \rangle \in \text{rpo}_i; (\mathbf{jf}_i \setminus \text{po}_i); (\text{po}_i \cup \mathbf{jf}_i)^+$. Let $E = \lfloor w \rfloor_{(\text{po}_i \cup \mathbf{jf}_i)^+}$. Note that $r, w \in E$. Because S_i is PROMISE-FREE by D.17 we have that $G = S|_E$ is RC11-consistent execution graph. By D.14 we also have that $(\text{po}_i \cup \mathbf{jf}_i)^+ = (\text{po}_i \cup \mathbf{rf}_i)^+$ and thus $\langle r, w \rangle \in \text{rpo}_i; (\mathbf{rf}_i \setminus \text{po}_i); (G.\text{po}_i \cup G.\mathbf{rf}_i)^+$. To show that $\langle r, w \rangle$ forms a load-buffering race in G we only have to prove that $\langle r, w \rangle \notin \mathbf{hb}_i$.

Suppose $\langle r, w \rangle \in \mathbf{hb}_i$. Since $\langle r, w \rangle \notin \text{po}_i$, by definition of \mathbf{hb}_i , there should exist either.

- Write $w_{\text{rel}} \in S.W^{\text{rel}}$, s.t. $\langle r, w_{\text{rel}} \rangle \in S.\text{po}_i$ and $\langle w_{\text{rel}}, w \rangle \in \mathbf{hb}_i \setminus \text{po}_i$. But in this case we have $\langle w_{\text{rel}}, e_{i+1} \rangle \in \mathbf{cf}_{i+1} \cap \mathbf{jo}_{i+1}$, i.e., w_{rel} is a promise. Then by lemma D.19 $w_{\text{rel}} \in S.W^{\text{rlx}}$ which contradicts the fact that w_{rel} is a release write.
- Fence $f_{\text{rel}} \in S.F^{\text{rel}}$, s.t. $\langle r, f_{\text{rel}} \rangle \in \text{po}_i$ and $\langle f_{\text{rel}}, w \rangle \in \mathbf{hb}_i$. But in this case $\langle f_{\text{rel}}, p \rangle \in \text{po}_i$ which contradicts the fact that $\langle r, p \rangle \in \text{rpo}_i$.

Therefore the pair $\langle r, w \rangle$ forms a load-buffering race in RC11-consistent execution G , which contradicts our assumption that P is LB-race-free. Hereby, by contradiction we have proved that $\mathbf{jo}_{i+1}; \mathbf{jf}_{i+1}; [e_{i+1}] = \emptyset$, which implies $\mathbf{cf}_{i+1} \cap \mathbf{jo}_{i+1} = \emptyset$. □

THEOREM D.22. *Let P be LB-race-free under RC11. Then the set of its RC11 consistent execution graphs coincides with the set of its Weakestmo2 consistent graphs.*

PROOF. By lemma D.2 every RC11 consistent graph of P is also Weakestmo consistent.

Let G be Weakestmo consistent, i.e., there exists Weakestmo consistent event structure S , s.t. $S \triangleright G$. By lemma D.21 S is PROMISE-FREE. Therefore, by lemma D.15, G is RC11 consistent execution graph of P . □

E EXTENDING THE LBRF THEOREM TO THE FULL WEAKESTM02 MODEL

In this section, we describe how the proof of the $\text{LBRF}_{\text{RC11}}$ theorem can be extended to account for the remaining features of Weakestmo2, namely read-modify-write atomics and sequentially consistent accesses/fences. We observe that there is a certain difficulty in achieving this goal and thus we propose two different solutions to overcome it. We admit that neither solution is ideal and requires a certain trade-off.

But first, let us present the version of RC11 consistency covering all the features of C11.

Definition E.1. Execution graph G is RC11 consistent if the following hold:

- $\text{po} \cup \mathbf{rf}$ is acyclic (RF-NO-THIN-AIR)
- $\mathbf{hb}; \mathbf{eco}^?$ is irreflexive. (COHERENT)
- $\mathbf{rmw} \cap (\mathbf{fr}; \mathbf{co}) = \emptyset$. (RMW-ATOMIC)

- **psc** is acyclic. (SEQUENTIAL-CONSISTENCY)

Next, note that in the proof of D.21 at some point we have to construct RC11 consistent execution graph from the $\text{po} \cup \text{jf}$ prefix of the event structure. There, the fact that $\text{po} \cup \text{jf}$ prefix would satisfy RC11 consistency constraints followed almost trivially. In the full model, however, the prefix should also satisfy **RMW-ATOMIC** and **IMM-SEQUENTIAL-CONSISTENCY**. There are two options how it can be achieved.

The first obvious option would be to strengthen the Weakestmo2 consistency further and explicitly require from each $\text{po} \cup \text{jf}$ conflict-free prefix of the event structure to satisfy these constraints. While this strengthening looks reasonable and harmless, we have observed that at least in the case of **RMW-ATOMIC** axiom, it has undesirable consequences. The model becomes strictly stronger and the following transformation ends up to be unsound:

$$a : R_{\text{ex}}^{rlx} ; b : W_{\text{ex}}^{rlx} ; c : W^{rlx} \rightsquigarrow c : W^{rlx} ; a : R_{\text{ex}}^{rlx} ; b : W_{\text{ex}}^{rlx}$$

The second option is to relax the definition of LB race and try to prove that for every LB-race-free program (according to the revised definition) each $\text{po} \cup \text{jf}$ conflict-free prefix of the event structure would satisfy **RMW-ATOMIC**. We managed to do that with the assumption that P has no **sc** accesses or fences. Below we present the proof.

Definition E.2. Given a well-formed execution graph G a pair of events $\langle e_1, e_2 \rangle$ is called a *load buffering race* if it satisfies the following:

- $G.\text{loc}(e_1) = G.\text{loc}(e_2)$
- $\langle e_1, e_2 \rangle \notin G.\text{hb} \cup G.\text{hb}^{-1}$
- Either
 - $\langle e_1, e_2 \rangle \in [R^{rlx}] ; \text{rpo} ; (\text{rf} \setminus \text{po}) ; (\text{po} \cup \text{rf})^* ; [W]$ or
 - $\langle e_1, e_2 \rangle \in [W^{rlx} \cup R_{\text{ex}}^{rlx}] ; (\text{po} \cup \text{rf})^+ \setminus \text{rf} ; [R_{\text{ex}}^{rlx}]$

Definition E.3. Event structure S is *prefix-rmw-atomic* if for every event $e \in S.E$ its $(\text{po} \cup \text{jf})$ prefix satisfies atomicity, i.e., for $E = \lfloor e \rfloor_{(\text{po} \cup \text{jf})^*}$ the following is true:

- $\text{rmw}|_E \cap (\text{fr}|_E ; \text{co}|_E) = \emptyset$ (PREFIX-RMW-ATOMIC)

Definition E.4. A triple of events $\langle w, r_{\text{ex}}, w_{\text{ex}} \rangle$ is an atomicity violation in an execution graph G if

- $\langle w_{\text{ex}}, r_{\text{ex}} \rangle \in G.\text{rmw}$,
- $\langle r_{\text{ex}}, w \rangle \in G.\text{fr}$, and
- $\langle w, w_{\text{ex}} \rangle \in G.\text{co}$.

Definition E.5. An atomicity violation $\langle w, r_{\text{ex}}, w_{\text{ex}} \rangle$ in execution graph G is *co-repairable* if the following conditions are met:

- $w \notin G.W_{\text{ex}}$.
- Either w or w_{ex} is $(\text{po} \cup \text{rf})^+$ maximal.

Definition E.6. Two executions are equivalent modulo **co**, written $G \cong_{\text{co}} G'$, if they agree on all components except **co**, that is:

- (1) $G'.E = G.E$.
- (2) $G'.\text{tid} = G.\text{tid}$
- (3) $G'.\text{lab} = G.\text{lab}$
- (4) $G'.\text{po} = G.\text{po}$.
- (5) $G'.\text{rf} = G.\text{rf}$.

LEMMA E.7. *Let G be an execution graph satisfying all the RC11-consistency constraints except **RMW-ATOMIC**. Suppose G contains exactly one atomicity violation, $\langle w, r_{\text{ex}}, w_{\text{ex}} \rangle$, and further suppose*

that the violation is **co**-repairable. Then there exists RC11-consistent execution graph G' equivalent to G modulo **co**.

PROOF. Let $G' \triangleq \langle G.E, G.tid, G.lab, G.po, G.rf, id, co' \rangle$ where

$$\Delta_{co} = \{ \langle w, w_{ex} \rangle \} \quad \text{and} \quad co' = G.co \setminus \Delta_{co} \cup \Delta_{co}^{-1}$$

That is, co' is the same as **co** except that we swapped a single edge from $\langle w, w_{ex} \rangle$ to $\langle w_{ex}, w \rangle$. Note that G' trivially satisfies condition $G' \cong_{co} G$.

Let us show that co' is well-formed, i.e., it is a total order. We only need to show that irreflexivity since the totality would follow from the totality of **co**.

Consider the following inequality:

$$\begin{aligned} & (G.co \setminus \Delta_{co} \cup \Delta_{co}^{-1}) ; (G.co \setminus \Delta_{co} \cup \Delta_{co}^{-1}) = \\ & = (G.co \setminus \Delta_{co}) ; (G.co \setminus \Delta_{co}) \cup \\ & \cup (G.co \setminus \Delta_{co}) ; \Delta_{co}^{-1} \cup \\ & \cup \Delta_{co} ; (G.co \setminus \Delta_{co}^{-1}) \stackrel{?}{\subseteq} \\ & \stackrel{?}{\subseteq} G.co \setminus \Delta_{co} \cup \Delta_{co}^{-1} \end{aligned}$$

First note that $(G.co \setminus \Delta_{co}) ; (G.co \setminus \Delta_{co}) \subseteq (G.co \setminus \Delta_{co})$. Indeed, suppose there exists w' , s.t. $\langle w, w' \rangle \in G.co$ and $\langle w', w_{ex} \rangle \in G.co$. But then, since $\langle w_{ex}, w \rangle \in fr$ we would also have $\langle w_{ex}, w' \rangle \in fr$. Therefore $\langle w', r_{ex}, w_{ex} \rangle$ would form an atomicity violation in G , which contradicts the assumption that $\langle w, r_{ex}, w_{ex} \rangle$ is single atomicity violation. For the same reasons $(G.co \setminus \Delta_{co}) ; \Delta_{co}^{-1} \subseteq G.co \setminus \Delta_{co}$ and $\Delta_{co} ; (G.co \setminus \Delta_{co}^{-1}) \subseteq G.co \setminus \Delta_{co}$. Therefore we have:

$$co' \triangleq (G.co \setminus \Delta_{co} \cup \Delta_{co}^{-1})^+ = G.co \setminus \Delta_{co} \cup \Delta_{co}^{-1}$$

Note that irreflexivity of relation r is equivalent to $r \cap id = \emptyset$ and thus irreflexivity of the union of relations is equivalent to irreflexivity of components. Irreflexivity $co \setminus \Delta_{co}$ follows from irreflexivity of **co**, and Δ_{co}^{-1} is trivially irreflexive. Therefore co' is well-formed and thus G' is well-formed as well.

Since $G'.po = G.po$, $G'.rf = G.rf$, and G satisfies **RF-NO-THIN-AIR** we have that G' also satisfies **RF-NO-THIN-AIR**.

Because $\langle w, r_{ex}, w_{ex} \rangle$ is a single atomicity violation in G and it is clearly not an atomicity violation in G' (because $\langle w, w_{ex} \rangle \notin co'$ as was shown above) then G' satisfies **RMW-ATOMIC**.

Also note that since **hb** is defined in terms of **lab**, **po**, and **rf** we have that $G'.hb = G.hb$. Therefore $G.hb'$ is irreflexive.

We need to show $G'.hb ; G'.eco^?$ is irreflexive.

First, note that for any well-formed execution graph (or, more general, event structure) the following equation can be proven:

$$\begin{aligned} eco & \triangleq (rf \cup co \cup fr)^+ = \\ & = rf \cup co ; rf^? \cup fr ; rf^? \end{aligned}$$

Then observe that:

$$G'.eco = G.rf \cup ; G'.co ; G.rf^? \cup G.rf^{-1} ; G'.co ; G.rf^?$$

Irreflexivity of the following components follows from the fact that G is **COHERENT**.

- $G.hb ; G.rf$ is irreflexive.

- $G.\text{hb}$; $((G.\text{co} \setminus \Delta_{\text{co}}); G.\text{rf}^?)$ is irreflexive.
- $G.\text{hb}$; $(G.\text{rf}^{-1}; (G.\text{co} \setminus \Delta_{\text{co}}); G.\text{rf}^?)$ is irreflexive.

We only left to show irreflexivity of the following components

- $G.\text{hb}$; $\Delta_{\text{co}}^{-1}; \text{rf}^?$.
- $G.\text{hb}$; $\text{rf}^{-1}; \Delta_{\text{co}}^{-1}; \text{rf}^?$.

Note that $\forall e \in G.E. \langle e, w_{\text{ex}} \rangle \in \text{hb} \Leftrightarrow \langle e, r_{\text{ex}} \rangle \in \text{hb}$. Suppose that $\langle w, w_{\text{ex}} \rangle \in G.\text{rf}^?; G.\text{hb}$. Then it would imply $\langle w, r_{\text{ex}} \rangle \in G.\text{rf}^?; G.\text{hb}$ and therefore we would have a cycle $\langle r_{\text{ex}}, r_{\text{ex}} \rangle \in G.\text{fr}; G.\text{rf}^?; G.\text{hb}$ which contradicts the assumption that G is **COHERENT**.

Next, suppose that $\langle w, w_{\text{ex}} \rangle \in G.\text{rf}^?; G.\text{hb}; G.\text{rf}^{-1}$. By our assumption either w or w_{ex} is $(G.\text{po} \cup G.\text{rf})^+$ maximal. If w_{ex} is $(G.\text{po} \cup G.\text{rf})^+$ maximal then $G.\text{rf}^{-1}; [w_{\text{ex}}]$ is empty. Similarly, $[w]; G.\text{rf}$ is empty if w is maximal. Thus we have only left to consider the case $\langle w, w_{\text{ex}} \rangle \in G.\text{hb}; G.\text{rf}^{-1}$ and w is $(G.\text{po} \cup G.\text{rf})^+$ maximal. But then notice that $G.\text{hb} \subseteq (G.\text{po} \cup G.\text{rf})^+$ and thus $[w]; G.\text{hb}$ is empty as well. \square

Definition E.8. An atomicity violation $\langle w'_{\text{ex}}, r_{\text{ex}}, w_{\text{ex}} \rangle$ in execution graph G is **rf-repairable** if:

- $w'_{\text{ex}} \in G.W_{\text{ex}}$ and
- w_{ex} is $\text{po} \cup \text{rf}$ maximal.

LEMMA E.9. *Let G be an execution graph satisfying all the RC11-consistency constraints except **RMW-ATOMIC**. Let $\langle w'_{\text{ex}}, r_{\text{ex}}, w_{\text{ex}} \rangle$ be a single atomicity violation and let it be a **rf-repairable**. Then there exists RC11 consistent execution graph G' equivalent to G modulo $\text{rf}(w'_{\text{ex}}, r_{\text{ex}})$, denoted as $G' \cong_{\text{rf}(w'_{\text{ex}}, r_{\text{ex}})} G$, in the following sense:*

- (1) $G'.E = G.E \setminus \{w_{\text{ex}}, r_{\text{ex}}\} \cup \{r'_{\text{ex}}\}$ where
 - $r'_{\text{ex}} \triangleq \langle G.\text{id}(r_{\text{ex}}), G.\text{tid}(r_{\text{ex}}), \text{lab} \rangle$
 - $\text{lab} \triangleq R_{\text{ex}}^o(l, v)$
 - $o \triangleq G.\text{mod}(r_{\text{ex}}) \quad l \triangleq G.\text{loc}(r_{\text{ex}}) \quad v \triangleq G.\text{val}(w'_{\text{ex}})$
- (2) $G'.\text{po} = G.\text{po}|_{G.E'}$.
- (3) $G'.\text{rf} = G.\text{rf}|_{G.E' \setminus \{r_{\text{ex}}\}} \cup \Delta_{\text{rf}}$ where $\Delta_{\text{rf}} \triangleq \{\langle w'_{\text{ex}}, r'_{\text{ex}} \rangle\}$
- (4) $G'.\text{co} = G.\text{co}|_{G.E'}$.

PROOF. It is easy to show that G' is well-formed.

Since w_{ex} is $\text{po} \cup \text{rf}$ maximal in G then r_{ex} is $\text{po} \cup \text{rf}$ maximal in G' . Thus adding Δ_{rf} cannot introduce any cycle in $G.\text{po} \cup G.\text{rf}$ and therefore $G'.\text{po} \cup G'.\text{rf}$ is acyclic, i.e., G' sat. **RF-NO-THIN-AIR**.

The fact that G' is **RMW-ATOMIC** follows from assumption that $\langle w'_{\text{ex}}, r_{\text{ex}}, w_{\text{ex}} \rangle$ is the only atomicity violation in G . Since we have removed w_{ex} from the graph G' has no other atomicity violations.

Thus we only left to show that $G'.\text{hb}; G'.\text{eco}^?$ is irreflexive.

- $G'.\text{hb}$ is irreflexive since $G'.\text{hb} \subseteq (G.\text{po} \cup G.\text{rf})^+$.
- Irreflexivity of $G'.\text{hb}; G'.\text{rf}$ follows from irreflexivity of $G.\text{hb}; G.\text{rf}$ and the fact that r'_{ex} is $(G'.\text{po} \cup G'.\text{rf})$ maximal.
- Irreflexivity of $G'.\text{hb}; G'.\text{co} = G'.\text{hb}; G.\text{co} = G.\text{hb}; G.\text{co}$ follows directly.
- Let us consider $G'.\text{hb}; G'.\text{co}; G'.\text{rf}^?$.

$$\begin{aligned} G'.\text{hb}; G'.\text{co}; G'.\text{rf}^? &= G.\text{hb}; G.\text{co}; G'.\text{rf} = \\ &G.\text{hb}; G.\text{co}; (G.\text{rf} \cup \Delta_{\text{rf}}) \end{aligned}$$

But the cycle $\langle r'_{\text{ex}}, r'_{\text{ex}} \rangle \in G.\text{hb}; G.\text{co}; \Delta_{\text{rf}}$ would contradict $G'.\text{po} \cup G'.\text{rf}$ maximality of r'_{ex} .

- Finally, suppose there is cycle $\langle r'_{\text{ex}}, r'_{\text{ex}} \rangle \in G'.\text{fr}; G'.\text{rf}^?; G'.\text{hb}$. Let w be the endpoint of the first edge in the cycle, i.e., $\langle r'_{\text{ex}}, w \rangle \in G'.\text{fr}$. Then we would have $\langle w'_{\text{ex}}, w \rangle \in G.\text{co}$ and therefore $\langle w, r_{\text{ex}} \rangle \in G.\text{fr}; G.\text{rf}^?; G.\text{hb}$ would also close a cycle in $G.\text{hb}; G.\text{eco}$.

□

Finally, we need to adjust the proof of the lemma D.21. Namely, along proving by induction that each S_i is **PROMISE-FREE**, we also need to prove that it is **PREFIX-RMW-ATOMIC**.

PROOF. The base of the induction is trivial again, since $S_0.\text{rmw} = \emptyset$. Thus we only consider the step of the induction. Note the following holds:

$$\begin{aligned} \text{rmw}_{i+1} &= \text{rmw}_i \cup \text{rmw}_{i+1}; [e_{i+1}] \\ \text{rmw}_{i+1} \cap (\text{fr}_{i+1}; \text{co}_{i+1}) &= \\ &= \text{rmw}_i \cap (\text{fr}_i; \text{co}_i) \\ &\cup \text{rmw}_i \cap (\text{fr}_i; [e_{i+1}]; \text{co}_i) \\ &\cup (\text{rmw}_{i+1}; [e_{i+1}]) \cap (\text{fr}_i; \text{co}_{i+1}; [e_{i+1}]) \end{aligned}$$

Using these equalities we prove **PREFIX-RMW-ATOMIC** by case analysis on e_{i+1} .

- $e_{i+1} \in S.F$

In this case $\text{fr}_{i+1}; [e_{i+1}] = \emptyset$, $\text{co}_{i+1}; [e_{i+1}] = \emptyset$, and $\text{rmw}_{i+1}; [e_{i+1}] = \emptyset$. Thus $\text{rmw}_{i+1} \cap (\text{fr}_{i+1}; \text{co}_{i+1}) = \text{rmw}_i \cap (\text{fr}_i; \text{co}_i) = \emptyset$. Therefore, for every event $e \in S.E_{i+1}$, $E = [e]_{(\text{po} \cup \text{jf})^*}$ the required property $\text{rmw}_{i+1}|_E \cap (\text{fr}_{i+1}|_E; \text{co}_{i+1}|_E) = \emptyset$ follows immediately.

- $e_{i+1} \in S.R$

This case is similar to the previous one.

- $e_{i+1} \in S.W$

For every $e \in S.E_i$ the **rmw**-atomicity of its prefix follows immediately from the inductive assumption. Thus we only need to show that prefix of e_{i+1} satisfies **rmw**-atomicity. Let $E = [e]_{(\text{po} \cup \text{jf})^*}$. We need to show:

$$\begin{aligned} \text{rmw}_{i+1}|_E \cap (\text{fr}_{i+1}|_E; \text{co}_{i+1}|_E) &= \\ &= \text{rmw}_i|_E \cap (\text{fr}_i|_E; \text{co}_i|_E) \\ &\cup \text{rmw}_i|_E \cap (\text{fr}_i|_E; [e_{i+1}]; \text{co}_i|_E) \\ &\cup (\text{rmw}_{i+1}|_E; [e_{i+1}]) \cap (\text{fr}_i|_E; \text{co}_{i+1}|_E; [e_{i+1}]) \end{aligned}$$

– $\text{rmw}_i|_E \cap (\text{fr}_i|_E; \text{co}_i|_E) = \emptyset$ by inductive assumption.

– Let $\langle r_{\text{ex}}, w_{\text{ex}} \rangle \in \text{rmw}_i|_E$, s.t. $\langle r_{\text{ex}}, e_{i+1} \rangle \in \text{fr}_{i+1}$ and $\langle e_{i+1}, w_{\text{ex}} \rangle \in \text{co}_{i+1}$. Note that by well-formedness we have that $r_{\text{ex}} \in S_i.R_{\text{ex}}$ and we have $\langle r_{\text{ex}}, e_{i+1} \rangle, \langle w_{\text{ex}}, e_{i+1} \rangle \in (\text{po} \cup \text{jf})^+$. Consider the cases for e_{i+1} .

- * $e_{i+1} \in S.W_{\text{ex}}$.

Let e_j be immediate **po** predecessor if e_{i+1} . Then by well-formedness $\langle e_j, e_{i+1} \rangle \in S_{i+1}.\text{rmw}$. Then $\langle w_{\text{ex}}, e_j, e_{i+1} \rangle$ forms **rf**-repairable atomicity violation in E . Note that from inductive assumption it follows that it is the only atomicity violation in E . Then, giving that S_{i+1} is **PROMISE-FREE**, using the lemmas D.17 and E.5, we can conclude that there exists RC11 consistent execution graph G , s.t. $G \cong_{\text{rf}(e_j)} S_{i+1}|_E$. We have $\langle r_{\text{ex}}, e_j \rangle \in [R_{\text{ex}}]$; $(G.\text{po} \cup G.\text{rf})^+; [R_{\text{ex}}]$, i.e., the pair $\langle r_{\text{ex}}, e_j \rangle$ forms load-buffering race in RC11-consistent execution G , which contradicts our assumption that P is LB-race-free.

- * $e_{i+1} \notin S.W_{\text{ex}}$.

Then $\langle e_{i+1}, r_{\text{ex}}, w_{\text{ex}} \rangle$ forms a **co**-repairable atomicity violation in E . Note that from inductive assumption it follows that it is the only atomicity violation in E . Then, giving that S_{i+1} is **PROMISE-FREE**, using the lemmas D.17 and E.5, we can conclude that there exists RC11 consistent execution graph G , s.t. $G \cong_{\text{co}} S_{i+1}|_E$. Since S_{i+1} is **PROMISE-FREE**, by D.14

we also have that

$$\begin{aligned} (\text{po}_{i+1} \cup \mathbf{jf}_{i+1})^+ &= (\text{po}_{i+1} \cup \mathbf{rf}_{i+1})^+ \\ &= (G.\text{po} \cup G.\mathbf{rf})^+ \end{aligned}$$

- Therefore the pair $\langle r_{\text{ex}}, e_{i+1} \rangle \in [\mathbf{R}_{\text{ex}}] ; (G.\text{po} \cup G.\mathbf{rf})^+ ; [\mathbf{R}_{\text{ex}}]$, forms a load-buffering race in RC11-consistent execution G , which contradicts our assumption that P is LB-race-free.
- Let $\langle e_j, e_{i+1} \rangle \in \mathbf{rmw}_{i+1}$. Let $w \in E$, s.t. $\langle e_j, w \rangle \in \mathbf{fr}_i$ and $\langle w, e_{i+1} \rangle \in \mathbf{co}_{i+1}$. Then similarly to the previous cases, we can show that there exists a RC11-consistent execution graph G of program P , that contains a load-buffering race.

□

F CORRECTNESS OF COMPILATION MAPPINGS FOR WEAKESTMO2

In this section, we present the proof of the correctness of compilation from Weakestmo2 to IMM.

THEOREM F.1. *Let P be a program, and G be an IMM-consistent execution graph of P . Then, there exists a Weakestmo2-consistent event structure S of P such that $S \triangleright G$.*

Our development is based on the proof of the same statement for the original Weakestmo model by Moiseenko et al. [2020].

F.1 Recap of the Original Proof Structure

Moiseenko et al. [2020] construct the required event structure S step by step following a *traversal* of the IMM graph G [Moiseenko et al. 2020, §2.3]. Traversal of the graph G induces operational small-step semantics $G \vdash TC \xrightarrow{e} TC'$ where TC and TC' are *traversal configurations* and e is an event being traversed. A traversal configuration is a tuple $\langle C, I \rangle$, where $C \subseteq G.E$ is a set of *covered events* and $I \subseteq G.W$ is a set of *issued writes*.

Giving the operational semantics of traversal $G \vdash TC \xrightarrow{e} TC'$, such that $G.\text{tid}(e) = t$, and the operational semantics of event structure construction $S \xrightarrow{e} S'$ the proof then proceeds using the standard *simulation* argument. Moiseenko et al. [2020, §4.2] define a *simulation relation* $I(P, T, G, TC, S, X)$ between the program P , a set of thread identifiers $T \subseteq \text{Tid}$, the current traversal configuration TC of execution graph G , the current Weakestmo-consistent event structure S , and a justified configuration X of S . They then prove the following three lemmas which state that (i) the initial traversal configuration and initial event structure are related by the simulation relation, (ii) each traversal step can be simulated by an event structure construction step, and (iii) it is possible to extract the traversed execution graph from an event structure corresponding to the terminal traversal configuration.

LEMMA F.2 (SIMULATION START). *Let P be a program, and G be an IMM-consistent execution graph of P . Then $I(P, \text{tid}(P), G, TC_{\text{init}}(G), S_{\text{init}}(P), \{e_0\})$ holds where*

- $\text{tid}(P)$ the set of threads of program P ;
- $TC_{\text{init}}(G) \triangleq \langle e_0, e_0 \rangle$ is initial traversal configuration containing only the initialization event;
- $S_{\text{init}}(P)$ is initial event structure also containing only the initialization event.

LEMMA F.3 (WEAK SIMULATION STEP). *If $I(P, T, G, TC, S, X)$ and $G \vdash TC \rightarrow TC'$ hold, then there exist S' and X' such that $I(P, T, G, TC', S', X')$ and $S \rightarrow^* S'$ hold.*

LEMMA F.4 (SIMULATION END). *If $I(P, \text{tid}(P), G, TC_{\text{final}}(G), S, X)$ holds, where $TC_{\text{final}}(G) \triangleq \langle G.E, G.E \rangle$ is a terminal traversal configuration, then the execution graph associated with X is isomorphic to G , or in other words, G can be extracted from S : $S \triangleright G$.*

Lemma F.3 is the most challenging lemma to prove. Simulating a traversal step might require to construct a new *certification branch* Br_{cert} in the event structure containing multiple events [Moiseenko et al. 2020, §4.3]. To construct this branch, one must (i) pick a suitable justification write for each read event [Moiseenko et al. 2020, §4.3.1], and (ii) chose the position of each write event in the event structure’s coherence order [Moiseenko et al. 2020, §4.3.2]. Once the certification branch Br_{cert} is constructed, it is used to replace the events of the corresponding thread in the configuration X .

$$X' \triangleq X \setminus Br_{\text{prev}} \cup Br_{\text{cert}}$$

$$Br_{\text{prev}} \triangleq X|_t \triangleq \{e \in X \mid S.\text{tid}(e) = t\}$$

The new constraint **NO-BAIT-AND-SWITCH** added to Weakestmo2 applies exactly to the read events in two adjacent branches of an event structure. Thus, in order to adjust the proof for Weakestmo2 model, we need to show that the justification writes chosen for reads in a certification branch match those justification writes in the previous branch that issued the promises.

F.2 Simulation Relation

Before going into the details of the modified proof for Weakestmo2, we highlight some invariants guaranteed by the simulation relation \mathcal{I} . We present here only some of invariants derived from \mathcal{I} rather than full simulation relation for two reasons. First, the presented invariants should help to develop some intuition about the construction without going into too much technical details. Second, the proof adaptation requires only this subset of simulation relation’s properties. Curious reader may find the full simulation relation in [Moiseenko et al. 2020, §A] or in the Coq developments accompanying that paper <https://github.com/weakmemory/weakestmoToImm/>.

The simulation relation $\mathcal{I}(P, T, G, TC, S, X)$ establishes a connection between the event structure S and execution graph G with the help of a function $s2g : S.E \rightarrow G.E$ which maps an event of S to an event of G . This function can be lifted to sets of events in the following way⁷:

$$\text{for } A_S \subseteq S.E : \llbracket A_S \rrbracket \triangleq \{s2g(e) \in G.E \mid e \in A_S\}$$

$$\text{for } A_G \subseteq G.E : \llbracket A_G \rrbracket \triangleq \{e \in S.E \mid s2g(e) \in A_G\}$$

The simulation relation \mathcal{I} contains the following properties.

- (1) Events of S restricted to threads T and events of X correspond to covered and issued events and their po-predecessors:
 - $\llbracket S.E|_T \rrbracket = \llbracket X \rrbracket = C \cup \text{dom}(G.\text{po}^? ; [I])$
- (2) Labels of events in S match labels of events in G modulo their values.
 - (a) $\forall e \in S.E. S.\{\text{tid}, \text{typ}, \text{loc}, \text{mod}\}(e) = G.\{\text{tid}, \text{typ}, \text{loc}, \text{mod}\}(\llbracket e \rrbracket)$
Labels of determined events match precisely.
 - (b) $\forall e \in X \cap \llbracket G.D_{(C,T)} \rrbracket. S.\text{val}(e) = G.\text{val}(\llbracket e \rrbracket)$
- (3) Program order in S corresponds to program order in G :
 - $\llbracket S.\text{po} \rrbracket \subseteq G.\text{po}$
- (4) Identity relation in G corresponds to identity or conflict relation in S :
 - $\llbracket \text{id} \rrbracket \subseteq S.\text{cf}^?$
- (5) Reads in S are justified by writes that have already been observed by the corresponding events in G .
 - (a) $\llbracket S.\text{jf} \rrbracket \subseteq G.\text{rf}^? ; G.\text{hb}^?$
Moreover, restriction of justified-from relation on X corresponds to *stable justification relation* (see F.3) of the current traversal configuration in graph.

⁷In a similar manner it is possible to lift $s2g$ to binary relations on events.

- (b) $\llbracket S.\text{jf}; [X] \rrbracket \subseteq G.\text{sjf}_{TC}$
 As a consequence it is possible to derive that justification for covered events in X corresponds to their justification in G : $\llbracket S.\text{jf}; [X \cap \llbracket C \rrbracket] \rrbracket \subseteq G.\text{rf}$.
 Finally, only issued events can be used for external justification.
- (c) $\text{dom}(S.\text{jfe}) \subseteq \text{dom}(S.\text{ew}^*; [X \cap \llbracket I \rrbracket])$
- (6) Equivalent writes in S are mapped to the same write in G .
- (a) $\llbracket S.\text{ew} \rrbracket \subseteq \text{id}$
 Also, every $S.\text{ew}^*$ equivalence class has a representative issued event in X .
- (b) $S.\text{ew}^* \subseteq (S.\text{ew}^*; [X \cap \llbracket I \rrbracket]; S.\text{ew}^*)^?$
- (7) Coherence edges in S correspond to coherence or identity edges in G .
- (a) $\llbracket S.\text{co} \rrbracket \subseteq G.\text{co}^?$
 Coherence edges in S ending in X restricted to threads T correspond precisely to coherence edges in G .
- (b) $\llbracket S.\text{co}; [X|_T] \rrbracket \subseteq G.\text{co}$
- (8) Synchronize-with and happens-before relations in S conform to the corresponding relations in G .
- (a) $\llbracket S.\text{sw} \rrbracket \subseteq G.\text{sw}$
- (b) $\llbracket S.\text{hb} \rrbracket \subseteq G.\text{hb}$

Let us summarize the description of the simulation relation with an informal explanation of its invariants. As we have mentioned, traversal of the execution graph captures a possible execution order of the program leading to the given graph. Covering an event corresponds to execution an instruction of the program in-order, while issuing a write corresponds to out-of-order execution.

The construction of the event structure reflects this execution strategy, with a notable difference being that in event structure each write event executed out-of-order must come together with a sequence of events that lead to this write. The first property of the simulation relation 1 formalizes this intuition. The receptiveness property [Podkopaev et al. 2019, §6.4] guarantees that it is always possible to construct an execution branch leading to a specific issued event, and that this branch will contain the same set of events with almost the same labels, except that values of some events might be different. Hence we have the properties 2 and 3. The property 4 further guarantees that events in the event structure which have the same image in graph G are either equal or in conflict.

The next few properties constrain the justification of reads. First, property 5a ensures that a read event in the event structure cannot be justified from a write located happens-before after the write this event reads-from in the graph. Property 5b asserts that justification of reads in the configuration X should match the *stable justification* relation in G . We postpone the detailed description of this relation until F.3. Property 5c assures that only issued events can be used to justify a read externally, i.e., by a write from another thread. This implies, in particular, that promises in the event structure correspond to issued events in the graph.

Properties 6a and 6b state that each equivalence class of writes in S correspond exactly to some issued write event in G . Finally, properties 7, 8a, and 8b relate $S.\text{co}$, $S.\text{sw}$, and $S.\text{hb}$ to their counterparts in the graph G .

To further develop intuition about the simulation process we recommend to consult an example in [Moiseenko et al. 2020, §2.4].

Remark. Note that in property 2b we use determined events $G.D_{(C,I)}$ (to be defined later in F.3), while Moiseenko et al. [2020] uses covered and issued events. As will become evident later $C \cup I \subseteq G.D_{(C,I)}$ and thus our property is stronger. Nevertheless, the construction of Moiseenko et al. [2020] actually allows to establish this stronger property, as we will see later. It just that for

the proof for Weakestmo model the weaker property was sufficient, while our modified proof will require the stronger property.

F.3 Stable Justification

To simulate a traversal step $G \vdash TC \xrightarrow{e} TC'$ of a thread $G.\text{tid}(e)$ that has outstanding promises (i.e., issued writes which are not covered yet), the event structure must take multiple steps to construct a certification branch. For each read in this branch, we must select a write event justifying this read. For this purpose, Moiseenko et al. [2020, §4.3.1] define the *stable justification relation* in a few steps.

First, *determined* events, $G.D_{\langle C, I \rangle}$, are events that must have the same label in G and in configuration X of S .

$$\begin{aligned} G.D_{\langle C, I \rangle} \triangleq & C \cup I \cup \\ & \cup G.W \setminus \text{rng}(G.\text{ppo}) \cup \\ & \cup \text{dom}(G.\text{rfi}^?; G.\text{ppo}; [I]) \cup \\ & \cup \text{rng}([I]; G.\text{rfi}) \cup \\ & \cup \text{rng}(G.\text{rfe}; [G.E^{\exists\text{acq}}]) \end{aligned}$$

Next, the *viewfront relation*, $G.\text{vf}_{TC}$, represents which write events have been observed by other events.

$$G.\text{vf}_{\langle C, I \rangle} \triangleq [G.W]; (G.\text{rf}; [C])^?; G.\text{hb}^? \cup G.\text{rf}; [G.D_{\langle C, I \rangle}]; G.\text{po}^?$$

Finally, *stable justification* picks **co**-maximal observed write.

$$G.\text{sjf}_{TC} \triangleq ([G.W]; (G.\text{vf}_{TC} \cap =_{G.\text{loc}}); [G.R]) \setminus (G.\text{co}; G.\text{vf}_{TC})$$

Remark. We have slightly modified the definition of determined events compared to the one in [Moiseenko et al. 2020, §4.3.1]. The modified version of the definition will help us later simplify our proofs. The original definition of determined events is shown below.

$$G.D_{\langle C, I \rangle} \triangleq C \cup I \cup \text{dom}(G.\text{rfi}^?; G.\text{ppo}; [I]) \cup \text{rng}([I]; G.\text{rfi}) \cup \text{rng}(G.\text{rfe}; [G.E^{\exists\text{acq}}])$$

The only difference of the modified version is that it also includes $G.W \setminus \text{rng}(G.\text{ppo})$, i.e., writes that do not depend on preceding read events, and so their values can be fixed in advance. We argue that this modification does not affect the proofs. The viewfront relation (and thus the stable justification relation too) in fact depends only on the set of determined reads. Making more write events determined does not affect it.

Viewfront and stable justification relations possess a number of useful properties.

PROPOSITION F.5. $G.\text{vf}_{TC}; G.\text{eco}$ is *irreflexive*.

PROPOSITION F.6. $G.\text{sjfi}_{TC} = G.\text{rfi}$

PROPOSITION F.7. $G.\text{sjf}_{TC}; [G.D_{TC}] = G.\text{rf}; [G.D_{TC}]$ In particular, it implies that given a traversal step $G \vdash TC \rightarrow TC'$, we have that:

$$G.\text{sjf}_{TC'}; [G.D_{TC}] = G.\text{rf}; [G.D_{TC}] = G.\text{sjf}_{TC}; [G.D_{TC}]$$

PROPOSITION F.8. Given a traversal step $G \vdash \langle C, I \rangle \rightarrow \langle C', I' \rangle$, we have that

$$G.\text{sjf}_{\langle C', I' \rangle} = [I]; G.\text{sjfe}_{\langle C', I' \rangle} \cup G.\text{sjfi}_{\langle C', I' \rangle}$$

Put simply, each stable external justification write is issued.

PROPOSITION F.9. Given a traversal step $G \vdash \langle C, I \rangle \rightarrow_t \langle C', I' \rangle$, we have that

$$G.\text{sjf}_{\langle C', I' \rangle}; [G.E \setminus G.E|_t] = G.\text{sjf}_{\langle C, I \rangle}; [G.E \setminus G.E|_t]$$

F.4 Simulation Relation for Certification Branch

Along with the simulation relation $\mathcal{I}(P, T, G, TC, S, X)$, the authors also define an auxiliary relation $\mathcal{I}^{\text{cert}}(P, T, G, TC, TC', S, X, Br_{\text{cert}})$, which we will call *certification simulation relation*. This relation holds during the construction of the certification branch [Moiseenko et al. 2020, §A]. It is also used at the end of the construction of the certification branch in order to restore the main simulation relation, i.e., to show that $\mathcal{I}(P, T, G, TC', S', X')$ holds. Below we list core invariants of $\mathcal{I}^{\text{cert}}(P, T, G, \langle C, I \rangle, \langle C', I' \rangle, S, X, Br_{\text{cert}})$. As in case of the main simulation relation we mention only a subset of properties which would be relevant to our modification of the proof.

- (1) $\mathcal{I}(P, T \setminus t, G, \langle C, I \rangle, S, X)$ holds.
- (2) $G \vdash \langle C, I \rangle \rightarrow_t \langle C', I' \rangle$, i.e., there is a traversal step, covering/issuing an event from thread t .
- (3) Certification branch is equal to a branch of the thread t in the configuration X up to a first uncovered event.
 - (a) $Br_{\text{cert}} \cap \llbracket C \rrbracket = X|_t \cap \llbracket C \rrbracket$.
In general, set of events of the certification branch is a subset of covered/issued events and their po-predecessors from the new traversal configuration.
 - (b) $\llbracket Br_{\text{cert}} \rrbracket \subseteq C' \cup \text{dom}(G.\text{po}^?; [I'])$.
- (4) Labels of determined events in Br_{cert} match labels of events in G .
 - (a) $\forall e \in Br_{\text{cert}} \cap \llbracket G.D_{\langle C', I' \rangle} \rrbracket. S.\text{val}(e) = G.\text{val}(\llbracket e \rrbracket)$
- (5) Reads in Br_{cert} are justified according to the stable justification relation.
 - (a) $\llbracket S.\text{jf}; [Br_{\text{cert}}] \rrbracket \subseteq G.\text{sjf}_{TC}$
- (6) For every issued event in Br_{cert} there exists $S.\text{ew}$ equivalent event in X .
 - (a) $Br \cap \llbracket I \rrbracket \subseteq \text{dom}(S.\text{ew}^?; [X])$
Symmetrically, every issued event in X within the processed part of the certification branch has an $S.\text{ew}$ equivalent event in Br_{cert} .
 - (b) $X \cap \llbracket I \cap \llbracket Br_{\text{cert}} \rrbracket \rrbracket \subseteq \text{dom}(S.\text{ew}^?; [Br_{\text{cert}}])$
- (7) The $S.\text{co}$ edges ending in Br_{cert} corresponds to $G.\text{co}$ edges.
 - (a) $\llbracket S.\text{co}; [Br_{\text{cert}}] \rrbracket \subseteq G.\text{co}$
The $S.\text{co}$ edges ending in $X|_t$ and not in the processed part of the certification branch correspond to $G.\text{co}$ edges.
 - (b) $\llbracket S.\text{co}; [X|_t \setminus \llbracket \llbracket Br_{\text{cert}} \rrbracket \rrbracket] \rrbracket \subseteq G.\text{co}$

F.5 Adapting the Proof

We now have all the ingredients to adapt the proof for Weakestmo2. Recall the **NO-BAIT-AND-SWITCH** axiom.

$$(\text{jf} \setminus (\text{jf}^?; \text{hb})); \text{po}; \text{ew} \subseteq \text{jf}; (\text{po} \cup \text{lbpatt})$$

Essentially, we need to show that if the simulation relation $\mathcal{I}(P, T, G, TC, S, X)$ holds, the event structure S satisfies **NO-BAIT-AND-SWITCH**, and the traversal makes a step $G \vdash TC \rightarrow TC'$, then the event structure can take multiple steps $S \rightarrow^* S'$ to construct a branch Br_{cert} , such that $\mathcal{I}(P, T, G, TC', S', X')$ holds and S' also satisfies **NO-BAIT-AND-SWITCH**.

Moiseenko et al. [2020, §4.3.3] describes how to construct a single branch Br_{cert} and S' such that $\mathcal{I}(P, T, G, TC', S', X')$ holds, where S' is equal to original event structure S extended with the new events from Br_{cert} . The branch Br_{cert} is built with the help of the receptiveness property using $G.\text{sjf}_{TC}$ to chose justification writes for read events. Unfortunately, if we take just this branch Br_{cert} , then S' might actually violate **NO-BAIT-AND-SWITCH**.

In order to fix the proof, we need to consider a *series* of certification branches $\{Br_1, \dots, Br_{n+1} = Br_{\text{cert}}\}$, where each branch is guided by a variant of the stable justification relation $G.\text{sjf}_{TC}^i$ (to be defined later). We can then show that each corresponding event structure S'_i satisfies **NO-BAIT-AND-SWITCH** and a weaker version of the simulation relation $I_{\text{weak}}^i(P, G, TC, TC', S'_i, X'_i)$. We then demonstrate that the final event structure S'_{n+1} in this series satisfies the main simulation relation $I(P, T, G, TC', S'_{n+1}, X'_{n+1})$.

The weak version of the simulation relation $I_{\text{weak}}^i(P, T, G, TC, TC', S, X)$ is parameterized by the natural number i and two consecutive traversal configurations $G \vdash TC \rightarrow TC'$. It has the following differences comparing to the regular simulation relation $I(P, T, G, TC, S, X)$.

- In **2b** instead of $G.D_{\langle C, I \rangle}$ we take $G.D_{\langle C', I' \rangle}^i$.
 - $\forall e \in X \cap \llbracket G.D_{\langle C', I' \rangle}^i \rrbracket. S.\text{val}(e) = G.\text{val}(\llbracket e \rrbracket)$
- In **5b** instead of $G.\text{sjf}_{TC}$ we take $G.\text{sjf}_{TC'}^i$.
 - $\llbracket S.\text{jf}; [X] \rrbracket \subseteq G.\text{sjf}_{TC'}^i$

We first briefly explain the intuition behind the construction of multiple branches. The newly issued write w might depend on several read events via the $G.\text{ppo}$ order. It means that the presence of w in certification branch and its label depend on the labels of these reads. The definition of $G.\text{sjf}_{TC'}$ ensures that event w would appear in certification branch and would have suitable label by choosing appropriate justification writes *for all* reads on which w depends via $G.\text{ppo}$. Changing justification for multiple reads at once violates **NO-BAIT-AND-SWITCH**, because the previous branch Br_{prev} of S and the certification branch Br_{cert} would have different set of incoming justification writes. It is possible, however, to replace justification writes for $G.\text{ppo}$ -preceding reads *incrementally* one-by-one, constructing a series of certification branches $\{Br_1, \dots, Br_{n+1}\}$. This way, any two consecutive branches in this sequence would only disagree on a single justification write located at the point of their immediate conflict, a situation which is allowed by axiom **NO-BAIT-AND-SWITCH** with the help of the load-buffering pattern relation lbpatt . Note that the newly constructed branches do not alter justification writes for reads $G.\text{ppo}$ -preceding previously issued writes I , because $\text{dom}(G.\text{ppo}; [I]) \subseteq G.D_{\langle C, I \rangle}$ and $G.\text{sjf}_{TC'}; [G.D_{TC}] = G.\text{sjf}_{TC}; [G.D_{TC}]$. In other words, all branches $\{Br_1, \dots, Br_{n+1}\}$ would contain writes equal to writes from I , which would allow us to certify all pending promises in each branch (remember that issued events correspond to promises).

Next we present a more formal version of the proof explained above. We would need several auxiliary definitions, propositions, and lemmas.

LEMMA F.10. *The set of determined reads can be characterized as follows.*

$$G.R \cap G.D_{\langle C, I \rangle} = G.R \cap C \cup \text{dom}(G.\text{ppo}; [I]) \cup \text{rng}([I]; G.\text{rfi}) \cup \text{rng}(G.\text{rfe}; [G.E^{\exists \text{acq}}])$$

PROOF. Follows directly from the definition of determined events and the following facts:

- $G.R \cap G.W = \emptyset$;
- $I \subseteq G.W$;
- $\text{dom}(G.\text{rfi}) \subseteq G.W$.

□

LEMMA F.11. *The viewfront relation can be equivalently defined as follows.*

$$G.\text{vf}_{\langle C, I \rangle} \triangleq [G.W]; (G.\text{rf}; [C])^2; G.\text{hb}^? \cup G.\text{rf}; [\text{dom}(G.\text{ppo}; [I])]; G.\text{po}^? \cup G.\text{rfe}; [G.E^{\exists \text{acq}}]; G.\text{po}^?$$

PROOF. Note that $\text{rng}(G.\text{rf}) \subseteq R$ and thus $G.\text{rf}; [G.D_{\langle C, I \rangle}] = G.\text{rf}; [R \cap G.D_{\langle C, I \rangle}]$. Using Lemma F.10 we can rewrite $G.\text{vf}_{\langle C, I \rangle}$ as follows.

$$G.\text{vf}_{\langle C, I \rangle} = [G.W]; (G.\text{rf}; [C])^2; G.\text{hb}^? \cup$$

$$\begin{aligned}
& \cup G.\mathbf{rf}; [C]; G.\mathbf{po}^? \cup \\
& \cup G.\mathbf{rf}; [\mathit{dom}(G.\mathbf{ppo}; [I])]; G.\mathbf{po}^? \cup \\
& \cup G.\mathbf{rf}; [\mathit{rng}([I]; G.\mathbf{rfi})]; G.\mathbf{po}^? \cup \\
& \cup G.\mathbf{rf}; [\mathit{rng}(G.\mathbf{rfe}; [G.E^{\exists\text{acq}}])]; G.\mathbf{po}^?
\end{aligned}$$

Applying the following (in)equalities we arrive to the conclusion.

$$G.\mathbf{rf}; [C]; G.\mathbf{po}^? \subseteq [G.W]; (G.\mathbf{rf}; [C])^?; G.\mathbf{hb}^?$$

$$G.\mathbf{rf}; [\mathit{rng}([I]; G.\mathbf{rfi})]; G.\mathbf{po}^? \subseteq [I]; G.\mathbf{rfi}; G.\mathbf{po}^? \subseteq [G.W]; G.\mathbf{po} \subseteq [G.W]; (G.\mathbf{rf}; [C])^?; G.\mathbf{hb}^?$$

$$G.\mathbf{rf}; [\mathit{rng}(G.\mathbf{rfe}; [G.E^{\exists\text{acq}}])]; G.\mathbf{po}^? = G.\mathbf{rfe}; [G.E^{\exists\text{acq}}]; G.\mathbf{po}^?$$

□

LEMMA F.12. *Let $G \vdash \langle C, I \rangle \rightarrow \langle C', I' \rangle$. Then we have:*

$$\begin{aligned}
G.D_{\langle C', I' \rangle} &= G.D_{\langle C, I \rangle} \cup \\
& \cup C' \setminus C \cup I' \setminus I \cup \\
& \cup \mathit{dom}(G.\mathbf{rfi}^?; G.\mathbf{ppo}; [I' \setminus I]) \cup \\
& \cup \mathit{rng}([I' \setminus I]; G.\mathbf{rfi})
\end{aligned}$$

PROOF. Follows immediately from the definition of determined events and the monotonicity of the traversal step. □

LEMMA F.13. *Let $G \vdash \langle C, I \rangle \rightarrow_t \langle C', I' \rangle$. Then we have:*

$$\begin{aligned}
G.\mathbf{vf}_{\langle C', I' \rangle} &= G.\mathbf{vf}_{\langle C, I \rangle} \cup \\
& \cup [G.W]; (G.\mathbf{rf}; [C' \setminus C])^?; G.\mathbf{hb}^? \cup \\
& \cup G.\mathbf{rf}; [\mathit{dom}(G.\mathbf{ppo}; [I' \setminus I])]; G.\mathbf{po}^? \cup
\end{aligned}$$

PROOF. Follows from the Lemma F.11 and the monotonicity of traversal step. □

LEMMA F.14. *For any valid traversal configuration the following is true.*

$$\mathit{dom}(G.\mathbf{rfi}; G.\mathbf{ppo}; [I]) \subseteq G.W \setminus \mathit{rng}(G.\mathbf{ppo}) \cup \mathit{rng}([\mathit{dom}(G.\mathbf{ppo}; [I])]; G.\mathbf{ppo})$$

PROOF. Consider $w_{\text{iss}} \in I$ and w, r s.t. $\langle w, r \rangle \in G.\mathbf{rfi}$ and $\langle r, w_{\text{iss}} \rangle \in G.\mathbf{ppo}$. Then consider the cases for w : it either has some $G.\mathbf{ppo}$ -preceding reads or not. In the later case $w \in G.W \setminus \mathit{rng}(G.\mathbf{ppo})$. In the former case consider some r_w , s.t. $\langle r_w, w \rangle \in G.\mathbf{ppo}$. But then

$$\langle r_w, w_{\text{iss}} \rangle \in G.\mathbf{ppo}; G.\mathbf{rfi}; G.\mathbf{ppo}; [I] \subseteq G.\mathbf{ppo}; [I]$$

(the last inequality follows from the definition of $G.\mathbf{ppo}$).

Therefore we have that $w \in \mathit{rng}([\mathit{dom}(G.\mathbf{ppo}; [I])]; G.\mathbf{ppo})$. □

Definition F.15. Let $G \vdash \langle C, I \rangle \rightarrow \langle C', I' \rangle$ and $\mathcal{E} \triangleq G.R \cap (C' \setminus C) \cup \mathit{dom}(G.\mathbf{ppo}; [I' \setminus I])$.

Since $\mathbf{ppo} \subseteq \mathbf{po}$ we have that all events from \mathcal{E} belong to the same thread, i.e., there exists $t \in \text{Tid}$ such that for all $e \in \mathcal{E}$ we have $G.\text{tid}(e) = t$. Therefore all the events in \mathcal{E} can be ordered according to their \mathbf{po} order. That is $\mathcal{E} = \{e_0, \dots, e_n\}$, such that $\langle e_i, e_{i+1} \rangle \in \mathbf{po}$ for all $0 \leq i \leq n$.

Then we define $G.D_{\langle C', I' \rangle}^i$, $G.\mathbf{vf}_{\langle C', I' \rangle}^i$ and $G.\mathbf{sjf}_{\langle C', I' \rangle}^i$ as follows.

- $G.D_{\langle C, I' \rangle}^0 = G.D_{\langle C, I \rangle}$
- $G.D_{\langle C, I' \rangle}^{i+1} = G.D_{\langle C, I' \rangle}^i \cup [e_i] \cup \text{dom}(G.\text{rf}i; G.\text{ppo}; [I' \setminus I]) \cap \text{rng}([e_i]; G.\text{ppo})$ for $0 < i < n$
- $G.D_{\langle C, I' \rangle}^{n+1} = G.D_{\langle C, I' \rangle}^n \cup [e_n] \cup \text{rng}([I' \setminus I]; G.\text{rf}i)$
- $G.\text{vf}_{\langle C, I' \rangle}^0 \triangleq G.\text{vf}_{\langle C, I \rangle}$
- $G.\text{vf}_{\langle C, I' \rangle}^{i+1} \triangleq G.\text{vf}_{\langle C, I' \rangle}^i \cup G.\text{rf}; [e_i]; G.\text{po}^?$ for $0 < i \leq n+1$
- $G.\text{sjf}_{TC}^i \triangleq ([G.W]; (G.\text{vf}_{TC}^i \cap =_{G.\text{loc}}); [G.R]) \setminus (G.\text{co}; G.\text{vf}_{TC}^i)$ for $0 \leq i \leq n+1$

LEMMA F.16. *We have that:*

- $G.D_{\langle C, I' \rangle}^{n+1} = G.D_{\langle C, I' \rangle}$
- $G.\text{vf}_{\langle C, I' \rangle}^{n+1} = G.\text{vf}_{\langle C, I' \rangle}$
- $G.\text{sjf}_{\langle C, I' \rangle}^{n+1} = G.\text{sjf}_{\langle C, I' \rangle}$

PROOF. Follows from the Def. F.15 and Lemmas F.12 to F.14. \square

LEMMA F.17. *For any $0 \leq i \leq n+1$ $G.\text{vf}_{TC}^i; G.\text{eco}$ is irreflexive.*

PROOF. By Prop. F.5 we know that $G.\text{vf}_{\langle C, I' \rangle}; G.\text{eco}$ is irreflexive. But since $G.\text{vf}_{\langle C, I' \rangle}^i \subseteq G.\text{vf}_{\langle C, I' \rangle}$ we arrive at the conclusion that $G.\text{vf}_{\langle C, I' \rangle}^i; G.\text{eco}$ should also be irreflexive. \square

LEMMA F.18. *For $0 \leq i \leq n$ let $\mathcal{R}_i \triangleq \text{rng}([e_i]; (\text{po} \cap =_{\text{loc}})^?); [R]$. Then the following equation holds.*

$$G.\text{sjf}_{TC}^{i+1}; [G.E \setminus \mathcal{R}_i] = G.\text{sjf}_{TC}^i; [G.E \setminus \mathcal{R}_i]$$

PROOF. By definition of $G.\text{sjf}_{TC}^{i+1}$, we have:

$$\begin{aligned} G.\text{sjf}_{TC}^{i+1}; [G.E \setminus \mathcal{R}_i] &= ([G.W]; (G.\text{vf}_{TC}^{i+1} \cap =_{G.\text{loc}}); [G.R]) \setminus (G.\text{co}; G.\text{vf}_{TC}^{i+1}); [G.E \setminus \mathcal{R}_i] = \\ &= ([G.W]; ((G.\text{vf}_{TC}^{i+1} \cap =_{G.\text{loc}}); [G.R \setminus \mathcal{R}_i])) \setminus (G.\text{co}; (G.\text{vf}_{TC}^{i+1} \cap =_{G.\text{loc}}); [G.R \setminus \mathcal{R}_i]) \end{aligned}$$

Therefore it is sufficient to show that $(G.\text{vf}_{TC}^{i+1} \cap =_{G.\text{loc}}); [G.R \setminus \mathcal{R}_i] = (G.\text{vf}_{TC}^i \cap =_{G.\text{loc}}); [G.R \setminus \mathcal{R}_i]$.

Thus by definition of $G.\text{vf}_{TC}^{i+1}$ we have:

$$G.\text{vf}_{TC}^{i+1}; [G.E \setminus \mathcal{R}_i] = (G.\text{vf}_{\langle C, I' \rangle}^i \cap =_{G.\text{loc}}); [G.E \setminus \mathcal{R}_i] \cup G.\text{rf}; [e_i]; (G.\text{po}^? \cap =_{G.\text{loc}}); [G.E \setminus \mathcal{R}_i]$$

Then it is easy to see that $G.\text{rf}; [e_i]; (G.\text{po}^? \cap =_{G.\text{loc}}); [G.E \setminus \mathcal{R}_i] = \emptyset$. \square

LEMMA F.19. *For any $0 \leq i \leq n$ we have the following.*

$$G.\text{sjf}_{TC}^{i+1}; [G.D_{TC}^i] = G.\text{sjf}_{TC}^i; [G.D_{TC}^i]$$

PROOF. By Lemma F.18 it is sufficient to show that $G.D_{TC}^i \subseteq G.E \setminus \mathcal{R}_i$ or equivalently $G.D_{TC}^i \cap \mathcal{R}_i = \emptyset$. Also $\mathcal{R}_i \subseteq G.R$ and thus the later is equivalent to $G.R \cap G.D_{TC}^i \cap \mathcal{R}_i = \emptyset$. Then by induction on i we can prove $G.R \cap G.D_{TC}^i \cap \mathcal{R}_i = G.R \cap G.D_{TC} \cap \mathcal{R}_i$. Finally, it remains to use Lemma F.10 and unfold the definition of \mathcal{R}_i . \square

LEMMA F.20. *Let $0 \leq i \leq n$. Then we have that:*

$$G.\text{sjf}_{\langle C, I' \rangle}^{i+1}; [e_i] = G.\text{rf}_{\langle C, I' \rangle}; [e_i]$$

Or in other words, stable justification for event e_i coincides with the reads-from relation.

PROOF. Since $G.\text{sjf}$ is functional and complete for read events it is sufficient to prove

$$G.\text{rf}_{\langle C, I \rangle}; [e_i] \subseteq G.\text{sjf}_{\langle C, I \rangle}^{i+1}; [e_i]$$

It is easy to see that $G.\text{rf}_{\langle C, I \rangle}; [e_i] \subseteq G.\text{vf}_{\langle C, I \rangle}^{i+1}; [e_i]$. Let $\langle w, e_i \rangle \in G.\text{rf}_{\langle C, I \rangle}; [e_i]$. Then it is left to show that $\langle w, e_i \rangle \notin G.\text{co}; G.\text{vf}_{\langle C, I \rangle}^{i+1}$. Suppose otherwise. Then there exists w' s.t. $\langle w, w' \rangle \in G.\text{co}$ and $\langle w', e_i \rangle \in G.\text{vf}_{\langle C, I \rangle}^{i+1}$. It cannot be the case that $\langle w', e_i \rangle \in G.\text{rf}; [e_i]; G.\text{po}^?$ so therefore $\langle w', e_i \rangle \in G.\text{vf}_{\langle C, I \rangle}^i$. But we also have $\langle e_i, w' \rangle \in G.\text{fr} \subseteq G.\text{eco}$ which contradicts Lemma F.17. \square

LEMMA F.21. For any $0 \leq i \leq n + 1$ we have the following.

$$G.\text{sjf}_{TC}^i; [G.D_{TC}^i] = G.\text{rf}; [G.D_{TC}^i]$$

PROOF. Proof goes straightforwardly by induction on i using Prop. F.7 and Lemmas F.19 and F.20. \square

LEMMA F.22. Let $0 \leq i \leq n$. Then we have that:

$$G.\text{sjf}_{\langle C, I \rangle}^{i+1}; [\text{rng}([e_i]; G.\text{po} \cap =_{G.1oc})] \subseteq G.\text{sjf}_{\langle C, I \rangle}^i \cup G.\text{sjf}_{\langle C, I \rangle}^{i+1}; [e_i]; G.\text{po}$$

PROOF. Consider $\langle w, r \rangle \in G.\text{sjf}_{\langle C, I \rangle}^{i+1} \subseteq G.\text{vf}_{\langle C, I \rangle}^{i+1} = G.\text{vf}_{\langle C, I \rangle}^i \cup G.\text{rf}; [e_i]; G.\text{po}^?$. If $\langle w, r \rangle \in G.\text{rf}; [e_i]; G.\text{po}^?$ then by Lemma F.20 $\langle w, r \rangle \in G.\text{sjf}_{\langle C, I \rangle}^{i+1}; [e_i]; G.\text{po}^?$ and we are done. Otherwise $\langle w, r \rangle \in G.\text{vf}_{\langle C, I \rangle}^i$. But then

$$\langle w, r \rangle \in G.\text{vf}_{\langle C, I \rangle}^i \setminus G.\text{co}; \text{vf}_{\langle C, I \rangle}^{i+1} \subseteq G.\text{vf}_{\langle C, I \rangle}^i \setminus G.\text{co}; \text{vf}_{\langle C, I \rangle}^i = G.\text{sjf}_{\langle C, I \rangle}^i$$

\square

LEMMA F.23. For any $0 \leq i \leq n$ the following is true.

$$G.\text{sjf}_{\langle C, I \rangle}^i; [e_i] \subseteq G.\text{sjf}_{\langle C, I \rangle}^i; G.\text{hb}$$

PROOF. Let $\langle w, e_i \rangle \in G.\text{sjf}_{\langle C, I \rangle}^i \subseteq G.\text{vf}_{\langle C, I \rangle}^i$. By Lemma F.11:

$$G.\text{vf}_{\langle C, I \rangle}^i = [G.W]; (G.\text{rf}; [C])^?; G.\text{hb}^? \cup G.\text{rf}; [\text{dom}(G.\text{ppo}; [e_0, \dots, e_{i-1}])]; G.\text{po}^? \cup G.\text{rfe}; [G.E^{\exists \text{acq}}]; G.\text{po}^?$$

It is left to notice the following facts

- $G.\text{rf}; [C] = G.\text{sjf}_{\langle C, I \rangle}^i; [C]$;
- $G.\text{rf}; [e_j] = G.\text{sjf}_{\langle C, I \rangle}^i; [e_j]$ for all $j < i$ (by Lemma F.21);
- $G.\text{rfe}; [G.E^{\exists \text{acq}}] = G.\text{sjfe}_{\langle C, I \rangle}^i; [G.E^{\exists \text{acq}}]$ (also by Lemma F.21).

\square

LEMMA F.24. Let $0 \leq i \leq n$. Then we have that:

$$\begin{aligned} & G.\text{sjf}_{TC}^{i+1} \setminus ((G.\text{sjf}_{TC}^{i+1})^?; G.\text{hb}); [\text{rng}([e_i]; G.\text{po} \cap =_{G.1oc})] = \\ & = G.\text{sjf}_{TC}^i \setminus ((G.\text{sjf}_{TC}^i)^?; G.\text{hb}); [\text{rng}([e_i]; G.\text{po} \cap =_{G.1oc})] \end{aligned}$$

PROOF. From left to right \subseteq it follows from the Lemma F.22. In the opposite direction, consider w, r s.t.

- $\langle w, r \rangle \in G.\text{sjf}_{TC}^i$;
- $\langle w, r \rangle \notin (G.\text{sjf}_{TC}^i)^?; G.\text{hb}$;
- $r \in \text{rng}([e_i]; G.\text{po} \cap =_{G.1oc})$.

Consider the write w' s.t. $\langle w', r \rangle \in G.\text{sjf}_{TC}^{i+1}$. By Lemma F.22 either

- $\langle w', r \rangle \in G.\text{sjf}_{TC}^i$ or

- $\langle w', r \rangle \in G.\mathbf{sjf}_{TC'}^{i+1}; [e_i]; G.\mathbf{po}$.

In the former case we conclude that $w' = w$ by functionality of $G.\mathbf{sjf}_{TC'}^i$.

In the latter case we notice that $\langle w', r \rangle \in G.\mathbf{sjf}_{TC'}^{i+1}; [e_i]; G.\mathbf{po}$ implies that the viewfront $G.\mathbf{vf}_{TC'}^{i+1}$ does not increase between e_i and r . But that means viewfront $G.\mathbf{vf}_{TC'}^i$ also cannot increase and thus it has to be $\langle w, r \rangle \in G.\mathbf{sjf}_{TC'}^i; [e_i]; G.\mathbf{po}$. From that using Lemma F.23 we can derive that $\langle w, r \rangle \in (G.\mathbf{sjf}_{TC'}^i)^?; G.\mathbf{hb}$ which contradicts our assumption. \square

LEMMA F.25. *For any $0 \leq i \leq n + 1$ we have the following.*

$$G.\mathbf{sjfi}_{TC'}^i = G.\mathbf{rfi}$$

PROOF. We prove this by induction on i . In the base case $i = 0$ we arrive at the conclusion by noticing that $G.\mathbf{sjf}_{TC'}^0 = G.\mathbf{sjf}_{TC}$ and by using Prop. F.6. For the induction step consider $\langle w, r \rangle \in G.\mathbf{sjf}_{TC'}^{i+1}$. Either $r \in \mathcal{R}_i$ or $r \notin \mathcal{R}_i$. In the latter case by Lemma F.18 we have that $G.\mathbf{sjf}_{TC'}^{i+1}; [r] = G.\mathbf{sjf}_{TC}^i; [r]$ and we arrive at the conclusion by using inductive assumption.

Otherwise consider the cases. If $r = e_i$ then by Lemma F.20 $G.\mathbf{sjfi}_{TC'}^{i+1}; [e_i] = G.\mathbf{rfi}; [e_i]$. If $r \in \mathit{rng}([e_i]; G.\mathbf{po} \cap =_{G.1oc})$ then by Lemma F.22 either

- $G.\mathbf{sjfi}_{TC'}^{i+1}; [r] = G.\mathbf{sjfi}_{TC'}^i; [r] = G.\mathbf{rfi}; [r]$ (the last equality due to inductive assumption), or
- $G.\mathbf{sjfi}_{TC'}^{i+1}; [r] \subseteq G.\mathbf{sjfi}_{TC'}^i; [e_i]; G.\mathbf{po}$, and given that $G.\mathbf{sjfi}_{TC'}^i; [e_i] = G.\mathbf{rfi}; [e_i]$ we conclude that $G.\mathbf{sjfi}_{TC'}^{i+1}; [r] = G.\mathbf{rfi}; [r]$

\square

LEMMA F.26. *For any $0 \leq i \leq n + 1$ we have the following.*

$$G.\mathbf{sjf}_{TC'}^i = [I]; G.\mathbf{sjfe}_{TC'}^i \cup G.\mathbf{sjfi}_{TC'}^i$$

PROOF. This lemma can be proven similarly as Lemma F.25 using Prop. F.8. \square

LEMMA F.27. *For any $0 \leq i \leq n$ we have the following.*

$$G.\mathbf{sjf}_{TC'}^{i+1}; [G.E \setminus G.E|_t] = G.\mathbf{sjf}_{TC'}^i; [G.E \setminus G.E|_t]$$

PROOF. This lemma can be proven similarly as Lemma F.25 using Prop. F.9. \square

Similarly as we defined a weaker version of the simulation relation $I_{\text{weak}}^i(P, T, G, TC, TC', S, X)$, we need to make some changes to the simulation relation for the certification branch and introduce $I_{\text{weak}}^{i+1 \text{ cert}}(P, T, G, \langle C, I \rangle, \langle C', I' \rangle, S, X, Br_{i+1})$.

- In **1** instead of $I(P, T \setminus t, G, \langle C, I \rangle, S, X)$ we take $I_{\text{weak}}^i(P, T \setminus t, G, \langle C, I \rangle, \langle C', I' \rangle, S, X)$
- In **3a** we say instead that the Br_{i+1} certification branch is equal to a branch of the thread t in the configuration X up to the event e_i .
 - $Br_{i+1} \cap \llbracket G.\mathbf{po}; [e_i] \rrbracket = X|_t \cap \llbracket G.\mathbf{po}; [e_i] \rrbracket$
- In **4a** instead of $G.D_{\langle C', I' \rangle}$ we take $G.D_{\langle C', I' \rangle}^{i+1}$.
 - $\forall e \in Br_{i+1} \cap \llbracket G.D_{\langle C', I' \rangle}^{i+1} \rrbracket. S.\mathbf{val}(e) = G.\mathbf{val}(\llbracket e \rrbracket)$
- In **5a** instead of $G.\mathbf{sjf}_{\langle C', I' \rangle}$ we take $G.\mathbf{sjf}_{\langle C', I' \rangle}^{i+1}$.
 - $\llbracket S.\mathbf{jf}; [Br_{i+1}] \rrbracket \subseteq G.\mathbf{sjf}_{\langle C', I' \rangle}^{i+1}$

Next, in order to modify the proof of F.3 we need to show the following.

- (1) Simulation relation $I(P, T, G, TC, S, X)$ implies weak simulation relation $I_{\text{weak}}^0(P, T, G, TC, TC', S, X)$.

- (2) Given that weak simulation relation holds, it is possible to start a construction of the next certification branch, i.e., from $\mathcal{I}_{\text{weak}}^i(P, T, G, TC, TC', S, X)$ we can prove that $\mathcal{I}_{\text{weak}}^{i+1 \text{ cert}}(P, T, G, TC, TC', S, X, Br_{i+1})$ holds, where $Br_{i+1} = X|_t \cap \llbracket G.\text{po} ; [e_i] \rrbracket$;
- (3) It is possible to add an event to the certification branch preserving the simulation relation and **NO-BAIT-AND-SWITCH**, i.e., if $\mathcal{I}_{\text{weak}}^i \text{ cert}(P, T, G, TC, TC', S, X, Br_i)$ holds and $S \xrightarrow{e} S'$, s.t. $\text{dom}(S'.\text{po} ; [e]) = Br_i$ then $\mathcal{I}_{\text{weak}}^i \text{ cert}(P, T, G, TC, TC', S', X, Br_i \cup \{e\})$ and S' satisfies **NO-BAIT-AND-SWITCH** assuming that S also satisfies it.
- (4) Given a complete certification branch it is possible to establish the weak simulation relation, i.e., from $\mathcal{I}_{\text{weak}}^i \text{ cert}(P, T, G, TC, TC', S, X, Br_i)$ where $\llbracket Br_i \rrbracket = C' \cup \text{dom}(G.\text{po}^? ; [I'])$ we can prove that $\mathcal{I}_{\text{weak}}^i(P, T, G, TC, TC', S, X')$ holds, where $X' \triangleq X \setminus X|_t \cup Br_i$.
- (5) Given the complete last certification branch it is possible to establish the regular simulation relation, i.e., from $\mathcal{I}_{\text{weak}}^{n+1 \text{ cert}}(P, T, G, TC, TC', S, X, Br_{n+1})$ where $\llbracket Br_{n+1} \rrbracket = C' \cup \text{dom}(G.\text{po}^? ; [I'])$ we can prove that $\mathcal{I}(P, T, G, TC', S, X')$ holds, where $X' \triangleq X \setminus X|_t \cup Br_{n+1}$.

In other words, in order to simulate a traversal step, we start the construction of the certification branches. We show that upon the completion of each certification branch we can restore weak version of the simulation relation and from that we can start the construction of the next branch. Finally, upon the completion of the last branch we restore the simulation relation.

The next few lemmas formalize this idea.

LEMMA F.28. $\mathcal{I}(P, G, TC, S, X)$ and $G \vdash TC \rightarrow_t TC'$ implies $\mathcal{I}_{\text{weak}}^0(P, G, TC, TC', S, X)$.

PROOF. Holds trivially since by definition $G.D_{\langle C, I' \rangle}^0 = G.D_{\langle C, I \rangle}$ and $G.\text{sjf}_{\langle C, I' \rangle}^0 = G.\text{sjf}_{\langle C, I \rangle}$. \square

LEMMA F.29. Suppose $\mathcal{I}_{\text{weak}}^i(P, T, G, TC, TC', S, X)$ holds, and $G \vdash TC \rightarrow_t TC'$. Then $\mathcal{I}_{\text{weak}}^{i+1 \text{ cert}}(P, T, G, TC, TC', S, X, Br_{i+1})$ holds, where $Br_{i+1} = X|_t \cap \llbracket \text{dom}(G.\text{po} ; [e_i]) \rrbracket$.

PROOF. We adapt the proof of a similar lemma for $\mathcal{I}^{\text{cert}}$ from [Moiseenko et al. 2020]. We need to show that each property of $\mathcal{I}_{\text{weak}}^i \text{ cert}$ holds.

- 1 and 2 hold trivially by our assumptions.
- 3a holds trivially by our choice of Br_{i+1} :

$$Br_{i+1} \cap \llbracket \text{dom}(G.\text{po} ; [e_i]) \rrbracket = X|_t \cap \llbracket \text{dom}(G.\text{po} ; [e_i]) \rrbracket$$

- 3b holds since $\text{dom}(\text{po} ; [e_i]) \subseteq C' \cup \text{dom}(G.\text{po}^? ; [I'])$.
- To see why 4a holds first consider the following equation.

$$Br_{i+1} \cap \llbracket G.D_{TC'}^{i+1} \rrbracket = X|_t \cap \llbracket \text{dom}(G.\text{po} ; [e_i]) \cap G.D_{TC'}^{i+1} \rrbracket = X|_t \cap \llbracket G.D_{TC}^i \rrbracket$$

and for $e \in X \cap \llbracket G.D_{TC}^i \rrbracket$ we know that $S.\text{lab}(e) = G.\text{lab}(\llbracket e \rrbracket)$ by the property 4a of certification simulation relation.

- 5a holds because

$$\llbracket S.\text{jf} ; [Br_{i+1}] \rrbracket = \llbracket S.\text{jf} ; [X \cap \llbracket \text{dom}(G.\text{po} ; [e_i]) \rrbracket] \rrbracket \subseteq G.\text{sjf}_{TC}^i ; [\text{dom}(G.\text{po} ; [e_i])] \subseteq G.\text{sjf}_{TC}^{i+1}$$

The last inequality follows from lemma F.18 and the fact that

$$\text{dom}(G.\text{po} ; [e_i]) \subseteq \mathcal{R}_i = G.E \setminus \text{rng}([e_i]; (G.\text{po} \cap =_{G.\text{loc}})^? ; [G.R])$$

- 6a holds by our choice of Br_{i+1} :

$$Br_{i+1} \cap \llbracket I \rrbracket = X|_t \cap \llbracket \text{dom}(G.\text{po} ; [e_i]) \rrbracket \cap \llbracket I \rrbracket \subseteq \text{dom}(S.\text{ew}^? ; [X])$$

- **6b** also holds by our choice of Br_{i+1} :

$$\begin{aligned} X \cap \llbracket I \cap \llbracket Br_{i+1} \rrbracket \rrbracket X \cap \llbracket I \cap \llbracket X|_t \cap \llbracket \text{dom}(G.\text{po} ; [e_i]) \rrbracket \rrbracket \rrbracket &\subseteq \\ &\subseteq X \cap \llbracket I \cap \llbracket X|_t \rrbracket \cap \text{dom}(G.\text{po} ; [e_i]) \rrbracket \subseteq \\ &\subseteq X|_t \cap \llbracket I \cap \text{dom}(G.\text{po} ; [e_i]) \rrbracket \subseteq Br_{i+1} \cap \llbracket I \rrbracket \subseteq \text{dom}(S.\text{ew}^? ; [Br_{i+1}]) \end{aligned}$$

- **7a** and **7b** follow from property **7b** of the simulation relation and the fact that $Br_{i+1} \subseteq X|_t$. □

LEMMA F.30. For $1 \leq i \leq n+1$ if $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, T, G, TC, TC', S, X, Br_i)$ holds and $S \xrightarrow{e} S'$, s.t.

- $\text{dom}(S'.\text{po} ; [e]) = Br_i$;
- $\llbracket e \rrbracket \in C' \cup \text{dom}(G.\text{po}^? ; [I'])$;
- $\llbracket S'.\text{jf} ; [e] \rrbracket \subseteq G.\text{sjf}_{TC'}^i$

then $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, T, G, TC, TC', S', X, Br_i \cup \{e\})$ holds.

PROOF. We need to show that each property of $\mathcal{I}_{\text{weak}}^{i \text{ cert}}$ holds.

- To show **1** we refer to the proof of the similar statement in [Moiseenko et al. 2020]. Since the only difference between $\mathcal{I}(P, T \setminus t, G, \langle C, I \rangle, S, X)$ and $\mathcal{I}_{\text{weak}}^i(P, T \setminus t, G, \langle C, I \rangle, \langle C', I' \rangle, S, X)$ is that the later uses $G.D_{\langle C', I' \rangle}^i$ and $G.\text{sjf}_{\langle C', I' \rangle}^i$ instead of $G.D_{\langle C, I \rangle}$ and $G.\text{sjf}_{\langle C, I \rangle}$ the proof proceeds in a similar vein. The only significant point we want to emphasize is the proof of **5c**. To show that $\text{dom}(S'.\text{jfe}) \subseteq \text{dom}(S.\text{ew}^* ; [X \cap \llbracket I \rrbracket])$ the original proof utilized the fact that $\llbracket S'.\text{jfe} ; [e] \rrbracket \subseteq G.\text{sjf}_{\langle C, I \rangle}$ and property **F.8**. We instead note that $\llbracket S'.\text{jfe} ; [e] \rrbracket \subseteq G.\text{sjf}_{\langle C', I' \rangle}^i$ and use similar Lemma **F.26**.
- **2** follow immediately from $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, G, TC, TC', S, X, Br_i)$.
- **3a** holds because

$$(Br_i \cup \{e\}) \cap \llbracket \text{dom}(G.\text{po} ; [e_{i-1}]) \rrbracket = Br_i \cap \llbracket \text{dom}(G.\text{po} ; [e_{i-1}]) \rrbracket = X|_t \cap \llbracket \text{dom}(G.\text{po} ; [e_{i-1}]) \rrbracket$$

3b holds trivially by the precondition $\llbracket e \rrbracket \in C' \cup \text{dom}(G.\text{po}^? ; [I'])$.

- **4** and **5** hold by the choice of e , as we chose labels of events and their justification according to $G.\text{sjf}_{TC'}^i$ and by lemma **F.19** and proposition **F.7** we have $G.\text{sjf}_{TC'}^i ; [G.D_{TC'}^i] = G.\text{rf} ; [G.D_{TC'}^i]$.
- To show **6** we note that $I \subseteq G.D_{\langle C, I \rangle} \subseteq G.D_{\langle C', I' \rangle}^i$ and thus if $\llbracket e \rrbracket \in I$ then $S'.\text{val}(e) = G.\text{val}(\llbracket e \rrbracket)$. Therefore, similarly as it is done in [Moiseenko et al. 2020], we can pick an equivalent write in $X|_t$, s.t. the labels of two writes would be equal. As we have shown above using the properties of $G.D_{\langle C', I' \rangle}^i$ both writes would have label equal to $G.\text{val}(\llbracket e \rrbracket)$.
- Since we update $S'.\text{co}$ order following similar rules as in [Moiseenko et al. 2020], the proof of **7** is not affected by our modifications and thus this property remains valid. □

LEMMA F.31. For $1 \leq i \leq n$ if $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, T, G, TC, TC', S, X, Br_i)$ holds and $\llbracket Br_i \rrbracket = C' \cup \text{dom}(G.\text{po}^? ; [I'])$ then $\mathcal{I}_{\text{weak}}^{i-1}(P, T, G, TC, TC', S, X')$ holds, where $X' \triangleq X \setminus X|_t \cup Br_i$.

PROOF. We adapt the proof of a similar lemma for $\mathcal{I}^{\text{cert}}$ from [Moiseenko et al. 2020]. We need to show that each property of $\mathcal{I}_{\text{weak}}^{i-1}$ holds.

- First we note that **2a**, **3**, **4**, **5a**, **5c**, **6**, **7a**, and **8** follow immediately from the same properties of $\mathcal{I}_{\text{weak}}^{i-1}(P, T, G, TC, TC', S, X')$ which is implied from $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, T, G, TC, TC', S, X, Br_i)$ by **1**.
- **1** holds trivially due to precondition $\llbracket Br_i \rrbracket = C' \cup \text{dom}(G.\text{po}^? ; [I'])$ and our choice of X' .

- **2b** holds because we have

$$X' \cap \llbracket G.D_{\langle C, I' \rangle}^i \rrbracket = (X \setminus X|_t) \cap G.D_{\langle C, I' \rangle}^{i-1} \cup Br_i \cap \llbracket G.D_{\langle C, I' \rangle}^i \rrbracket$$

For left subset we derive the required property from **2b** of $\mathcal{I}_{\text{weak}}^{i-1}(P, T, G, TC, TC', S, X')$, for the right subset we derive it from **4a** of $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, T, G, TC, TC', S, X, Br_i)$.

- To see why **5b** holds first note that.

$$\llbracket S.\mathbf{jf}; [X'] \rrbracket = \llbracket S.\mathbf{jf}; [X \setminus X|_t] \rrbracket \cup \llbracket S.\mathbf{jf}; [Br_i] \rrbracket$$

Then $\llbracket S.\mathbf{jf}; [Br_i] \rrbracket \subseteq G.\mathbf{sjf}_{TC}^i$ is implied by **5a** of $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, T, G, TC, TC', S, X, Br_i)$. Next using **5b** of $\mathcal{I}_{\text{weak}}^{i-1}(P, T, G, TC, TC', S, X')$ and Lemma **F.27** we prove the following.

$$\llbracket S.\mathbf{jf}; [X \setminus X|_t] \rrbracket \subseteq G.\mathbf{sjf}^{i-1}; [G.E \setminus G.E|_t] \subseteq G.\mathbf{sjf}^i$$

- **7b** can be proven by the similar reasoning using property **7a** of $\mathcal{I}_{\text{weak}}^{i \text{ cert}}(P, T, G, TC, TC', S, X, Br_i)$ and **7b** of $\mathcal{I}_{\text{weak}}^{i-1}(P, T, G, TC, TC', S, X')$.

□

LEMMA F.32. *If $\mathcal{I}_{\text{weak}}^{n+1 \text{ cert}}(P, T, G, TC, TC', S, X, Br_{n+1})$ holds and $\llbracket Br_{n+1} \rrbracket = C' \cup \text{dom}(G.\text{po}^?; [I'])$ then $\mathcal{I}(P, T, G, TC', S, X')$ holds, where $X' \triangleq X \setminus X|_t \cup Br_{n+1}$.*

PROOF. We note that due to lemma **F.16** $\mathcal{I}_{\text{weak}}^{n+1 \text{ cert}}(P, T, G, TC, TC', S, X, Br_{n+1})$ is equivalent to $\mathcal{I}^{\text{cert}}(P, T, G, TC, TC', S, X, Br_{n+1})$ and thus we can reuse the proof of the similar lemma from [Moiseenko et al. 2020]. □

We are now moving to the main lemma of our development, that is we are going to show that the axiom **NO-BAIT-AND-SWITCH** is preserved during the construction of the event structure.

LEMMA F.33. *Suppose S satisfies **NO-BAIT-AND-SWITCH**. Assume $S \xrightarrow{e} S'$ and $S'.\text{ew} = S.\text{ew}$. Then S' also satisfies **NO-BAIT-AND-SWITCH**.*

PROOF. Note that $S.E$ is prefix-closed w.r.t. po and \mathbf{jf} , i.e., $S.\text{po}'$; $[S.E] = S.\text{po}$ and $S.\mathbf{jf}'$; $[S.E] = S.\mathbf{jf}$. Therefore $S'.\mathbf{jf}$; $S'.\text{po}$; $S'.\text{ew} = S.\mathbf{jf}$; $S.\text{po}$; $S.\text{ew}$. From that we can derive the following.

$$\begin{aligned} & (S'.\mathbf{jf} \setminus (S'.\mathbf{jf}'^?; S'.\text{hb})); S'.\text{po}; S'.\text{ew} = \\ & = (S.\mathbf{jf} \setminus (S.\mathbf{jf}'^?; S.\text{hb})); S.\text{po}; S.\text{ew} \subseteq S.\mathbf{jf}; (S.\text{po} \cup S.\text{lbp}at) \subseteq S'.\mathbf{jf}; (S'.\text{po} \cup S'.\text{lbp}at) \end{aligned}$$

The last inequality follows from the fact that po , \mathbf{jf} and $\text{lbp}at$ are monotone w.r.t. addition of new events $S \xrightarrow{e} S'$. □

LEMMA F.34. *For $0 \leq i \leq n$ if $\mathcal{I}_{\text{weak}}^{i+1 \text{ cert}}(P, T, G, TC, TC', S, X, Br_{i+1})$ holds, $S \xrightarrow{e} S'$, s.t.*

- $\text{dom}(S'.\text{po}; [e]) = Br_{i+1}$;
- $\llbracket e \rrbracket \in C' \cup \text{dom}(G.\text{po}^?; [I'])$;
- $\llbracket S'.\mathbf{jf}; [e] \rrbracket \subseteq G.\mathbf{sjf}_{TC}^i$

*and if S satisfies **NO-BAIT-AND-SWITCH** then S' also satisfies **NO-BAIT-AND-SWITCH**.*

PROOF. Let us note that it is sufficient to consider only the case when $e \in W$. Otherwise by construction $S'.\text{ew} = S.\text{ew}$ and using Lemma **F.33** we can immediately conclude that S' satisfies **NO-BAIT-AND-SWITCH**. So let $e \in W$.

By the construction of [Moiseenko et al. 2020, §4.3.2] the ew relation is updated during the step $S \xrightarrow{e} S'$ when there exists some issued write event w belonging to the configuration X such that its images of w and e in the graph G are equal, i.e., $w \in X \cap \llbracket I \rrbracket$ and $\llbracket e \rrbracket = \llbracket e \rrbracket$. In this case e is

attached to the $S.\text{ew}^*$ equivalence class of w by adding a new ew -edge from w to e (see also property 6 of simulation relation for certification branch to get the idea about this construction).

Therefore, in order to prove that S' satisfies **NO-BAIT-AND-SWITCH** we need to consider justification writes for reads in the two conflicting branches $Br_i = X|_t$ and Br_{i+1} . By Lemma F.30 we know that $J_{\text{weak}}^{i+1 \text{ cert}}(P, T, G, TC, TC', S', X, Br_{i+1} \cup \{e\})$ holds and thus by property 5b of weak simulation relation and property 5a of simulation relation for certification branch we know that

- $\llbracket S.\text{jf} ; [Br_i] \rrbracket \subseteq G.\text{sjf}_{\langle C, I' \rangle}^i$, and
- $\llbracket S.\text{jf} ; [Br_{i+1}] \rrbracket \subseteq G.\text{sjf}_{\langle C, I' \rangle}^{i+1}$.

By Lemma F.18 we also know that:

$$G.\text{sjf}_{TC'}^{i+1} ; [G.E \setminus \mathcal{R}_i] = G.\text{sjf}_{TC'}^i ; [G.E \setminus \mathcal{R}_i]$$

where $\mathcal{R}_i \triangleq \text{rng}([e_i \cup \text{rng}([I' \setminus I] ; G.\text{rfi})] ; (\text{po} \cap =_{\text{loc}})^? ; [\mathbf{R}])$. Put simply, for all read events except those in \mathcal{R}_i two conflicting branches Br_i and Br_{i+1} pick the same justification writes and thus for these reads the conditions of **NO-BAIT-AND-SWITCH** are met.

Therefore it is sufficient to consider justification for reads in \mathcal{R}_i .

Let us first consider events po -succeeding event e_i . Let r be such event: $\langle e_i, r \rangle \in (G.\text{po} \cap =_{G.\text{loc}}) ; [G.R]$. Using Lemma F.24 we derive that

$$G.\text{sjf}_{\langle C, I' \rangle}^{i+1} \setminus ((G.\text{sjf}_{\langle C, I' \rangle}^{i+1})^? ; G.\text{hb}) ; [r] = G.\text{sjf}_{\langle C, I' \rangle}^i \setminus ((G.\text{sjf}_{\langle C, I' \rangle}^i)^? ; G.\text{hb}) ; [r]$$

But that means justification writes for reads in Br_{i+1} and Br_i corresponding to r should either be taken from $S'.\text{jf} ; S'.\text{hb}$ prefix or be equal in two branches, therefore the conditions of **NO-BAIT-AND-SWITCH** are met for r .

It is left to consider the event e_i itself. Let r_i and r_{i+1} be corresponding events in Br_i and Br_{i+1} respectively.

First, we consider justification write for r_i . By Lemma F.23 it should be the case that

$$\llbracket S'.\text{jf} ; [r_i] \rrbracket = \llbracket S.\text{jf} ; [r_i] \rrbracket \subseteq G.\text{sjf}_{\langle C, I' \rangle}^i ; [e_i] \subseteq [G.W] ; (G.\text{rf} ; [C])^? ; G.\text{hb}$$

But then it should also be the case that:

$$S.\text{jf} ; [r_i] \subseteq ([X] ; S.\text{rf} ; [X \cap \llbracket C \rrbracket])^? ; S.\text{hb} = ([X] ; (S.\text{ew}^* ; S.\text{jf}) \setminus S.\text{cf} ; [X \cap \llbracket C \rrbracket])^? ; S.\text{hb}$$

From that we can conclude that

$$S'.\text{jf} ; [r_i] = S.\text{jf} ; [r_i] \subseteq S.\text{jf} ; S.\text{hb} \subseteq S'.\text{jf} ; S'.\text{hb}$$

Therefore r_i is justified locally and thus **NO-BAIT-AND-SWITCH** is satisfied for r_i .

Finally, we need to show that r_{i+1} and w form load buffering pattern, i.e., $\langle r_{i+1}, w \rangle \in \text{lbpatt}$ (recall that w is an equivalent write we picked for e , i.e., $\langle w, e \rangle \in S'.\text{ew}$). If we do this then clearly conditions of **NO-BAIT-AND-SWITCH** are met for r_{i+1} .

Recall the definition of the load buffering pattern

$$\text{lbpatt} \triangleq \text{cf}_{\text{imm}} ; [\text{rng}(\text{jf} \cap (\text{jf}^? ; \text{hb}))] ; \text{po}$$

By property 3a of simulation relation for certification branch we know that Br_i and Br_{i+1} are equal up to event e_i and thus r_i and r_{i+1} should be in immediate conflict: $\langle r_{i+1}, r_i \rangle \in S'.\text{cf}_{\text{imm}}$. Moreover, as we have shown r_i is locally justified, i.e., $r_i \in \text{rng}(S'.\text{jf} \cap (S'.\text{jf}^? ; S'.\text{hb}))$. Therefore $\langle r_{i+1}, w \rangle \in S'.\text{cf}_{\text{imm}} ; [\text{rng}(\text{jf} \cap (\text{jf}^? ; \text{hb}))] ; S'.\text{po}$.

□

G SOUNDNESS OF PROGRAM TRANSFORMATIONS FOR WEAKESTM02

In this section we present the proofs of the soundness of program transformations in Weakestmo2. Our development is based on the proofs for the original Weakestmo model [Chakraborty et al. 2019].

G.1 Reordering of Independent Instructions

We start with the soundness of reorderings. We consider four kind of reorderings: store/load, store/store, load/load, and load/store.

$$\begin{array}{ll}
 x := r_1 ; r_2 := y & \rightsquigarrow & r_2 := y ; x := r_1 & \text{store/load} \\
 x := r_1 ; y := r_2 & \rightsquigarrow & y := r_2 ; x := r_1 & \text{store/store} \\
 r_1 := x ; r_2 := y & \rightsquigarrow & r_2 := y ; r_1 := x & \text{load/load} \\
 r_1 := x ; y := r_2 & \rightsquigarrow & y := r_2 ; r_1 := x & \text{load/store}
 \end{array}$$

We assume that all accesses are relaxed. Handling other kinds of accesses would significantly complicate the proofs with a lot of technical details that are orthogonal to the main idea behind our reasoning.

THEOREM G.1. *Reordering transformation $P_{\text{src}} \xrightarrow{\text{reord}(t,a,b)} P_{\text{tgt}}$ where instructions a and b form a store/load, store/store, load/load, or load/store pair of independent accesses is sound.*

G.1.1 Recap of the Proof Structure. Let us first remind the proof structure for the original version of Weakestmo [Chakraborty et al. 2019, §F].

Given a Weakestmo2 consistent execution graph G_{tgt} of program P_{tgt} one need to construct a Weakestmo2 consistent execution graph G_{src} of P_{src} such that $\mathcal{B}(G_{\text{src}}) = \mathcal{B}(G_{\text{tgt}})$. To do so Chakraborty et al. [2019, §F] consider target event structure S_{tgt} , s.t. $S_{\text{tgt}} \triangleright G_{\text{tgt}}$, and build source event structure S_{src} and execution graph G_{src} , s.t. $S_{\text{src}} \triangleright G_{\text{tgt}}$, following the construction of S_{tgt} by the operational semantics $S_{\text{init}}(P_{\text{tgt}}) \rightarrow^* S_{\text{tgt}}$.

Similarly to the proof of the correctness of compilation mappings §F.1, the proof of the soundness of reorderings relies on the *simulation* argument. Chakraborty et al. [2019, §F] define the simulation relation $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S_{\text{src}}, S_{\text{tgt}}, X_{\text{src}}, \mu)$ between the source and target programs $P_{\text{src}}, P_{\text{tgt}}$; target execution graphs G_{tgt} ; source and target Weakestmo consistent event structures $S_{\text{src}}, S_{\text{tgt}}$; justified configuration of the source event structure X_{src} , and function $\mu : S_{\text{src}}.E \rightarrow S_{\text{tgt}}.E$ that maps events of the source event structure to the events of the target event structure⁸.

The three simulation lemmas are presented below ensure that: (i) initial source and target event structures are bound by the simulation relation; (ii) each construction step of the target event structure can be simulated by the construction step of the source event structure; (iii) at the end of the simulation process, selected configuration of source event structures forms an execution graph with the same behavior as the given target execution graph.

LEMMA G.2 (SIMULATION START). *Let P_{src} and P_{tgt} be source and target programs, s.t. $P_{\text{src}} \xrightarrow{\text{reord}(t,a,b)} P_{\text{tgt}}$, and let G_{tgt} be Weakestmo consistent execution graph, corresponding to the target program P_{tgt} . Then $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S_{\text{init}}(P_{\text{src}}), S_{\text{init}}(P_{\text{tgt}}), \{e_0^s\}, \mu_0)$ holds where*

- $S_{\text{init}}(P_{\text{src}})$ is initial source event structure containing only the initialization event $\{e_0^s\}$;

⁸In order to simplify the proof, we have slightly modified the simulation relation presented in Chakraborty et al. [2019, §F]. In particular, instead of the relation $\mathcal{M} \subseteq S_{\text{src}}.E \times S_{\text{tgt}}.E$ that connects events of the source and target event structure, we use function $\mu : S_{\text{src}}.E \rightarrow S_{\text{tgt}}.E$. We then use notations $\llbracket \cdot \rrbracket$ and $\llbracket \cdot \rrbracket$, similarly as in §F.2, to lift function μ to subsets of events and component relations (such as po, jf, co) of the two event structures.

- $S_{\text{init}}(P_{\text{tgt}})$ is initial target event structure containing only the initialization event $\{e_0^t\}$;
- μ_0 is the initial event mapping, s.t. $\mu_0(e_0^s) = e_0^t$ and $\mu_0(e) = \perp$ for all $e \neq e_0^s$.

LEMMA G.3 (WEAK SIMULATION STEP). *If $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S_{\text{src}}, S_{\text{tgt}}, X_{\text{src}}, \mu)$ and $S_{\text{tgt}} \rightarrow^* S'_{\text{tgt}}$ hold then there exist $S'_{\text{src}}, X'_{\text{src}}$, and μ' such that $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S'_{\text{src}}, S'_{\text{tgt}}, X'_{\text{src}}, \mu')$ and $S_{\text{src}} \rightarrow^* S'_{\text{src}}$ hold.*

LEMMA G.4 (SIMULATION END). *If $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S_{\text{src}}, S_{\text{tgt}}, X_{\text{src}}, \mu)$ and $S_{\text{tgt}} \triangleright G_{\text{tgt}}$ hold, then the execution graph associated with X_{src} , denoted as G_{src} , has the same behavior as G_{tgt} , i.e., $\mathcal{B}(G_{\text{src}}) = \mathcal{B}(G_{\text{tgt}})$.*

Similarly to the case of Theorem 3.14, the main objective of our modification of the proof is to show that the simulation of the construction step (i.e., Lemma G.3) preserves the new axiom **NO-BAIT-AND-SWITCH**.

G.1.2 *Simulation Relation.* We next describe the simulation relation \mathcal{I} .

The simulation relation \mathcal{I} distinguishes three subsets of events of the target event structure S_{tgt} .

Giving that $P_{\text{src}} \xrightarrow{\text{reord}(t,a,b)} P_{\text{tgt}}$ let

- $A \subseteq S_{\text{tgt}}.E$ be a subset of events obtained as a result of executing instruction a;
- $B \subseteq S_{\text{tgt}}.E$ be a subset of events obtained as a result of executing instruction b;
- $C \triangleq S_{\text{tgt}}.E \setminus (A \cup B)$ contain the remaining events.

The simulation \mathcal{I} relation establishes the following properties

- (1) For each non-reordered event in S_{tgt} there exists a corresponding event in S_{src} :
 - (a) $S_{\text{tgt}}.E \cap C \subseteq \llbracket S_{\text{src}}.E \rrbracket$
- (2) Labels of events in S_{src} match labels of events in S_{tgt} :
 - (a) $\forall e \in S_{\text{src}}.E. S_{\text{src}}.\text{lab}(e) = S_{\text{tgt}}.\text{lab}(\llbracket e \rrbracket)$
- (3) Program order in S_{tgt} between events from A and C corresponds to program order in S_{src} :
 - (a) $[A \cup C]; S_{\text{tgt}}.\text{po}; [A \cup C] = [A \cup C]; \llbracket S_{\text{src}}.\text{po} \rrbracket; [A \cup C]$
Similarly, program order in S_{tgt} between events from B and C corresponds to program order in S_{src} :
 - (a) $[B \cup C]; S_{\text{tgt}}.\text{po}; [B \cup C] = [B \cup C]; \llbracket S_{\text{src}}.\text{po} \rrbracket; [B \cup C]$
 For immediate program order in S_{src} between events from A and B the following is true:
 - (a) $[A]; \llbracket S_{\text{src}}.\text{po}_{\text{imm}} \rrbracket; [B] \subseteq S_{\text{tgt}}.\text{ew}^?; S_{\text{tgt}}.\text{po}_{\text{imm}}^{-1}; S_{\text{tgt}}.\text{cf}_{\text{imm}}^?$
Immediate po predecessor of an event from B in S_{tgt} corresponds to immediate po predecessor of an event from A in S_{tgt} :
 - (a) $[C]; S_{\text{tgt}}.\text{po}_{\text{imm}}; [B] = [C]; \llbracket S_{\text{src}}.\text{po}_{\text{imm}} \rrbracket; \llbracket [A] \rrbracket; S_{\text{src}}.\text{po}_{\text{imm}} \rrbracket; [B]$
- (4) Identity relation in S_{tgt} corresponds to identity or conflict relation in S_{src} :
 - (a) $\llbracket \text{id} \rrbracket \subseteq S_{\text{src}}.\text{cf}^?$
- (5) Justified-from relation in S_{tgt} from an event belonging to A or C to an event belonging to C corresponds to justified-from relation in S_{src} :
 - (a) $[A \cup C]; S_{\text{tgt}}.\text{jf}; [C] = [A \cup C]; \llbracket S_{\text{src}}.\text{jf} \rrbracket; [C]$
Justified-from relation in S_{tgt} from an event belonging to B or C to an event belonging to B or C corresponds to justified-from relation in S_{src} :
 - (a) $[B \cup C]; S_{\text{tgt}}.\text{jf}; [B \cup C] = [B \cup C]; \llbracket S_{\text{src}}.\text{jf} \rrbracket; [B \cup C]$
For each justification of an event from A in S_{tgt} there exists corresponding justification in S_{src} :
 - (a) $S_{\text{tgt}}.\text{jf}; [A] \subseteq \llbracket S_{\text{src}}.\text{jf} \rrbracket$
Each justification of an event from A in S_{src} is either local (i.e., from $\text{jf}^?; \text{hb}$ preceding write) or the same as in S_{tgt} :

- (d) $S_{\text{src}}.\text{jf} ; [\llbracket A \rrbracket] \subseteq S_{\text{src}}.\text{jf}^? ; S_{\text{src}}.\text{hb} \cup \llbracket S_{\text{tgt}}.\text{jf} \rrbracket$
- (6) Each pair of equivalent writes from in S_{src} between events from B corresponds to the same event or a pair of equivalent writes in S_{tgt} :
- (a) $[B] ; \llbracket S_{\text{src}}.\text{ew} \rrbracket ; [B] \subseteq S_{\text{tgt}}.\text{ew}^?$
Equivalent-writes relation in S_{tgt} between write events from C corresponds to Equivalent-writes relation in S_{src} :
- (b) $[C] ; S_{\text{tgt}}.\text{ew} ; [C] = [C] ; \llbracket S_{\text{src}}.\text{ew} \rrbracket ; [C]$
- (7) Coherence relation in S_{tgt} corresponds to the coherence relation in S_{src} :
- (a) $S_{\text{tgt}}.\text{co} = \llbracket S_{\text{src}}.\text{co} \rrbracket$
- (8) Extracted execution graph G_{src} , corresponding to configuration X_{src} , is equal to execution graph G_{tgt} modulo reordering of events $a \in A \cap G_{\text{tgt}}.\text{E}$ and $b \in B \cap G_{\text{tgt}}.\text{E}$
- (a) $G_{\text{src}}.\text{E} = G_{\text{tgt}}.\text{E} \cap S_{\text{tgt}}.\text{E}$
- (b) $\forall e \in G_{\text{src}}.\text{E}. G_{\text{src}}.\text{lab}(e) = G_{\text{tgt}}.\text{lab}(\llbracket e \rrbracket)$
- (c) $\llbracket G_{\text{src}}.\text{po} \rrbracket = G_{\text{tgt}}.\text{po}|_{S_{\text{tgt}}.\text{E}} \setminus \{(b, a)\} \cup \{(a, b)\}$
- (d) $\llbracket G_{\text{src}}.\text{rf} \rrbracket = G_{\text{tgt}}.\text{rf}|_{S_{\text{tgt}}.\text{E}}$
- (e) $\llbracket G_{\text{src}}.\text{co} \rrbracket = G_{\text{tgt}}.\text{co}|_{S_{\text{tgt}}.\text{E}}$

G.1.3 Adapting the proof. We need to show that simulation step respects **NO-BAIT-AND-SWITCH** axiom.

$$(\text{jf} \setminus (\text{jf}^? ; \text{hb})) ; \text{po} ; \text{ew} \subseteq \text{jf} ; (\text{po} \cup \text{lbpatt})$$

That is, we need to show that if the simulation relation $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S_{\text{src}}, S_{\text{tgt}}, X_{\text{src}}, \mu)$ holds, the source event structure S_{src} satisfies **NO-BAIT-AND-SWITCH**, and the construction of the target event structure makes a step $S_{\text{tgt}} \xrightarrow{e_t} S'_{\text{tgt}}$, s.t. S_{tgt} and S'_{tgt} both satisfy **NO-BAIT-AND-SWITCH**, then there exist $S'_{\text{src}}, X'_{\text{src}}$, and μ' such that $S_{\text{src}} \rightarrow^* S'_{\text{src}}$, $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S'_{\text{src}}, S'_{\text{tgt}}, X'_{\text{src}}, \mu')$ holds, and S'_{src} also satisfies **NO-BAIT-AND-SWITCH**.

In order to prove this we consider each kind of reordering separately.

Store/Load: Consider the cases.



Fig. 8. A fragment of the event structure construction that justifies store/load reordering. Semi-transparent events may or may not be presented in the event structures when event b_2^t is added.

- $e_t \in B$

We first consider the most challenging case. Let us consult Fig. 8. The target event structure adds event $e_t = b_t : R(y, \beta)$. To simulate this step, source event structure should be first updated with an event $a_s : W(x, \alpha)$ (if there exists a_t' , s.t. $\langle a_t', a_t \rangle \in S'_{\text{tgt}}.\text{ew}$, then S_{src} should already contain a_s , s.t. $\mu(a_s) = a_t'$), and then with an event $b_s : R(y, \beta)$:

$$S_{\text{src}} \xrightarrow{a_s^?} S''_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

Note that it is indeed possible to perform the first step, because instructions a and b are independent (thus the written value α of the event $a_t : W(x, \alpha)$ cannot depend on the value of the previous event $b_t : R(y, \beta)$), and because it is always possible to add a write event to

the event structure (because write events do not need to be justified via **jf** relation). Finally, note that if event b_t depends on some promises in S'_{tgt} , then in S'_{src} it should depend on the same or smaller set of promises. Indeed, suppose that in S'_{tgt} event b_t depends via **jo** path on event a'_t , s.t. $\langle a'_t, a_t \rangle \in S'_{tgt}.ew$. In S'_{src} this dependency will be lost, since b_s will depend on a_s . All other promises will remain the same. Similarly, write event b_s and all subsequent write events will depend on the same set of external writes (since b_s and b_t have the same justification writes). Therefore, if b_s depends on some promises it will be possible to certify them by the same certification writes as in target event structure S'_{tgt} without violation of **NO-BAIT-AND-SWITCH**.

- $e_t \in A$
Consult Fig. 8 again, and let $e_t = a_t$. In this case we can conclude that S_{src} should already have a matching event a_s . Thus $S'_{src} = S_{src}$.
- $e_t \in C$
In this case construction of the source event structure can simulate the step by adding event e_s with the same label $S_{src} \xrightarrow{e_s} S'_{src}$. Since S'_{tgt} satisfies **NO-BAIT-AND-SWITCH** then so S'_{src} , because if e belongs to some certification branch in S'_{tgt} then in S'_{src} this branch should depend on the smaller or equal set of promises (as was explained above) and on the same external writes.

Store/Store: Consider the cases.



Fig. 9. A fragment of event structure construction that justifies store/store reordering.

- $e_t \in B$
Let us consult Fig. 9. The target event structure adds event $e_t = b_t : W(y, \beta)$. To simulate this step, events $a_s : W(x, \alpha)$ and $b_s : W(y, \beta)$ are added to the source event structure:

$$S_{src} \xrightarrow{a_s} S''_{src} \xrightarrow{b_s} S'_{src}$$

The events a_s and b_s in S'_{src} depend on the same set of promises and external writes as b_t in S'_{tgt} . Therefore, each of these promise will be certified by the same certification writes as in target event structure S'_{tgt} without violation of **NO-BAIT-AND-SWITCH**.

- $e_t \in A$
Consult Fig. 8 again, and let $e_t = a_t$. In this case we can conclude that S_{src} should already have a matching event a_s . Thus $S'_{src} = S_{src}$.
- $e_t \in C$
In this case construction of the source event structure can simulate the step by adding event e_s with the same label $S_{src} \xrightarrow{e_s} S'_{src}$. Since S'_{tgt} satisfies **NO-BAIT-AND-SWITCH** then so S'_{src} .

Load/Load: Consider the cases.

- $e_t \in B$
Let $e_t = b_t : R(y, \beta)$ and let $a_t :: R(x, \alpha)$ be the subsequent event. To simulate this step, events

$a_s : R(x, \alpha)$ and $b_s : R(y, \beta)$ are added to the source event structure:

$$S_{\text{src}} \xrightarrow{a_s} S''_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

We need to consider the cases with respect to promises which b_t and a_t depend on.

- Suppose that a_t, b_t and their po predecessor up to the read event c_t depend on the same set of promises, (including the case of empty set, i.e., when these events do not depend on any promises at all). This case is depicted on Fig. 10. In this case we can conclude that all promises will be certified by the same certification writes as in target event structure S'_{tgt} without violation of **NO-BAIT-AND-SWITCH**. Indeed, each of these promises will depend on the same external writes, because a_s and b_s have the same justification writes as a_t and b_t ; moreover all other reads in the certification branch remain unchanged.

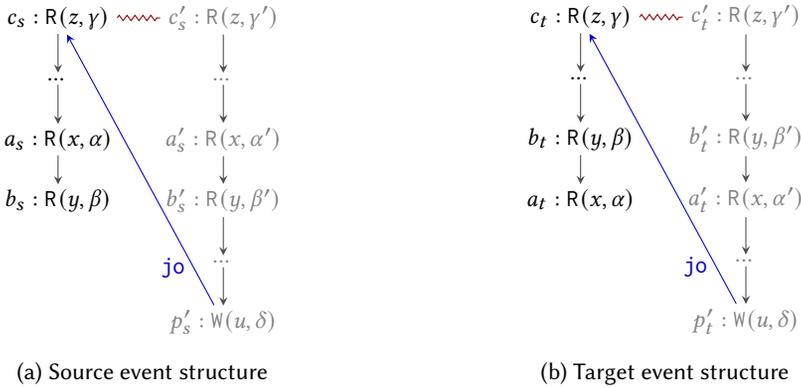


Fig. 10. A fragment of event structure construction that justifies load/load reordering. Case when both reads b_t, a_t and their po predecessors up to c_t depend on the same set of promises.

- Suppose that a_t, b_t depend on the same set of promises, but their po predecessor do not depend on these promises. Also suppose that p'_t is po maximal promise. This case is depicted on Fig. 11. First, let us prove that in this case $\alpha = \alpha'$ and thus a_t and a'_t should have the same label. Indeed, since a_t and a'_t belong to neighboring branches due to **NO-BAIT-AND-SWITCH** it has to be that either these reads are local, i.e., they read from **jf**? ; **hb** prior **co**-maximal write, or they are external and thus should read from the same write. Therefore, it has to be the case that S_{src} already contains a_s , s.t. $\mu(a_s) = a'_t$, and thus $S''_{\text{src}} = S_{\text{src}}$. Note that b_s in S'_{src} depends on the same set of promises as b_t in S'_{tgt} . All these promises in S'_{src} will depend on the same or lesser set of external writes compared to S'_{tgt} (because the justification write of read event a_s will not be considered as external write with respect to certification branch containing b_s). Therefore these promises will be certified by the same certification writes as in the target event structure S'_{tgt} without violation of **NO-BAIT-AND-SWITCH**.
- Suppose that a_t depends on some set of promises, but b_t does not depend on these promises. Also suppose that p'_t is po maximal promise and a'_t is read event in immediate conflict with a_t . This case is depicted on Fig. 12. Event a_s in S'_{src} depends on the same set of promises as a_t in S'_{tgt} . These promises will be certified by the same certification writes as in the target event structure. However, if the read event b_s is justified from external write, then in the source event structure these promises will also depend on this external write. Nevertheless, note that event b'_s in the branch that have issued the promises should read from the same write as b_s , since $\mu(b'_s) = \mu(b_s) = b_t$. Therefore, in the source event structure the promise

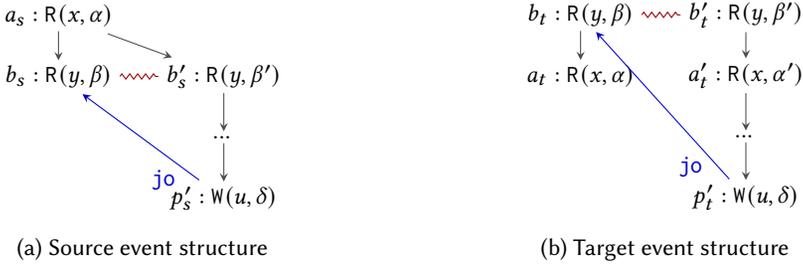


Fig. 11. A fragment of event structure construction that justifies load/load reordering. Case when both reads b_t, a_t depend on the same set of promises but their po predecessors do not depend on these promises.

p'_s and its certification write will depend on the same set of external writes, thus satisfying **NO-BAIT-AND-SWITCH**.

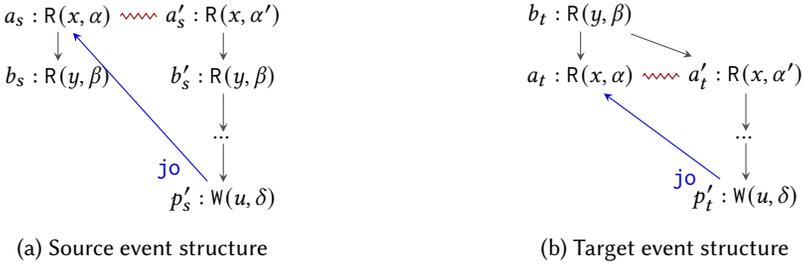


Fig. 12. A fragment of event structure construction that justifies load/load reordering. Case when read a_t depends on some set of promises, but b_t does not depend on these promises.

- $e_t \in A$

Consult Figures 10 to 12 again, and let $e_t = a_t$. In this case we can conclude that S_{src} should already have a matching event a_s . Thus $S'_{\text{src}} = S_{\text{src}}$.

- $e_t \in C$

In this case construction of the source event structure can simulate the step by adding event e_s with the same label $S_{\text{src}} \xrightarrow{e_s} S'_{\text{src}}$. Since S'_{tgt} satisfies **NO-BAIT-AND-SWITCH** then so S'_{src} .

Load/Store: Consider the cases.

- $e_t \in B$

Let $e_t = b_t : W(y, \beta)$ and let $a_t : R(x, \alpha)$ be the subsequent event. Then we need to consider the following cases.

- Suppose that a_t does not depend on b_t via **jo** path. This case is depicted on Fig. 13. Then it is possible to update source event structure by adding events $a_s : R(x, \alpha)$ and $b_s : W(y, \beta)$:

$$S_{\text{src}} \xrightarrow{a_s} S''_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

Note that event a_s is justified by the same write in S'_{src} as a_t in S'_{tgt} . Therefore, if a_s (and, possibly, b_s) depend on some set of promises, it will be possible to certify them by the same writes as in S'_{tgt} . Moreover, the set of external writes of these promises remain the same. Therefore, **NO-BAIT-AND-SWITCH** holds.



Fig. 13. A fragment of event structure construction that justifies load/store reordering. Case when a_t does not depend on b_t via **jo** path.

- Suppose that a_t depend on b_t via **jo** path and that b_t, a_t do not depend on promises. This case is depicted on Fig. 14. In this case the source event structure is updated by addition of events $a'_s : R(x, \gamma)$ and $b'_s : W(y, \beta)$:

$$S_{\text{src}} \xrightarrow{a'_s} S''_{\text{src}} \xrightarrow{b'_s} S'_{\text{src}}$$

As a justification write for a'_s the construction chooses event $c_s : W(x, \gamma)$, that is **co** maximal non-conflicting write from **jf**? ; **hb** prefix of a'_s . Since a'_s and b'_s do not depend on promises, the condition **NO-BAIT-AND-SWITCH** is satisfied trivially.

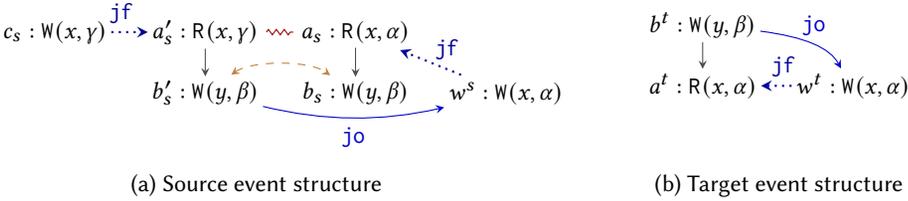


Fig. 14. A fragment of event structure construction that justifies load/store reordering. Case when b_t and a_t do not depend on promises.

- Suppose that a_t depend on b_t via **jo** path, and that a_t depends on some set of promises, but b_t does not depend on these promises. This case is depicted on Fig. 15. In this case, due to **NO-BAIT-AND-SWITCH**, event a_t should be in immediate conflict with event a'_t which is justified locally, i.e., from **jf**? ; **hb** preceding write $c_t : W(x, \gamma)$. We can conclude that S_{src} should contain events $a'_s : R(x, \gamma)$ and $b'_s : W(y, \beta)$, s.t. $\mu(a'_s) = a'_t$ and $\mu(b'_s) = b_t$. Therefore $S'_{\text{src}} = S_{\text{src}}$.

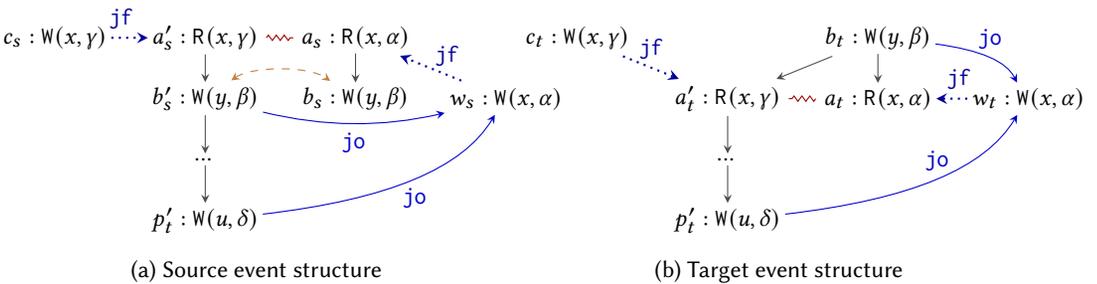


Fig. 15. A fragment of event structure construction that justifies load/store reordering. Case when a_t depends on promises but b_t does not.

- Suppose that a_t depend on b_t via **jo** path, and also that a_t, b_t depend on the same set of promises. This case is depicted on Fig. 16. Consider write w_t that justifies a_t . Note that w_t is external write, i.e., $\langle w_t, a_t \rangle \notin \text{jf}^? ; \text{hb}$, or otherwise we would get $\text{po} \cup \text{jf}$ cycle $b_t \xrightarrow{\text{jo}} w_t \xrightarrow{\text{jo}} b_t$. Also note that it cannot be the case that $\langle w_t, p_t'' \rangle \in \text{jf}^? ; \text{hb}$, as required by **NO-BAIT-AND-SWITCH**, because it also would imply $\text{po} \cup \text{jf}$ cycle $w_t \xrightarrow{\text{jo}} p_t'' \xrightarrow{\text{jo}} w_t$. Therefore, there should exist another certification branch neighboring to the branch containing a_t . In other words, there should exist event a_t' in immediate conflict with a_t , and all subsequent events up to the event p_t' certifying promise p_t' . Thus, the source event structure is updated by addition of a branch $a_s' \rightarrow b_s' \rightarrow \dots \rightarrow p_s'$:

$$S_{\text{src}} \rightarrow^* S'_{\text{src}}$$

Since this new branch in S'_{src} mimics the corresponding branch in S'_{tgt} , it also satisfies **NO-BAIT-AND-SWITCH**.

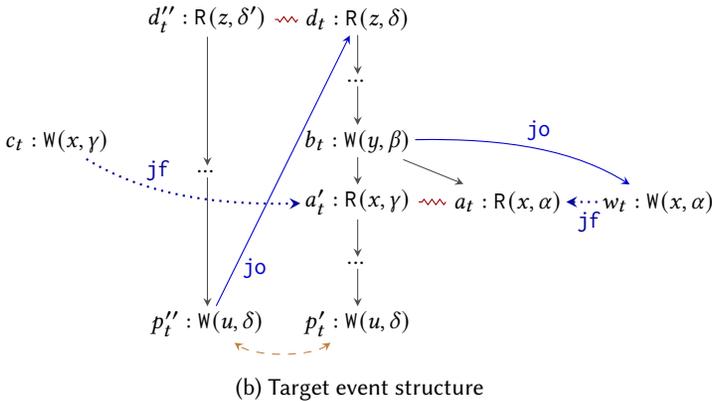
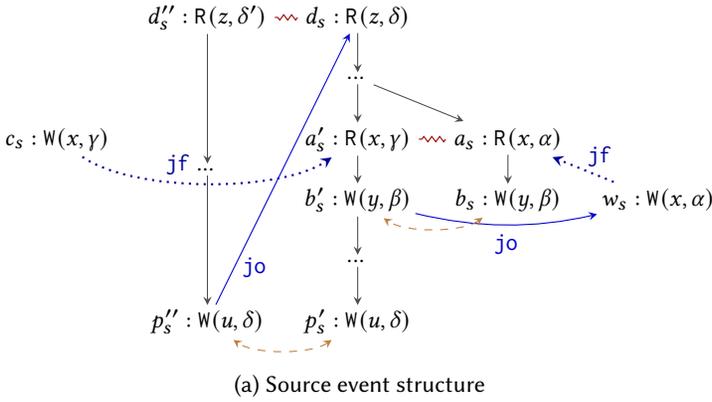


Fig. 16. A fragment of event structure construction that justifies load/store reordering. Case when a_t and b_t depends on the same set of promises.

- $e_t \in A$

Let $e_t = a_t : R(x, \alpha)$. We need to consider the same cases as above.

- Consult Fig. 13 again. It can be seen that S_{src} should already contain suitable event a_s . Therefore $S'_{\text{src}} = S_{\text{src}}$.
- Consult Fig. 14 again. In this case the source event structure is updated by addition of events $a_s : R(x, \alpha)$ and $b_s : W(y, \alpha)$:

$$S_{\text{src}} \xrightarrow{a_s} S''_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

These events form a new certification branch. This branch satisfies **NO-BAIT-AND-SWITCH**, because it only differs from the neighboring branch at the point of immediate conflict, and moreover a'_s is justified locally, i.e., from **jf**? ; **hb** prior write.

- Consult Fig. 15 again. In this case the source event structure is updated by addition of events $a_s : R(x, \alpha)$ and $b_s : W(y, \alpha)$:

$$S_{\text{src}} \xrightarrow{a_s} S''_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

These events form a new certification branch mimicking the branch starting at a_t in S'_{tgt} . All promises on which this branch depends will be certified by the same certification writes as in S'_{tgt} . Additionally, in S'_{src} this branch also depends on the promise b'_s . This promise is certified by b_s . Therefore, the new branch satisfies **NO-BAIT-AND-SWITCH**.

- Consult Fig. 16 again. In this case the source event structure is updated by addition of events $a_s : R(x, \alpha)$ and $b_s : W(y, \alpha)$:

$$S_{\text{src}} \xrightarrow{a_s} S''_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

Similarly to the previous case, these events form a new certification branch mimicking the branch starting at a_t in S'_{tgt} . By the same reasoning, the new branch satisfies **NO-BAIT-AND-SWITCH**.

- $e_t \in C$
In this case construction of the source event structure can simulate the step by adding event e_s with the same label $S_{\text{src}} \xrightarrow{e_s} S'_{\text{src}}$. Since S'_{tgt} satisfies **NO-BAIT-AND-SWITCH** then so S'_{src} .

G.2 Elimination of Redundant Accesses

Next we prove soundness of eliminations.

Suppose that thread t of a program P contains two consecutive instructions $a ; b$ such that these instructions access the same location and one of the instructions is redundant. For example, a and b can be a pair of stores to the same location, then the first store is redundant because its value is overwritten by the second one. Then elimination transformation $\text{elim}(t, a, b)$ removes or replaces redundant instruction, and leaves other threads unchanged.

We consider three kind of eliminations: store/load, store/store, load/load⁹:

$$\begin{aligned} x := r_1 ; r_2 := x &\rightsquigarrow x := r_1 ; r_2 := r_1 && \text{store/load} \\ x := r_1 ; x := r_2 &\rightsquigarrow x := r_2 && \text{store/store} \\ r_1 := x ; r_2 := x &\rightsquigarrow r_1 := x ; r_2 := r_1 && \text{load/load} \end{aligned}$$

Again, we assume that all accesses are relaxed.

THEOREM G.5. *Elimination transformation $P_{\text{src}} \xrightarrow{\text{elim}(t,a,b)} P_{\text{tgt}}$ where instructions a and b form a store/load, store/store, or load/load pair of redundant accesses is sound.*

⁹Note that load/store elimination is unsound for atomic accesses both in the original Weakestmo and in C11 models [Chakraborty et al. 2019; Vafeiadis et al. 2015] and thus we do not consider it

G.2.1 Recap of the Proof Structure. Proof of the soundness of eliminations mimics proof of the soundness of transformations [Chakraborty et al. 2019, §G]. Similarly as in §G.1.1 the construction of the source event structure simulates constructions of the target event structure. The main objective of our modification of the proof is to show that the simulation of the construction step preserves the new axiom **NO-BAIT-AND-SWITCH**.

G.2.2 Simulation Relation. We next describe the simulation relation \mathcal{I} used in the proof of the soundness of elimination transformations. This simulation relation has a similar form compared to the one presented in §G.2.2.

That is, the simulation relation $\mathcal{I}(P_{\text{src}}, P_{\text{tgt}}, G_{\text{tgt}}, S_{\text{src}}, S_{\text{tgt}}, X_{\text{src}}, \mu)$ establish a connection between the source and target programs $P_{\text{src}}, P_{\text{tgt}}$; target execution graphs G_{tgt} ; source and target Weakestmo consistent event structures $S_{\text{src}}, S_{\text{tgt}}$; justified configuration of the source event structure X_{src} , and function $\mu : S_{\text{src}}.E \rightarrow S_{\text{tgt}}.E$ that maps events of the source event structure to the events of the target event structure.

The simulation relation \mathcal{I} distinguishes three subsets of events of the target event structure S_{tgt} . Giving that $P_{\text{src}} \xrightarrow{\text{elim}(t,a,b)} P_{\text{tgt}}$ let

- $A \subseteq S_{\text{tgt}}.E$ be a subset of events obtained as a result of executing instruction a ;
- $B \subseteq S_{\text{tgt}}.E$ be a subset of events obtained as a result of executing instruction b ;
- $C \triangleq S_{\text{tgt}}.E \setminus (A \cup B)$ contain the remaining events.

Note that depending on the particular kind of elimination either A or B is empty, because the target program contains only one of the redundant instructions:

- in case of store/store elimination A is empty;
- in case of store/load and load/load eliminations B is empty.

The simulation \mathcal{I} relation establishes the following properties

- (1) For each event in S_{tgt} there exists a corresponding event in in S_{src} :
 - (a) $S_{\text{tgt}}.E \subseteq \llbracket S_{\text{src}}.E \rrbracket$
Function μ is injective on non-eliminated events:
 - (b) $\forall e_1, e_2 \in S_{\text{src}}.E. \mu(e_1) = \mu(e_2) \in C \implies e_1 = e_2$
- (2) Label of at least one of the events forming a redundant access pair in S_{src} matches the label of its image in S_{tgt} :
 - (a) $\forall \langle e_a, e_b \rangle \in \llbracket [A \cup B] \rrbracket ; S_{\text{src}}.\text{po}_{\text{imm}} ; \llbracket [A \cup B] \rrbracket$ either
 - $A \neq \emptyset$ and $S_{\text{src}}.\text{lab}(e_a) = S_{\text{tgt}}.\text{lab}(\llbracket e_a \rrbracket)$, or
 - $B \neq \emptyset$ and $S_{\text{src}}.\text{lab}(e_b) = S_{\text{tgt}}.\text{lab}(\llbracket e_b \rrbracket)$
 Labels of events from C match in both event structures:
 - (b) $\forall e \in \llbracket C \rrbracket. S_{\text{src}}.\text{lab}(e) = S_{\text{tgt}}.\text{lab}(\llbracket e \rrbracket)$
- (3) Program order in S_{tgt} between events from A and C corresponds to program order in S_{src} :
 - (a) $[A \cup C] ; S_{\text{tgt}}.\text{po} ; [A \cup C] = [A \cup C] ; \llbracket S_{\text{src}}.\text{po} \rrbracket ; [A \cup C]$
Similarly, program order in S_{tgt} between events from B and C corresponds to program order in S_{src} :
 - (b) $[B \cup C] ; S_{\text{tgt}}.\text{po} ; [B \cup C] = [B \cup C] ; \llbracket S_{\text{src}}.\text{po} \rrbracket ; [B \cup C]$
Pair of redundant accesses in S_{src} maps to the same event in S_{tgt} :
 - (c) $[A \cup B] ; \llbracket S_{\text{src}}.\text{po}_{\text{imm}} \rrbracket ; [A \cup B] \subseteq \text{id}$
- (4) Justified-from edge in S_{src} to an event belonging to A or B corresponds either to the same event in S_{tgt} or to a justified-from edge in S_{tgt} :
 - (a) $\llbracket S_{\text{src}}.\text{jf} \rrbracket ; [A \cup B] \subseteq S_{\text{tgt}}.\text{jf}^?$
Justified-from edge in S_{tgt} to an event belonging to A or B corresponds either to a justified-from edge in S_{src} :

Store/Load: Consider the cases.



Fig. 18. A fragment of event structure construction that justifies store/store elimination.

- $e_t \in A$

Let us consult Fig. 18. The target event structure adds event $e_t = b_t : W(x, \beta)$. To simulate this step, the construction adds events $a_s : W(x, \alpha)$ and $b_s : R(x, \beta)$ to the source event structure:

$$S_{\text{src}} \xrightarrow{a_s} S'_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

Since a_s is a write event it is not a subject of the constraint **NO-BAIT-AND-SWITCH**. In other words, if b_s depends on some set of promises in S'_{src} , they will be certified similarly as promises of b_t in S'_{tgt} , because the set of external writes that justify the certification branch does not change.

- $e_t \in C$

In this case construction of the source event structure can simulate the step by adding event e_s with the same label $S_{\text{src}} \xrightarrow{e_s} S'_{\text{src}}$. Since S'_{tgt} satisfies **NO-BAIT-AND-SWITCH** then so S'_{src} .

Load/Load: Consider the cases.



Fig. 19. A fragment of event structure construction that justifies load/load elimination.

- $e_t \in A$

Let us consult Fig. 19. The target event structure adds event $e_t = a_t : R(x, \alpha)$. Let $c_t : W(x, \alpha)$ be a write event chosen to justify a_t . To simulate this step, the construction adds events $a_s : R(x, \alpha)$ and $b_s : R(x, \alpha)$ to the source event structure:

$$S_{\text{src}} \xrightarrow{a_s} S'_{\text{src}} \xrightarrow{b_s} S'_{\text{src}}$$

As a justification write for both events the construction takes write event c_s , s.t. $\mu(c_s) = c_t$. It implies that b_s is justified locally, i.e., from $\text{jf} ; \text{po} \subseteq \text{jf}^? ; \text{hb}$ preceding write. Therefore, it is not a subject of the constraint **NO-BAIT-AND-SWITCH**. In other words, if a_s depends on some set of promises in S'_{src} , they will be certified similarly as promises of a_t in S'_{tgt} , because the set of external writes that justify the certification branch does not change.

- $e_t \in C$

In this case construction of the source event structure can simulate the step by adding event e_s with the same label $S_{\text{src}} \xrightarrow{e_s} S'_{\text{src}}$. Since S'_{tgt} satisfies **NO-BAIT-AND-SWITCH** then so S'_{src} .