Model Checking for Weakly Consistent Libraries

Michalis Kokologiannakis
MPI-SWS
michalis@mpi-sws.org

Azalea Raad
MPI-SWS
azalea@mpi-sws.org

Viktor Vafeiadis
MPI-SWS
viktor@mpi-sws.org

Abstract

We present GENMC, a model checking algorithm for concurrent programs that is parametric in the choice of memory model and can be used for verifying clients of concurrent libraries. Subject to a few basic conditions about the memory model, our algorithm is sound, complete and optimal, in that it explores each consistent execution of the program according to the model exactly once, and does not explore inconsistent executions or embark on futile exploration paths. We implement GENMC as a tool for verifying C programs. Despite the generality of the algorithm, its performance is comparable to the state-of-art specialized model checkers for specific memory models, and in certain cases exponentially faster thanks to its coarse partitioning of executions.

CCS Concepts: • Theory of computation → Verification by model checking; • Software and its engineering → Software testing and debugging.

Keywords: Model checking, weak memory models

1 Introduction

Suppose that we have a concurrent program, e.g.,

$$\begin{array}{l} x := 1 \\ y := 1 \end{array} \parallel \begin{array}{l} a := y \\ b := x \\ \text{assert}(a \leq b) \end{array} \quad (\text{MP})$$

with x and y initialized with 0, and we want to check whether its assertions are always satisfied. An effective way of doing so is using *stateless model checking* (SMC) [18, 19, 34], which enumerates all executions of the program and checks each execution individually. SMC has two major challenges.

The first challenge is associated with the *memory model* under which the program is executed, as it determines the program outcomes. For example, in the MP program above, the assertion ($a \leq b$) holds under SC [28] and TSO [36], but not under PSO [40], or RC11 with ‘relaxed’ accesses [27], because the latter two models allow for writes to distinct locations to be reordered. Thus, under PSO and RC11, $y := 1$ may execute before $x := 1$, thereby violating the assertion.

The second challenge is that any non-trivial concurrent program has a large number of executions that need to be explored (typically, *exponential* in the size of the program). To tackle this, *partial order reduction* techniques [1, 12, 16, 20, 42] have been developed, and try to partition the executions into equivalence classes and explore exactly one execution per equivalence class.

However, while there exist efficient techniques that target *specific* memory models [1–4, 12, 15, 16, 20–22, 26, 35, 39, 42], a *generic* technique that combats both these challenges is yet to be developed.

The goal of this paper is to develop such a model checking algorithm that is parametric in the choice of the memory model. Our algorithm, GENMC (Generic Model Checker), can be used not only for traditional memory models supporting reads, writes, and read-modify-write (RMW) instructions, but also for models incorporating high-level libraries, such as mutual exclusion locks, as primitive operations.

Our contributions can be summarized as follows:

- Through a series of examples, we present an intuitive account of our algorithm for verifying concurrent programs, using execution graphs and axiomatic semantics for *any* memory model (§2), so long as it satisfies four basic assumptions: *porf*-acyclicity, extensibility, prefix-closedness and well-blocking (§3).
- Our approach distinguishes executions based solely on the *program-order* and *reads-from* relations (§2.5), which can lead to exponentially fewer explorations compared to approaches that maintain a total *coherence* order between conflicting writes (§6.3).
- We demonstrate how our technique can verify programs under memory models that incorporate high-level *libraries*, such as mutual exclusion locks (§2.7).
- We describe our algorithm in detail (§4), and *prove* that it is (a) sound: produces no false positives; (b) complete: explores all possible program behaviours; and (c) optimal: explores each behaviour exactly once.
- We implement GENMC into a tool for verifying C programs (§5), and demonstrate that it has comparable or better performance than the state-of-the-art specialized tools for specific memory models (§6).

2 Overview

In the literature of axiomatic memory models [6, 27], the traces of shared memory accesses generated by a program are represented as a set of *execution graphs*, where each graph G comprises: (i) a set of events (graph nodes); and (ii) a few relations on events (graph edges). The two kinds of edges present in all memory models are the *program order* (po) and the *reads-from* relation (*rf*), which relates each read event r in G to a write event w in G , from which r obtains its value. The *semantics* of a program P is then given by the set of executions that satisfy a certain *consistency* predicate.

Figure 1. Execution graphs of $W+RW+W$.

For example, *sequential consistency* (SC) [28] requires the existence of a total order on all events extending the program order such that each read reads from the most recent prior write to same location in that total order. Equivalently, SC can be defined in terms of a *modification order*, mo , also known as the *coherence order*. An execution is SC-consistent if there exists mo such that for every location x , mo totally orders all writes to x and $po \cup rf \cup mo \cup (rf^{-1}; mo)$ is acyclic, where rf^{-1} denotes the inverse of rf , and ‘;’ denotes *relational composition*: $(r, w) \in rf^{-1}; mo \Leftrightarrow \exists w'. (w', r) \in rf \wedge (w', w) \in mo$.

Other models are weaker and deem more executions consistent. TSO [36] allows loads to execute before po -earlier stores, while PSO [40] further allows stores of a thread to execute out of order. RC11 [27] supports various access modes ranging from SC to ones even weaker than what PSO offers.

Let us now consider a simple example program:¹

$$x := 1 \parallel \begin{array}{l} a := x; \\ \text{if } a = 0 \text{ then } y := 1 \end{array} \parallel x := 2 \quad (W+RW+W)$$

Under SC, as depicted by the executions in Fig. 1, the read in thread 2 may read either 0 (from the initialization write), 1 (from the write in thread 1), or 2 (from thread 3).

Our goal is to *enumerate* such executions systematically. A simple approach taken, e.g., by HERD [6] and CPPMEM [8], is to enumerate *all* possible executions and filter them according to the consistency predicate of the memory model.

To do this, we require that the underlying memory model satisfy the following (see §3):

MM1: $porf$ is irreflexive, where $porf \triangleq (po \cup rf)^+$

This requirement is satisfied by several models (e.g., SC, TSO, PSO, and RC11), and ensures that loop-free programs have finitely many executions. Without this requirement, we can easily run into problems as the following program illustrates:

$$x := y \parallel y := x \quad (LB+DEP)$$

Under the (arguably useless) memory model that deems every execution graph consistent, the program can return $x = y = v$, for *any* value v , by having both threads read v and write v in a circular fashion as shown below:

In the weak memory literature, such executions are considered problematic because they generate values “out of thin air” (OOTA) [10, 31, 41] and inhibit compositional reasoning.

Remark 1. While restricting OOTA behaviours, **MM1** also precludes models allowing the outcome $a = b = 1$ for the following “load buffering” litmus test:

$$\begin{array}{l} a := y; \\ x := 1 \end{array} \parallel \begin{array}{l} b := x; \\ y := 1 \end{array} \quad (LB)$$

A few models allow this outcome and yet avoid OOTA executions. The Power [6] and ARM [37] models record (syntactic) dependencies in executions and forbid dependency cycles, while the Promising [23] and WeakestMO [11] models are not even defined in terms of execution graphs. Handling these models is beyond the scope of this paper.

2.1 Checking Consistency at Every Step

Even without OOTA executions, generating *all* executions and then checking consistency does not scale [24].

A much better approach, followed by most tools (e.g., [1, 2, 4, 24, 35]), is to construct executions *incrementally* by adding events one at a time and checking for consistency at each step, thereby avoiding the exploration of inconsistent graphs. For this approach to work, the underlying memory model must satisfy the following condition:

Every non-empty consistent graph has a po -maximal event that, if removed, yields a consistent graph.

This condition ensures that each execution can be generated by adding its events in *some* total extension of the $porf$ order, and checking for consistency after each step. For instance, execution ② in Fig. 1 can be generated by adding its events in the following order: $W(x, 0)$, $W(x, 1)$, $R(x)$, and $W(x, 2)$.

2.2 Fixing the Graph Construction Order

To generate all executions of a program following the condition of §2.1 one must in principle consider *all* possible extensions of $porf$. This, however, very often leads to duplicate explorations.

Therefore, ideally, one would generate all executions without considering all possible extensions of $porf$, regardless of the memory model. In fact, we can do this for all well-known memory models. In particular, models such as SC, TSO, PSO, and RC11 all satisfy an even stronger guarantee, namely *prefix-closedness*:

MM2: There exists a partial order R that includes reads-from and (preserved) program order such that, if a graph is consistent, so is every R -prefix of it.

This ensures that to generate a particular execution, it is sufficient to consider *any* total extension of $porf$.

¹In all our examples, we use x, y, z as global (shared) variables and a, b, c as local variables. All variables are implicitly initialized to 0.

As we demonstrate below, we can leverage this fact and *fix* an order in which we add execution events one at a time, thus generating all executions of a program systematically.

2.3 GENMC: A First Example

Let us run our model checking algorithm, GENMC, to generate the executions of $\mathbf{w+rw+w}$ by adding its events in a *fixed* order given by thread identifiers: first the events of (the left-most) thread 1, then the events of thread 2, and so forth.

We start with an initial graph G_0 containing only the initialization write $W(x, 0)$ (see below). First, we add the $W(x, 1)$ write of the first thread to G_0 , simultaneously adding the appropriate *po* edge between the events:

Continuing in thread order, we next add the $R(x)$ read of thread 2, which may read from either of the writes in the graph, yielding two distinct graphs (one for each case):

However, recording both graphs is inefficient: in the general case, we need to record one graph for each of the reads-from options of each read. Note that the two graphs are identical up to the read, which is the point of divergence. As such, each time we add a read that can read from more than one place, we proceed with one of the options, e.g., $W(x, 0)$, and record the alternative(s), i.e., $W(x, 1)$, into a work list W for later exploration. W maps each read to a list of writes it can also read from; in this case, the current graph along with W is given below. We refer to revisit options such as $W(x, 1)$ as *forward* revisits since they are already in the graph when the read ($R(x)$) is added to the graph.

Since the value read is 0, we next add the $W(y, 1)$ write of thread 2. Finally, we add $W(x, 2)$ of thread 3 which yields the graph below. Note that, as it is consistent for the read to read 2 from this newly added write, we also record this new reads-from as a revisit option in W . We refer to revisit options such as $W(x, 2)$ as *backward* revisits since they are added to the graph after the corresponding read ($R(x)$).

This first execution is now *completed* (denoted by the highlighted background): it corresponds to execution ① of

Fig. 1. To generate the remaining executions, we *revisit* the graph by picking an alternative reads-from option from W .

Suppose that we next pick $W(x, 2)$ from W . To continue, we restrict the graph to contain only the events added to the graph *prior* to (and including) the read (i.e., $W(x, 0)$, $W(x, 1)$ and $R(x)$), as well as the events that led up to (in *porf* order) the revisiting write $W(x, 2)$. This yields the complete graph below, corresponding to execution ② in Fig. 1. The $W(x, 2)$ option is marked as ✓ to denote that it has been considered.

When considering the alternative reads-from option $W(x, 2)$, we restrict the graph to contain only the events added *prior* to the read. Restricting the graph is important because events added after the read may depend on its value. For instance, it is crucial to remove $W(y, 1)$ as it is only present when 0 is read from x . Similarly, we must retain the events added before (in *porf* order) the alternative reads-from option $W(x, 2)$. For instance, if $x := 2$ in $\mathbf{w+rw+w}$ is wrapped in the conditional **if** $y = 1$ **then**, the presence of the $W(x, 2)$ event in the graph depends on the value read for y , i.e., the events before $W(x, 2)$ in *porf* order.

To generate the last execution, we revisit the graph once again by picking the remaining option $W(x, 1)$ in W . We then restrict the graph as before, yielding the graph below:

To continue, we add the $W(x, 2)$ write arriving at the graph below, corresponding to execution ③ in Fig. 1. Note that we do not re-add this write as an entry in W , as this option has already been explored (✓-marked).

Finally, as the graph is complete, and all options in W are explored, the algorithm terminates.

Avoiding Duplication. When revisiting a read event, write events may be removed from the graph and later re-added. As such, additional care is required to avoid duplicate backward revisits. For instance, continuing from (Ex2), by picking the next option in W ($W(x, 1)$), we removed $W(x, 2)$ arriving at (Pre3). We later re-added $W(x, 2)$ and obtained (Ex3). In doing so, we did not re-add $W(x, 2)$ as a (backward) revisit option to W as this option had already been explored before. Rather, by having previously marked $W(x, 2)$ as explored (✓-marked), we ascertained that $W(x, 2)$ is indeed a duplicate revisit. To

this end, as we describe in §4, backward options are not removed from W ; instead they are marked as explored (e.g., in (Pre3) and (Ex3)). By contrast, forward revisits do not lead to duplication. This is because when revisiting a read (e.g., $R(x)$), only events added after the read are removed from the graph. As such, since a forward option (e.g., $W(x, 1)$) is added to the graph *before* the read, it is not removed from the graph, and therefore not re-added, avoiding duplication. For efficiency, we thus remove forward options from W once explored (e.g., $W(x, 1)$ is removed in (Pre3) and (Ex3)).

2.4 GENMC: Extensible Memory Models

Note that as described in §2.3, GENMC generates all executions, even though it does not add events in *porf* order. This is because in cases where a read is added before the write it reads from, e.g., reading from $W(x, 2)$ in ②, the *rf* edge is recorded as an option in W once the write is added.

This then leads to the question, could events added after a read affect the consistency of the execution in a way that the write is never added and hence the alternative *rf* option is never considered? Perhaps surprisingly, the answer is yes. For example, consider the following program under a (contrived) memory model that dictates “if a read of y reads 0, then there cannot be a read of x that also reads 0”:

$$a := x \parallel b := y \parallel x := 42 \quad (\text{R+R+W})$$

In this case, adding the events in thread order results in a graph where both x and y read 0, which is then dropped as inconsistent, and thus we cannot generate the execution where the first thread reads 42. This brings us to our third requirement on memory models, *extensibility*:

MM3: Given a consistent execution G , a po-maximal event can always be added to G to yield a consistent execution (with an appropriate *rf* edge when applicable).

This requirement holds for all well-known memory models, and excludes “nonsensical” memory models such as that above. In particular, under that model, the consistent execution of R+R+W comprising the initialization events and $R(x)$ of the first thread reading 0 cannot be extended by adding $R(y)$ for any choice of *rf*.

2.5 GENMC: Modification Order and Writes-Before

Recall that using GENMC, we generated all three executions of W+RW+W under SC in Fig. 1. These executions, however, do not exactly correspond to the notion of executions in the formal definition of SC: as discussed above, SC executions additionally record the modification order *mo*, which totally orders all writes to a given memory location. We refer to such execution (which record *mo*) as *mo-executions*.

As such, the three executions in Fig. 1 correspond to the six *mo-executions* depicted in Fig. 2. In this program, each execution corresponds to two *mo-executions* representing the two ways $W(x, 1)$ and $W(x, 2)$ could be ordered by *mo*.

Figure 2. *mo-executions* of W+RW+W under SC.

One can of course adapt GENMC to enumerate all *mo-executions*, as e.g., in [24]; but doing so is wasteful because while the choice of *mo* can affect the consistency of an execution, it is not directly observable by the program. As long as checking for consistency is reasonably efficient, enumerating only (plain) executions is better because it searches through a space that is up to exponentially smaller.²

Now, how can we check consistency of an execution besides naively enumerating all *mo* possibilities? The idea is to compute the “writes-before” (*wb*) relation, which records the set of *mo*-edges whose direction is forced because of the *rf*-edges. Let us consider the following executions under SC:

In execution ①, the $W(x, 1)$ must *write-before* $W(x, 2)$: otherwise, the read may only read 1, due to coherence. Of course, the initialization write writes before the writes of both threads, as it is *po*-before them. By contrast, in execution ②, the writes of the three threads are not *wb*-ordered, as there is no causal ordering amongst them.

Computing *wb* can be done in cubic time, and yields a complete procedure for checking consistency for RC11 without SC features. For SC, while checking consistency of an execution is NP-complete [17], a *wb*-based check can approximate it extremely well³.

2.6 GENMC: Handling *rf*-Functionality Constraints

Memory models may prescribe *rf-functionality* constraints requiring that certain writes be read by at most one read. For instance, in case of the RMW (read-modify-write) instructions, e.g., CAS (compare-and-swap) or FAI (fetch-and-increment), to ensure their atomicity, two RMW events may

²To see that, consider an extension of W+RW+W with n parallel writes and one reader: that program has $n + 1$ executions and $(n + 1)!$ *mo-executions*.

³The definition of *wb* can be found in our technical appendix [25].

Figure 3. The executions of the **FAI/2** program.

not read from the same write. These constraints are, however, not exclusive to RMWs. For instance, as we discuss in the upcoming section, a lock library may require **rf**-functionality to ensure mutual exclusion. Indeed, as shown in [38], many well-known concurrent libraries require **rf**-functionality to ensure correct synchronization.

Handling such constraints requires additional care. Consider the program below and its executions depicted in Fig. 3:

$$a : \text{FAI}(x) \parallel b : \text{FAI}(x) \quad (\text{FAI/2})$$

Execution ① captures the case where thread 1 increments x first, while ② captures the case where thread 2 increments x first. Let us run GENMC on this example. We proceed by adding the RMW instruction of thread 1 (a) which reads from the initialization write. When we next add the RMW instruction of thread 2 (b), to ensure atomicity, there is only one consistent option for b to read from, namely a . However, this poses a problem: when b is added, it cannot add b to W as a revisit option for a (since that would create a **porf** cycle). As such, the algorithm fails to generate execution ②.

To remedy this, we allow for *temporary inconsistency* in the graph. More specifically, we push to W options that break such consistency constraints.

When this inconsistent execution is eventually picked from W , during the course of its exploration, we may encounter events that can revisit one of the events responsible for inconsistency, thus obtaining a consistent graph. The inconsistent execution is then dropped.

In our example, we push to W an entry for b to read from the initial write, and continue with the consistent option:

Figure 4. Executions of the **LOCK/2** program.

Finally, we pick the remaining revisit option in $W[a]$, restrict the graph as before and arrive at execution ②.

2.7 GENMC: Model Checking for Libraries

We next explain how GENMC generalizes to models incorporating high-level (abstract) libraries. To do so, let us consider a mutex library with *lock* and *unlock* instructions.

Although the mutex library does not have conventional read and write operations, its primitives behave very much like reads and writes. Intuitively, *unlock* can be viewed as a write, while *lock* can be viewed as a read that may either read from an initial value (i.e., acquiring the mutex immediately after it is initialized), or read from an *unlock* instruction (i.e., acquiring the mutex after it has been released by its previous holder). As with RMWs, the mutex library requires **rf**-functionality: no two *lock* events read from the same place, capturing the exclusivity of the mutex while held.

An interesting feature of the mutex library is that the calls to *lock* may *block* if the mutex is taken. Put formally, when all writes (initialization and unlocks) in an execution have already been read-from, due to **rf**-functionality, when adding a read (*lock*) event e to the graph, there may not exist a write from which e could read. When this is the case, the read event e blocks in that its thread cannot make progress and thus e has no **porf** successors. Note that in such libraries **rf** is not necessarily total on reads. However, lock events may not block arbitrarily: a lock may block only when all writes are read from; i.e., when the mutex is taken. This brings us to our final requirement on memory models, *well-blocking*:

MM4: Given a consistent execution G : 1) blocking reads in G have no **porf** successors; and 2) if G contains a blocking read, then all writes in G are read from.

Consider the program below with its executions in Fig. 4:

$$a : \text{lock}(l); \parallel b : \text{lock}(l); \quad (\text{LOCK/2})$$

$$a' : \text{unlock}(l); \parallel b' : \text{unlock}(l);$$

Note that neither lock call may block as the program contains sufficient writes: two unlocks and the implicit initialization.

Running our algorithm on this example, we add the events in order (a, a', b, b') and obtain execution ①. As with **FAI/2**, when adding b to the graph, we also consider inconsistent reads-from options and add them to the work list, arriving at the following configuration:

We then pick the next option for b and restrict the graph as before. As in the **FAI/2** example, we check whether b may itself generate backward revisit options for the reads in the graph. However, since $b : \text{lock}(l)$ is only a read event (in contrast to $b : \text{RMW}(x, 2)$ in **FAI/2** which is also a write), a cannot read from b . Nonetheless, by reading from the initialization event, b causes a to *block*. That is, blocking ($\perp$) is added as a revisit option for a . This graph is subsequently dropped as inconsistent (violating *rf*-functionality):

We next pick $\perp$ as a revisit option for a . Since a is now blocking, its thread cannot proceed and its subsequent events are skipped. We thus next add b' to the graph. As b' is a write, it may revisit a and is added as an option in $W[a]$. However, adding b' renders the graph inconsistent (a is blocking despite the available b') and is thus dropped:

Finally, we consider the last revisit option (b') for a . After restricting the graph, we add event a' and obtain ② in Fig. 4.

Note that running GENMC on **LOCK/2** was no different from running it on **FAI/2** and required no special treatment: we merely used the lock library consistency check rather than that of RC11. Indeed, the main difference between the two examples is the blocking behaviour of locks, which is prescribed by the lock library specification. As such, GENMC can be adapted to *any* memory model that meets the conditions in **MM1-MM4**. We next formalize these conditions.

3 Formal Model

We describe a framework for axiomatic memory models (MMs) and instantiate it to specify a mutex library. In the technical appendix [25], we present the SC [28], TSO [36] and RC11 [27] models as instances of this framework.

Execution Graphs. The traces of a program are represented as a set of *execution graphs*, where each graph G comprises: (i) a set of events; and (ii) a number of relations on events.

An event is a tuple of the form $\langle i, n, l \rangle$, where $i \in \text{Tid} \cup \{0\}$ is a *thread identifier* (0 for initialization events) with $\text{Tid} \subseteq \mathbb{N}$, $n \in \mathbb{N}$ is the *serial number* inside a thread, and $l \in \text{Lab}$ is an event *label*. The serial number of an event denotes its index (from 1) within its thread; e.g., the first event of a thread has serial number 1. Serial number 0 is reserved for initialization events. A label may be either: (i) the *error* label *error* (denoting assertion violations); or (ii) the *stuck* label *stuck* (e.g., due to a failed **assume** statement); or (iii) a memory model-specific label, e.g., the write label $W(x, 1)$ for writing 1 to x under the SC model. The *label function* lab returns the label of an event. We assume a set of locations Loc ; the loc function returns the location of a label.

Definition 3.1 (Executions). Given designated sets of *read* (R) and *write* (W) events, an *execution* is a tuple $G = \langle E, \text{rf} \rangle$, where E is a sequence of *events*, and $\text{rf} : E \cap R \rightarrow E \cap W$ is the *reads-from* function.

The sets of read and write events are designated by the memory model and are not necessarily low-level reads/writes. For instance, in case of the mutex library, *lock* and *unlock* events constitute read and write events, respectively.

Recall from §2.2 that to generate program executions using our algorithm, it suffices to fix the construction order. This is given by the order of events in the sequence E .

Given an execution G , we write $G.E$ and $G.\text{rf}$ for its components, and write $G.R$ (resp. $G.W$) for $G.E \cap R$ (resp. $G.E \cap W$). We write $G.E_i$ for $\{\langle i', -, - \rangle \in G.E \mid i = i'\}$; and write $G.\text{po}$ for the *program order* defined as follows:

$$G.\text{po} \triangleq G.E_0 \times (G.E \setminus G.E_0) \cup \left\{ \langle \langle i_1, n_1, l_1 \rangle, \langle i_2, n_2, l_2 \rangle \rangle \mid \langle i_1, n_1, l_1 \rangle, \langle i_2, n_2, l_2 \rangle \in G.E \setminus G.E_0 \wedge i_1 = i_2 \wedge n_1 < n_2 \right\}$$

In general, $G.\text{rf}$ may not be a total function: read events that do not read from any event are used to model blocking library events, such as a blocking lock event that is awaiting the release of a mutex. We write $G.B \triangleq G.R \setminus \text{dom}(G.\text{rf})$ for the set of *blocked events*. Finally, although $G.\text{rf}$ is a function, we often implicitly coerce it to a relation on $W \times R$.

Notation. Given a relation r and a set A , we write $r^?$, r^+ and r^* for the reflexive, transitive and reflexive-transitive closure of r , respectively. We write $\text{dom}(r)$ and $\text{rng}(r)$ for the domain and range of r , respectively. We write r^{-1} for the inverse of r ; $r|_A$ for $r \cap (A \times A)$; and $[A]$ for the identity relation on A : $\{(a, a) \mid a \in A\}$. Given relations r_1 and r_2 , we write $r_1; r_2$ for $\{(a, b) \mid \exists c. (a, c) \in r_1 \wedge (c, b) \in r_2\}$, i.e., their relational composition. Given an event set E , we write E_x for $\{e \in E \mid \text{loc}(e) = x\}$, and $G|_E$ for $\langle E', G.\text{rf}|_{E'} \rangle$ with $E' \triangleq G.E \cap E$. We write $G.\text{porf}$ for $(G.\text{po} \cup G.\text{rf})^+$, and write $G.\text{rf}[r \mapsto w]$ for the graph obtained from mapping $G.\text{rf}(r)$ to w . Finally, we write $\#$ for sequence concatenation.

Extension. We define graph *extension* in Def. 3.2, used by the incremental construction in GENMC, which describes

adding an *available* event to an execution. Given an execution G , an event $\langle i, n, - \rangle$ is available when thread i contains $n - 1$ events, none of which are blocking. As executions are constructed incrementally by adding one available event at a time, it follows that the events of each thread i are indexed with adjacent integers $1 \cdots |G.E_i|$.

Definition 3.2 (Extension). An event $e = \langle i, n, l \rangle$ is *available* for an execution G if $|G.E_i| = n - 1$ and $G.E_i \cap G.B = \emptyset$. The *extension* of G with an available event e , written $\text{Add}(G, e)$, denotes the execution $\langle E \uplus [e], G.\text{rf} \rangle$.

Consistency and Memory Model Assumptions. Given a program P , the admissible behaviours of P are commonly described as a set of *consistent* executions. Consistency of an execution is memory model (MM)-specific; as such, MMs often define a consistency predicate that prescribes the conditions required for consistency. As our model checking technique is MM-parametric, we assume the existence of such a consistency predicate: given an execution G , we write $\text{cons}_m(G)$ to denote that G is consistent under memory model m .

Recall from §2 that we require underlying memory models to satisfy certain properties as outlined by **MM1**–**MM4**. In what follows, we formally define these conditions.

The first condition (**MM1**) is captured by Def. 3.3. This well-formedness condition additionally requires that the MM be agnostic to the order in which events are added to the graph, as it constitutes auxiliary instrumentation used by our algorithm. As such, execution consistency must be independent of this order: if $\langle E, \text{rf} \rangle$ is consistent then $\langle E', \text{rf} \rangle$ is also consistent, where $E' \in \text{perm}(E)$ is a permutation of E .

Definition 3.3 (Well-formedness). An execution G is *well-formed* if $G.\text{porf}$ is irreflexive. A memory model m is *well-formed* iff for all G , if $\text{cons}_m(G)$ holds, then G is well-formed, and $\forall E \in \text{perm}(G.E). \text{cons}_m(\langle E, G.\text{rf} \rangle)$.

The prefix-closedness condition (**MM2**) is captured by Def. 3.4. A consistency model m is commonly considered *prefix-closed* iff: given a consistent execution G and a porf -closed set of events $E \subseteq G.E$ (i.e., $\text{dom}(G.\text{porf}; [E]) \subseteq E$), restricting the graph to those events in E yields a consistent execution, i.e., $\text{cons}_m(G|_E)$. However, this definition is too strong due to blocking reads.

To see this, consider the program $l_1 : \text{lock}(l) \parallel l_2 : \text{lock}(l)$. Under the mutex specification described in §2.7, one consistent execution of this program is a graph G in which l_1 reads from mutex initialization, whilst l_2 blocks. Let $E = \{l_2, \text{init}\}$; if we now restrict G to E , the resulting graph is inconsistent since l_2 blocks despite the available initialization event.

We thus weaken prefix-closedness by requiring that there exist a set of blocking events $B \subseteq E$ such that the graph restricted to $E \setminus B$ is consistent: $\text{cons}_m(G|_{E \setminus B})$. For instance, in the example above we can pick $B = \{l_2\}$. Note that for well-known memory models such as SC, TSO and RC11, the

strong and weak notions of prefix-closedness coincide, as these models do not contain blocking events.

Definition 3.4 (Prefix-closedness). A memory model m is *prefix-closed* iff for all $G, E \subseteq G.E$, if $\text{dom}(G.\text{porf}; [E]) \subseteq E$ and $\text{cons}_m(G)$, then there exists $B \subseteq G.B$ such that $\text{cons}_m(G|_{E \setminus B})$.

Memory model extensibility (**MM3**) is captured in Def. 3.5 and requires that a memory model be *read*-, *write*- and *rw*-*extensible*. The first two requirements are intuitive and stipulate that a consistent execution can always be extended by a read or write event, respectively. The *rw*-extensibility imposes certain conditions on events that are both read and write events (e.g., RC11 RMW events). These requirements are rather technical and are necessary for the correctness of our algorithm (see the technical appendix [25]).

Definition 3.5 (Extensibility). A memory model m is *read-extensible* iff for all $G, r \in R$ and $G' = \text{Add}(G, r)$, if $\text{cons}_m(G)$, there exists $w \in G.W \cup \{\perp\}$ such that $\text{cons}_m(G'.\text{rf}[r \mapsto w])$.

A memory model m is *write-extensible* iff for all $G, w \in G.W$, if $\text{cons}_m(G|_{G.E \setminus \{w\}})$ and $\text{rng}([w]; G.\text{porf}) = \emptyset$, then $\text{cons}_m(G)$.

A memory model m is *rw-extensible* iff for all G, r, w, u , if $\text{cons}_m(G)$, $u, u' \in G.R \cap G.W$ and $\text{rng}([u]; G.\text{po}) = \emptyset$, then:

- if $\langle u, r \rangle \in G.\text{rf}$ and $\text{rng}([r]; G.\text{po}) = \emptyset$, then there exists $w \in G.E \setminus \{u\}$ such that $\text{cons}(G.\text{rf}[r \mapsto w])$; and
- if $\langle w, u \rangle, \langle u, u' \rangle \in G.\text{rf}$ and $\text{rng}([u']; G.\text{porf}) = \emptyset$, then $\text{cons}(G|_{G.E \setminus \{u\}}.\text{rf}[u' \mapsto w])$.

A model is *extensible* iff it is read-, write- and rw-extensible.

Finally, the well-blocking condition (**MM4**) is captured by Def. 3.6. It stipulates that consistent executions satisfy two conditions with respect to blocking reads. First, blocking reads must be maximal in $G.\text{porf}$: if an event blocks then it cannot proceed. Second, reads may block only when all writes are *matched*. That is, if there is a blocking read on x ($G.R_x \not\subseteq \text{dom}(G.\text{rf})$), then all writes on x have already been read-from ($G.W_x \subseteq \text{rng}(G.\text{rf})$). Note that when $G.\text{rf}$ is a total function, this stipulation is trivially satisfied. As such, this is not a strong requirement: in all well-known memory models as well as the concurrent libraries specified in [38], $G.\text{rf}$ is specified to be total.

Definition 3.6 (Well-blocking). A memory model m is *well-blocked* iff for all G , if $\text{cons}_m(G)$ holds, then G is well-blocked.

An execution G is *well-blocked* iff 1) $[G.B]; G.\text{porf} = \emptyset$; and 2) $\forall x \in \text{Loc}. G.R_x \subseteq \text{dom}(G.\text{rf}) \vee G.W_x \subseteq \text{rng}(G.\text{rf})$.

From Programs to Executions. Given a concurrent program, we use the same technique as [24] to pre-process it to a program of the form $P = \parallel_{i \in \text{tid}} P_i$, where each P_i is a sequential loop-free deterministic program. The *set of executions* associated with P is then defined by induction over the structure of sequential programs P_i . We omit this formal construction here as it is standard in the literature e.g., [41].

Mutex Library. We formulate the notion of mutex library executions and their consistency predicate in Def. 3.7 below. For each mutex at location $l \in \text{Loc}$, the mutex events on l comprise lock and unlock events, where the set of unlock events contains a single initialization event. Given a mutex execution $G = \langle E, rf \rangle$, we define the *mutex consistency predicate* such that it holds on G if: 1) G is well-formed (Def. 3.3); 2) G is well-blocked (Def. 3.6); 3) E comprises mutex events; 4) rf is injective; and 5) rf maps lock events on to unlocks.

Intuitively, rf describes the order of mutex acquisition. For each lock event b with $\langle a, b \rangle \in rf$, if a is an unlock event, then a denotes the event releasing the mutex immediately before it is acquired by b ; when a is the initialization event, then b corresponds to the very first *lock* call on the mutex. As such, rf must be an injection.

Note that not all locks may be matched in rf . Unmatched locks are *blocked*, waiting for the mutex release. However, well-formedness ensures that an execution contains blocking locks only when all unlocks are matched (see (2) in Def. 3.6).

Definition 3.7. The *mutex event set* on l is $\text{MX}_l \triangleq L_l \uplus U_l$ with $L_l \triangleq \{e \mid \text{lab}(e) = \text{lock}(l)\}$, $U_l \triangleq \{e \mid \text{lab}(e) = \text{unlock}(l)\}$.

Execution G is *mutex-consistent*, written $\text{cons}_{\text{mx}}(G)$, iff: 1) G is well-formed; 2) G is well-blocked; 3) $G.E = \bigcup_{l \in \text{Loc}} \text{MX}_l$; 4) $G.rf$ is injective; and 5) $G.rf = \bigcup_{l \in \text{Loc}} rf_l$ for some given $rf_l \subseteq U_l \times L_l$.

It is straightforward to show that $\text{cons}_{\text{mx}}(\cdot)$ is well-formed, prefix-closed, extensible and well-blocked.

4 GENMC: The Generic Model Checker

In this section, we present a version of our model checking algorithm, GENMC, that does not record mo . It can be instantiated for any memory model by replacing the consistency checks in the code with MM-specific consistency predicates. We refer the reader to our technical appendix [25] for a version of GENMC that also tracks mo .

Configurations. Given a program P , recall from §2 that GENMC maintains a *configuration* comprising an execution G of P , and a work list W which stores revisit options both explored or otherwise. As described in §2.3, the options in W are categorized as forward or backward revisits; forward options are removed from W once explored, whilst backwards options are never removed and simply marked as explored.

Formally, we define a configuration as a tuple $\langle G, T, U, S \rangle$, where G is an execution of P ; T denotes a set of *revisitable reads*; U is a map from reads to *backward* revisits (both explored or otherwise); and S is a map from reads to both forward and backward revisits *yet to be explored*. As such, when a new revisit candidate is encountered, if it is a forward option, it is added only to S , whereas if it is a backward option then it is added to both S and U . That is, S serves as a work set (the W map in §2 limited to entries not $\checkmark$ -marked). Analogously, when a revisit is explored, it is only removed

Algorithm 1 Main exploration algorithm

```

1: procedure VERIFY( $P$ )
2:    $\langle G, T, U, S \rangle \leftarrow \langle G_0, \emptyset, \emptyset, \emptyset \rangle$ 
3:   VISITONE( $P, G, T, U, S$ )
4:   while  $\langle r, G' \rangle \leftarrow \text{RemoveMax}(S)$  do
5:      $\langle E_1, r, E_2 \rangle \leftarrow \text{split}(G.E, r)$ 
6:      $T \leftarrow T \cup E_2$ 
7:      $U \leftarrow U \cup \{U[r'] \mid r' \in E_2\}$ 
8:     if  $G'.rf[r] \neq \perp$  then
9:       CALCREVISITS( $G', T, U, S, r$ )
10:    VISITONE( $P, G', T, U, S$ )

```

from S and not U , and thus U retains all backward revisits. For efficiency, the revisitable set T tracks those reads whose incoming rf edges may be changed, i.e., revisit candidates.

Each entry in $S[r]$ (and $U[r]$) is a graph G' representing the effect of revisiting r by a write w . As we discuss later in §5, our implementation records only a *portion* of G' necessary for constructing it from G when r is revisited by w . However, for better readability, in our presentation here we record in G' the *entire* graph resulting from w revisiting r .

The next_P Function. Recall that we construct graphs by adding events in a *fixed* order (§2). We define a function, next_P , such that given a program P and an execution G of P , $\text{next}_P(G)$ returns an available event (Def. 3.2) of *any* thread i in G such that i is not stuck (e.g., due to a failed **assume** statement) and has not finished execution. When no such thread exists (i.e., all threads are stuck or finished), next_P returns *false*. We implement next_P to choose the left-most such thread, i.e., one with the smallest thread identifier.

4.1 The Main VERIFY Procedure

Given a program P , we begin exploring the executions of P by calling $\text{VERIFY}(P)$. This routine creates an initial configuration comprising the G_0 graph (containing only the initialization writes), an empty revisit set $T = \emptyset$, and empty maps $U = S = \emptyset$ (Line 2). It then generates the executions of P one at a time. This is done by calling $\text{VISITONE}(P, G, T, U, S)$ on Line 3, which fully explores *one* execution extending G , and pushes alternative reads-from options encountered to the work set S . Once $\text{VISITONE}(P, G, T, U, S)$ returns the full execution generated, remaining executions are generated by exploring the options in the work list S Lines 4–10.

To do this, an option G' is picked from $S[r]$ (Line 4) such that r is the *maximal* entry in S : r is added to the current graph G *after* all other reads in the domain of S . When $S[r]$ holds multiple options, an arbitrary entry is chosen. Picking the maximal entry in S makes it easier to update the current configuration and enables a key optimization (see §5).

We split G at r (Line 5) such that E_1 contains events in G added before r and E_2 contains those added after r . By construction, E_2 comprises events that either are not in G'

Algorithm 2 Explore one program execution

```

1: procedure VISITONE( $P, G, T, U, S$ )
2:   while  $\text{cons}(G) \wedge a \leftarrow \text{next}_P(G)$  do
3:     if  $a \in \text{error}$  then exit("erroneous program")
4:      $G \leftarrow \text{Add}(G, a)$ 
5:     if  $a \in R$  then
6:        $W \leftarrow G.E \cap W_{\text{loc}(a)} \cup \{\perp\}$ 
7:       choose some  $w_0 \in W$ 
8:        $G.\text{rf}[r] \leftarrow w_0$ 
9:        $T \leftarrow T \cup \{r\}$ 
10:       $S[a] \leftarrow \{G.\text{rf}[a \mapsto w] \mid w \in W \setminus \{w_0\}\}$ 
11:       $\text{CALCREVISITS}(G, T, U, S, a)$ 

```

or belong to the porf prefix of the event a that revisited r to generate G' . These latter events are responsible for the addition of a to the graph, and consequently the reason why r is revisited. As such, revisiting any of these latter events would “undo” the revisit of r . For this reason, we remove the events in E_2 from the set T of revisitable reads (Line 6).

Analogously, Line 7 removes the E_2 entries from U . Note that no such entries exist in S : all events in E_2 have been added to the graph *after* r , while we picked r to be the maximal entry in S ; i.e., $S[r']$ contains no entries for r' in E_2 .

Recall from §2 that when revisiting a (non-blocked) read, we check whether the read being revisited may itself generate backward revisit options for existing reads in the graph. For instance, in the **FAI/2** and **LOCK/2** examples, revisiting b generated additional revisit options for a —see (**FAI- $\perp$**) and (**LOCK- $\perp$**). This is done by calling **CALCREVISITS** on Line 9. Finally, on Line 10 we explore the updated configuration.

4.2 The VISITONE Procedure

The VISITONE procedure is the workhorse of the exploration algorithm. In each iteration of this loop, while the current graph (G) is consistent, it is extended with its next event a (given by $\text{next}_P(G)$, see page 8). When $\text{next}_P(G)$ returns *false*, VISITONE terminates. If the next event a is an assertion violation, then an error is reported (Line 3), and the algorithm terminates. Otherwise, we add a to the graph (Line 4). As before, we check whether the newly added event a generates backward revisit options for the existing reads in the graph by calling **CALCREVISITS** on Line 11.

If the new event a is a write, no additional work is required. However, if a is a read, we must calculate its incoming rf edge. We first calculate the set of writes W that a could read from, i.e., its forward revisit options (Line 6), choose a write w_0 for the current exploration (Line 7), set a to read from w_0 in G (Line 8), add the new read to the revisit set (Line 9), and push the remaining revisit options to S (Line 10).

Algorithm 3 Calculate which reads should be revisited

```

1: procedure CALCREVISITS( $G, T, U, S, a$ )
2:    $p_a \leftarrow \text{dom}(G.\text{porf}^2; [a])$ 
3:   for  $r \in T \cap R_{\text{loc}(a)} \setminus p_a$  do
4:      $\langle E_1, r, E_2 \rangle \leftarrow \text{split}(G.E, r)$ 
5:      $G' \leftarrow G|_{E_1 \# [r] \# (E_2 \cap p_a)}$ 
6:      $G'.\text{rf}[r] \leftarrow \text{if } a \in W \text{ then } a \text{ else } \perp$ 
7:     if  $G' \notin U[r]$  then
8:        $S[r] \leftarrow S[r] \cup \{G'\}$ 
9:        $U[r] \leftarrow U[r] \cup \{G'\}$ 

```

4.3 The CALCREVISITS Procedure

As described in §4.1–§4.2, the CALCREVISITS routine calculates the set of backward revisits that a can generate and pushes them to U and S unless they have already been considered, i.e., are in U (Lines 7–9).

To calculate the set of revisit graphs, we iterate through all revisitable reads on the same location as a (Line 3), excluding those reads whose revisit would violate porf -irreflexivity; i.e., those in the porf prefix of a calculated in p_a (Line 2). Recall that when a revisits r , in the resulting graph we retain r , the events added to the graph before r (E_1), as well as the events in the porf prefix of a that are added after r . To this end, we compute the sets E_1 and E_2 (Line 4), respectively comprising the events added before and after r , and set G' to contain E_1 , r , and the events in both E_2 and p_a (Line 5).

If a is a write event, we finally set r to read from a in G' (Line 6). If, however, a is a *read*, it cannot revisit existing reads in the graph itself but it may cause them to block (cf. **LOCK- $\perp$**), which is why we instead set the incoming rf edge of r to the blocking option $\perp$.

4.4 GENMC: Soundness, Completeness & Optimality

The GENMC algorithm (Algorithm 1) is *sound*, *complete* and *optimal*. Given a program P and a memory model m , soundness ensures that if GENMC generates G for P under m , then $\text{cons}_m(G)$ holds; completeness ensures that if G is an execution of P under m and $\text{cons}_m(G)$ holds, then GENMC generates G for P ; and optimality ensures that the P executions generated by GENMC under m are pair-wise distinct.

This is captured in the theorem below. The soundness proof is straightforward: GENMC checks consistency after each step, dropping inconsistent executions (Line 2 of VISITONE); as such, it only outputs consistent executions. The completeness and optimality proofs are non-trivial and are given in full in the technical appendix [25]; we proceed with an intuitive argument.

To show that GENMC is complete, we show that it generates *all* executions of a given program P . As discussed in §2, **MM1** and **MM2** ensure that every execution of P can be generated incrementally, by adding one event at a time. We

then demonstrate that each execution G of P generated incrementally can also be generated by GENMC if we *reshuffle* the order in which its events are added. That is, for each execution G generated by adding events in the order $S=e_1, e_2, \dots, e_n$, there exists a permutation S' of S , such that GENMC adds events in the S' order and generates G . To show that such a reshuffling exists, we often need to remove events from G and re-add them later (capturing the revisit step). This can always be done thanks to the extensibility property (MM3) ensuring that GENMC never gets stuck.

To show that GENMC is optimal, we observe that duplication can arise only when *revisiting* a read. As discussed in §2 (see “Avoiding Duplication” on page 3), forward revisits never cause duplication since they are never removed from the graph, while backward revisits may lead to duplication and thus the already-considered backward revisits are recorded in the map U . The optimality of GENMC is thus guaranteed by the properties of forward/backward revisits, the map U and the check on Line 7 of Algorithm 3.

Theorem 4.1 (Correctness). *The GENMC algorithm is sound, complete and optimal.*

5 Implementation

We have implemented GENMC as an open-source verification tool for C programs over the LLVM interpreter lli. GENMC is available at <http://github.com/mps-sws/genmc>.

We have implemented three variants of GENMC:

- LIB**: a generic variant that performs model checking on libraries, based on specifications provided by the user;
- WB**: an instantiation for the full RC11 memory model [27], using a consistency check based on **wb**; and
- MO**: an alternate RC11 instantiation that records the **mo** order during exploration. That is, whenever a write is added to a graph, we consider all its possible placements in **mo**, and create subexplorations for each case.

Naturally, the generic variant is slower than the RC11 ones because the latter have more optimized consistency checks; it is, however, still optimal.

Further, we have implemented some optimizations over the algorithm described in §4, which we will describe below.

The first key optimization has to do with the representation of the graphs to be revisited in S . In §4, each entry in $S[r]$ (and $U[r]$) is a full graph G' generated by a forward or backward revisit of the read r . For better space efficiency, rather than recording the entire graph G' , we store only the portion of G' of events after r , because that suffices for reconstructing the entire G' when the revisit takes place. The reason is that GENMC revisits executions from S by always choosing the *maximal* read in S when removing an entry from S (Line 4 of Algorithm 1). The effect of the revisit order is that the current graph projected to the events before r (i.e., E_1 in Algorithm 1) is exactly the same as the recorded graph $G' \in S[r]$ projected to the same events. As a result, it

suffices to record in each graph in S only the revisited read and the events after it.

Similarly, we store the graphs in U in a compressed form. Since we do not ever need to restore the graphs from U , we do not need to store all the events after r ; it suffices to record only their incoming **rf** edges because those determine the values read and hence the event labels.

Finally, in the WB and MO variants of GENMC, we use optimized consistency checks when adding a new event to the graph. We exploit the fact that the graph prior to adding the event was consistent, so it suffices to check only that the new event does not lead to any consistency violation.

6 Evaluation

Verification Tools. In the following, we compare the performance of GENMC to three other stateless model checkers: NIDHUGG [2], RCMC [24], and TRACER [4]. Initially, we also considered other tools—namely, CBMC [5, 14], CDS-CHECKER [35], and HERD [6]—but exclude them from head-to-head comparisons because they are typically significantly slower than NIDHUGG and RCMC and do not scale well (see, e.g., the evaluation in [24]): HERD because it was meant for experimenting only with small “litmus test” programs, CBMC because of the SAT solver, CDS-CHECKER because of its suboptimal partial order reduction technique.

NIDHUGG [2] is a state-of-the-art stateless model checker supporting SC, TSO, and PSO.⁴ It enumerates **mo**-executions (a.k.a. Mazurkiewicz traces [32]) and can operate both under an optimal mode (optimal-DPOR) and a non-optimal mode (source-DPOR). In our benchmarks, we use the source-DPOR version because it is typically faster than the optimal version. Under SC, NIDHUGG can also operate under a coarser equivalence partitioning (denoted SC^o – NIDHUGG with observers) [7]. This equivalence can be exponentially coarser than **mo**-executions, but remains exponentially finer than plain executions. We used version 0.3 of NIDHUGG, and ran it with the `-c11` switch, which makes the SC version of NIDHUGG noticeably faster.

RCMC [24] targets RC11 and WRC11, a weaker RC11 variant that does not record **mo** and does not enforce coherence. RCMC-RC11 also enumerates **mo**-executions of a program, though not optimally in the presence of RMW or SC accesses.

TRACER [4] targets RA (the release-acquire fragment of RC11), and enumerates plain executions. It is, however, built over the CDS-CHECKER infrastructure, which makes it quite difficult to apply it fairly to our benchmarks (e.g., it does not support assume statements, and requires manual instrumentation for programs with loops). For this reason, we apply it only to our synthetic benchmarks.

Benchmarks. We took as benchmarks all the programs from the benchmark suites of NIDHUGG and RCMC, together

⁴NIDHUGG also provides some very limited support for POWER, which we do not evaluate because it cannot encode most of our benchmarks.

Table 1. Lamport’s fast mutex algorithm [29]

	NIDHUGG		RCMC		GENMC		
	SC	SC ^o	RC11	WRC11	MO	WB	LIB
lamport(2)	0.13	0.10	0.04	∞	0.03	0.03	0.09
lamport(3)	7.53	4.49	5.40	∞	6.87	1.36	0.09
lamport(4)	–	–	–	∞	–	–	0.09

with some additional larger programs (e.g., seqlock, chase-lev) from open-source code. In total, we have assembled 127 benchmark programs, some of which are parametric in the number of operations/threads. For suitable values for their parameters, we have generated 202 test cases in total.

First (§6.1), we focus on the generic GENMC variant, and demonstrate how it is used to model check libraries. We conduct a case study for a lock library, and show that abstracting over its implementation has substantial runtime benefits.

Next (§6.2), we evaluate the overall performance of the RC11 variant of GENMC in both synthetic and real-world benchmarks. Our benchmarks highlight the importance of our optimality result, and show that GENMC verifies code currently deployed in production within seconds.

Finally (§6.3), we perform an extensive comparison between the WB and MO variants of GENMC. We show that the WB variant can explore *exponentially fewer* executions than MO, and the overhead due to its more expensive consistency checks is usually negligible.

Experimental Setup. We conducted all experiments on a Dell PowerEdge M620 blade system, running a custom Debian-based distribution, with two Intel Xeon E5-2667 v2 CPU (8 cores @ 3.3 GHz), and 256GB of RAM. We used LLVM 3.8.1 for RCMC and NIDHUGG. Unless explicitly noted otherwise, all reported times are in seconds.

6.1 Model Checking a Lock Library

As a simple demonstration of the benefits of parametricity and compositional verification, we consider a C implementation of Lamport’s fast mutual-exclusion algorithm [29] (see Table 1). We could have considered any correct lock implementation (e.g., the ones used in §6.2), but we chose Lamport’s algorithm because it has write-write races, which are rare in non-synthetic programs and highlight the differences between the various tools. NIDHUGG under TSO and PSO are excluded from this table for brevity, as they are slower than NIDHUGG-SC.

The first observation is that RCMC does not terminate under WRC11. This is because this test case has writes that are never ordered under WRC11, which makes the threads’ reads “oscillate” between the values of these writes ad infinitum. This behaviour is ruled out by RC11 and stronger memory models. Additionally, both RCMC and GENMC outperform NIDHUGG-SC (even though they explore more executions), with RCMC-RC11 being faster than GENMC-MO (see §6.2).

Table 2. Some synthetic benchmarks

	NIDHUGG		TRACER	RCMC		GENMC	
	SC	SC ^o	RA	RC11	WRC11	MO	WB
cinc(4)	2.98	3.11	1.13	0.69	0.67	0.43	0.45
cinc(5)	436.40	466.65	165.54	134.87	132.11	69.23	69.98
Nw1r(5)	1.25	0.17	0.01	0.11	0.05	0.08	0.03
Nw1r(8)	991.80	0.74	0.01	79.68	0.04	24.35	0.03

However, by feeding the axiomatic definition of the lock library to GENMC, and abstracting the inner working of the locks, GENMC is much faster than the other tools (shown in column LIB). For $N = 4$, for example, all other tools take more than 3 days to complete, whereas the generic variant of GENMC terminates almost instantly.

6.2 Overall Performance

Table 2 reports two synthetic benchmarks, which demonstrate the importance of optimality (Table 2).

In the cinc program, all threads perform a series of RMW operations. Since RCMC is not optimal in the presence of RMWs, it can explore many more executions than necessary, which leads to some runtime overhead. For 4 threads RCMC explores 45% more executions than GENMC, while for 5 threads, it explores almost twice as many executions as GENMC, and this is reflected in the running time. All other tools explore the same number of executions, but NIDHUGG is significantly slower than the other tools.

The Nw1r program has $N + 1$ concurrent writers and one concurrent reader of a shared variable, and thus has $(N + 2)!$ mo-executions versus only $(N + 2)$ plain executions. It is therefore not surprising that tools enumerating mo-executions (NIDHUGG-SC, RCMC-RC11, and GENMC-MO) do not scale well. NIDHUGG-SC^o explores 193 executions for $N=5$ and 2305 for $N=8$, and so also does not scale particularly well. In contrast, TRACER, RCMC-WRC11, and GENMC-WB finish almost instantly. Recall, however, that RCMC-WRC11 fails to terminate on other benchmarks (§6.1).

Next, we move to two sets of benchmarks extracted from real programs. Since NIDHUGG-SC^o does not reduce the number of executions and is in fact slower than NIDHUGG-SC on these benchmarks, we exclude it from further comparisons.

Table 3 compares the tools on the implementations of concurrent data structures from [13, 35]. We do not show the number of executions explored because all tools explore the same number of distinct executions⁵, excluding possible redundant executions explored by NIDHUGG (under 5%). These benchmarks have the same number of distinct executions regardless of the memory model (i.e., they are robust), which is expected since they only use non-SC accesses for performance reasons. The only exception is chase-lev, for

⁵NIDHUGG counts the number of executions that contain a failed `assume()` statement, while RCMC does not; we take this discrepancy into account.

Table 3. Data structure benchmarks from [13, 35]

	NIDHUGG			RCMC		GENMC	
	SC	TSO	PSO	RC11	WRC11	MO	WB
barrier(2)	0.12	0.14	0.16	0.04	0.04	0.04	0.03
barrier(3)	1.29	1.94	2.93	0.23	0.19	0.14	0.14
ms-queue(2)	0.36	0.63	0.76	0.10	0.11	0.07	0.07
ms-queue(3)	11.62	25.64	33.12	2.93	2.98	1.58	1.67
chase-lev(2)	3.06	7.46	29.95	0.79	0.80	0.32	0.32
chase-lev(3)	255.82	670.06	1.35h	79.79	81.44	19.82	19.40
linuxrwlocks(2)	0.28	0.33	0.42	0.06	0.07	0.05	0.06
linuxrwlocks(3)	26.93	50.40	64.23	5.09	5.03	3.13	4.66
mpmc-queue(2)	0.15	0.11	0.11	0.05	0.05	0.04	0.04
mpmc-queue(3)	135.46	265.13	339.30	69.55	70.13	50.77	54.45

barrier(N): A barrier implemented as a global flag with N threads that spinning and continuing only when all threads have reached the barrier.

ms-queue(N): The Michael-Scott queue with N threads, each enqueueing and (possibly) dequeuing an item.

chase-lev(N): An implementation of the Chase-Lev deque with one thread pushing and popping, and N threads stealing from the deque.

linuxrwlocks(N): A reader-writer lock ported from the Linux kernel. N threads read and/or write a shared variable while holding the lock.

mpmc-queue(N): A multiple-producer, multiple-consumer queue with N threads that enqueue and (possibly) dequeue.

which NIDHUGG explores more executions under PSO due to the absence of a store-store fence, which renders the precise modeling of acquire-release operations utilized by this benchmark difficult.

On these benchmarks, RCMC and GENMC outperform NIDHUGG, even though they operate under a weaker memory model. By contrast, NIDHUGG gets slower as the memory model gets weaker, which is expected due to the way it models TSO and PSO, and agrees with the observations in [2, 24]. GENMC performs similarly in terms of time under WB and MO, and explores the same number of executions. For linuxrwlocks, however, the WB verification requires much more time than MO. This is due to the calculation of **wb** as part of RC11’s consistency check, which is particularly slow when there are long chains of RMW events. (In general, calculating **wb** can take up to $O(n^3)$ time in the size of the execution graph, and achieves its worst-case complexity, when there are many writes to the same location.)

Table 4 summarizes the performance of the tools in lock implementations extracted verbatim from the Linux kernel (v4.13.6, v4.19.1). Headers, kernel primitives definitions, macros, and Kconfig options have been provided for all benchmarks as necessary. The test cases involve N threads accessing shared variables while holding the respective locks.

For all benchmarks, except mcs_spinlock, all tools explore the same number of executions, modulo a few redundant explorations for NIDHUGG, and the seqlock test case, where NIDHUGG-PSO again explores more executions due to the absence of a store-store fence. As shown, RCMC and GENMC outperform NIDHUGG by a large factor.

Table 4. Benchmarks extracted from the Linux-kernel

	NIDHUGG			RCMC		GENMC	
	SC	TSO	PSO	RC11	WRC11	MO	WB
mcs_spinlock(2)	0.12	0.09	0.10	0.05	0.05	0.05	0.05
mcs_spinlock(3)	2.98	6.84	12.54	0.84	0.67	0.89	0.78
mcs_spinlock(4)	0.68h	1.51h	3.32h	0.16h	0.15h	0.42h	0.26h
qspinlock(2)	0.17	0.11	0.11	0.04	0.04	0.04	0.04
qspinlock(3)	10.93	18.20	23.43	2.13	2.08	1.10	1.12
seqlock(2)	0.10	0.09	0.10	0.04	0.04	0.04	0.04
seqlock(3)	1.64	3.07	11.00	0.49	0.51	0.37	0.37

mcs_spinlock(N): An implementation of an MCS lock [33].

qspinlock(N): Queued spinlocks (1.2 KLOC) are the basic spinlock implementation currently used in the Linux kernel, rendering the code in this test case heavily deployed in production. The implementation is non-trivial, as it is based on an MCS lock, but tweaked in order to further reduce cache contention and the spinlock size (it fits in only 32 bits).

seqlock(N): Sequenced locks [9] (1.0 KLOC)

Figure 5. Comparison between GENMC-WB and MO

The mcs_spinlock benchmark is rather interesting for several reasons. First, it allows some relaxed behaviours to take place, and so NIDHUGG-PSO, GENMC-MO, and RCMC-MO explore more executions than NIDHUGG-SC and NIDHUGG-TSO (approximately 15% more). Nonetheless, GENMC and RCMC outperform NIDHUGG by a large factor. Second, GENMC-WB and RCMC-WRC11 explore *fewer* executions than GENMC-MO and RCMC-RC11, and shows the benefit of not recording **mo** in terms of verification time. Last, GENMC is slower than RCMC on this particular benchmark. This is because GENMC’s revisit procedure removes more events from the graph during backward revisits than RCMC. The extra events must then be re-added resulting in runtime overhead. Of course, this also depends on the nature of the benchmark, and the backward revisits that take place.

6.3 Modification Order vs Writes-Before

We next compare GENMC-WB and GENMC-MO more thoroughly. Admittedly, calculating **wb** for consistency is much more expensive ($O(n^3)$) than using the total order readily given by **mo**. As we show, however, (a) it can lead to exploring *exponentially fewer* executions than recording **mo**; and (b) the overhead imposed by the **wb** calculation is usually negligible.

To see (a), consider Fig. 5 (left), depicting the number of executions explored by GENMC-WB and GENMC-MO on some synthetic benchmarks. As shown, for 7 threads, GENMC-MO can visit up to 10^6 more executions than GENMC-WB, which is also reflected in the running time.

To see (b), consider Fig. 5 (right). This scatter diagram contains all 202 benchmarks that we used (including those of §6.2). With the only noticeable exception being `linuxrwlocks` (see §6.2), we can see that GENMC-WB is never much slower than GENMC-MO. On the other hand, there are many test cases where GENMC-WB is much faster than GENMC-MO.

The speedup is due to the presence of unordered concurrent writes in the program. Kokologiannakis et al. [24] argue that concurrent writes seldom appear in correct real-world programs, and our benchmarks confirm that claim.

However, there are two observations worth mentioning. First, there are real-world benchmarks (e.g., `lamport` and `mcs_spinlock`) where there is a difference (although not exponential) in the number of explored executions between GENMC-WB and GENMC-MO, and this difference is reflected in the running time. Second, while correct programs should not have concurrent unordered writes, this may happen in *incorrect* programs, and observing the difference between the `wb` and `mo` executions can be beneficial to spot such errors.

7 Conclusions and Related Work

We have presented GENMC as an effective model checking approach that is parametric in the choice of memory model and supports high-level concurrent libraries. Our approach relies on four basic assumptions about the underlying memory model: `porf`-acyclicity, extensibility, prefix-closedness, and well-blocking. In the future, we plan to investigate whether we can relax these assumptions to enable verification under hardware memory models such as Power [6] and ARM [37] (that do not satisfy `porf`-acyclicity) and library specifications such as queues [38] (that are not prefix-closed).

Amongst the verification tools handling weak memory models (MMs), the only properly MM-parametric tool is HERD [6], a memory model simulator that allows users to experiment with different consistency predicates on small “litmus test” programs. Unlike GENMC, HERD does not require models to satisfy conditions **MM1**–**MM4**, and so accepts a wider range of models than GENMC. Nevertheless, it follows the simple approach of enumerating *all* possible executions and filtering them according to the user-supplied consistency predicate, and thus is not scalable when applied to larger programs. It would be worth extending HERD to use the GENMC approach whenever the user-supplied model can be shown to satisfy conditions **MM1**–**MM4**.

As discussed in §2, several tools based on *stateless model checking* [18, 19, 34] combined with (dynamic) partial order reduction (DPOR) techniques [1, 16] have targeted specific memory models [2–4, 15, 24, 35, 42]. Unfortunately, all

of them use somewhat different ideas, making it difficult to get a model checking algorithm that is MM-parametric. Amongst these tools, the only ones enumerating plain executions (as opposed to `mo`-executions) are: TRACER [4] for the release-acquire fragment of (R)C11; DC-DPOR [12] for SC; and RCMC [24] for the WRC11 model.

GENMC follows the general design of RCMC, but uses a revisit procedure akin to that of TRACER, i.e., when in an execution graph G a write w revisits a read r , it removes from G all events that were added to G after r and are not `porf`-before w as opposed to removing only the events `porf`-after r . As a result, the completeness proof of GENMC (unlike that of RCMC) does not require “prefix-determinacy” [24, Lemma 3.9], which does not hold for the entire RC11 model: the weaker “prefix-closedness” suffices. So, while RCMC is optimal only in the absence of RMW and SC accesses, GENMC achieves optimality for the full RC11 model.

Other tools, such as CBMC [14], encode all executions of a program together with the memory model in a SAT/SMT formula and query a dedicated solver for its satisfiability [5]. This approach should in principle be able to handle models such as RC11; however, it is currently limited to SC, TSO, and PSO. The main drawback of this approach is its SAT/SMT component, which can be slow and highly unpredictable. As a result, CBMC tends to be significantly slower than NIDHUGG on relevant benchmarks [26, 30].

Another approach is *maximal causality reduction* (MCR) [20, 21], which introduces an even coarser equivalence partitioning than `porf`, based on *values* and not the places reads read-from. This approach fundamentally assumes “multi-copy atomicity” (i.e., that writes propagate simultaneously to all other processors), and thus cannot work for RC11 [27]. It does, however, work well for SC, TSO, and PSO.

Finally, unfolding-based techniques [22, 39] have obtained similar optimality results with some DPOR algorithms for SC. It remains to be seen whether they can be generalized or achieve optimality under a coarser equivalence partitioning.

Acknowledgments

We would like to thank Michael Emmi, Konstantinos Sagonas and the PLDI reviewers for their feedback. The second author was supported in part by a European Research Council (ERC) Consolidator Grant for the project “RustBelt”, under the European Union Horizon 2020 Framework Programme (grant agreement number 683289).

References

- [1] Parosh Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. 2014. Optimal dynamic partial order reduction. In *POPL 2014*. ACM, New York, NY, USA, 373–384. <https://doi.org/10.1145/2535838.2535845>
- [2] Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonardsson, and Konstantinos Sagonas. 2015. Stateless Model Checking for TSO and PSO. In *TACAS 2015 (LNCS)*, Vol. 9035.

- Springer, Berlin, Heidelberg, 353–367. https://doi.org/10.1007/978-3-662-46681-0_28
- [3] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Carl Leonardsson. 2016. Stateless Model Checking for POWER. In *CAV 2016 (LNCS)*, Vol. 9780. Springer, Berlin, Heidelberg, 134–156. https://doi.org/10.1007/978-3-319-41540-6_8
- [4] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Tuan Phong Ngo. 2018. Optimal Stateless Model Checking Under the Release-acquire Semantics. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 135 (Oct. 2018), 29 pages. <https://doi.org/10.1145/3276505>
- [5] Jade Alglave, Daniel Kroening, and Michael Tautschnig. 2013. Partial Orders for Efficient Bounded Model Checking of Concurrent Software. In *CAV 2013 (LNCS)*, Vol. 8044. Springer, Berlin, Heidelberg, 141–157. https://doi.org/10.1007/978-3-642-39799-8_9
- [6] Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory. *ACM Trans. Program. Lang. Syst.* 36, 2, Article 7 (July 2014), 74 pages. <https://doi.org/10.1145/2627752>
- [7] Stavros Aronis, Bengt Jonsson, Magnus Lång, and Konstantinos Sagonas. 2018. Optimal Dynamic Partial Order Reduction with Observers. In *TACAS (2) (LNCS)*, Vol. 10806. Springer, 229–248. https://doi.org/10.1007/978-3-319-89963-3_14
- [8] Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ Concurrency. In *POPL 2011* (Austin, Texas, USA). ACM, New York, NY, USA, 55–66. <https://doi.org/10.1145/1926385.1926394>
- [9] Hans-Juergen Boehm. 2012. Can seqlocks get along with programming language memory models?. In *MSPC 2012*. ACM, 12–20. <https://doi.org/10.1145/2247684.2247688>
- [10] Hans-Juergen Boehm and Brian Demsky. 2014. Outlawing Ghosts: Avoiding Out-of-thin-air Results. In *MSPC 2014* (Edinburgh, United Kingdom). ACM, New York, NY, USA, Article 7, 6 pages. <https://doi.org/10.1145/2618128.2618134>
- [11] Soham Chakraborty and Viktor Vafeiadis. 2019. Grounding thin-air reads with event structures. *Proc. ACM Program. Lang.* 3, POPL, Article 70 (Jan. 2019), 28 pages. <https://doi.org/10.1145/3290383>
- [12] Marek Chalupa, Krishnendu Chatterjee, Andreas Pavlogiannis, Nishant Sinha, and Kapil Vaidya. 2017. Data-centric Dynamic Partial Order Reduction. *Proc. ACM Program. Lang.* 2, POPL, Article 31 (Dec. 2017), 30 pages. <https://doi.org/10.1145/3158119>
- [13] David Chase and Yossi Lev. 2005. Dynamic circular work-stealing deque. In *SPAA 2005*. ACM, 21–28. <https://doi.org/10.1145/1073970.1073974>
- [14] Edmund M. Clarke, Daniel Kroening, and Flavio Lerda. 2004. A Tool for Checking ANSI-C Programs. In *TACAS 2004 (LNCS)*, Vol. 2988. Springer, Berlin, Heidelberg, 168–176. https://doi.org/10.1007/978-3-540-24730-2_15
- [15] Brian Demsky and Patrick Lam. 2015. SATCheck: SAT-directed Stateless Model Checking for SC and TSO. In *OOPSLA 2015* (Pittsburgh, PA, USA). ACM, New York, NY, USA, 20–36. <https://doi.org/10.1145/2814270.2814297>
- [16] Cormac Flanagan and Patrice Godefroid. 2005. Dynamic partial-order reduction for model checking software. In *POPL 2005*. ACM, New York, NY, USA, 110–121. <https://doi.org/10.1145/1040305.1040315>
- [17] Phillip B. Gibbons and Ephraim Korach. 1997. Testing Shared Memories. *SIAM J. Comput.* 26, 4 (Aug. 1997), 1208–1244. <https://doi.org/10.1137/S0097539794279614>
- [18] Patrice Godefroid. 1997. Model Checking for Programming Languages using VeriSoft. In *POPL 1997* (Paris, France). ACM, New York, NY, USA, 174–186. <https://doi.org/10.1145/263699.263717>
- [19] Patrice Godefroid. 2005. Software Model Checking: The VeriSoft Approach. *Formal Methods in System Design* 26, 2 (March 2005), 77–101. <https://doi.org/10.1007/s10703-005-1489-x>
- [20] Jeff Huang. 2015. Stateless model checking concurrent programs with maximal causality reduction. In *PLDI 2015*. ACM, New York, NY, USA, 165–174. <https://doi.org/10.1145/2737924.2737975>
- [21] Shiyu Huang and Jeff Huang. 2016. Maximal Causality Reduction for TSO and PSO. In *OOPSLA 2016*. ACM, New York, NY, USA, 447–461. <https://doi.org/10.1145/2983990.2984025>
- [22] Kari Kähkönen, Olli Saarikivi, and Keijo Heljanko. 2015. Unfolding Based Automated Testing of Multithreaded Programs. *Autom. Softw. Eng.* 22, 4 (Dec. 2015), 475–515. <https://doi.org/10.1007/s10515-014-0150-6>
- [23] Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A promising semantics for relaxed-memory concurrency. In *POPL 2017* (Paris, France). ACM, New York, NY, USA, 175–189. <https://doi.org/10.1145/3009837.3009850>
- [24] Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. 2017. Effective Stateless Model Checking for C/C++ Concurrency. *Proc. ACM Program. Lang.* 2, POPL, Article 17 (Dec. 2017), 32 pages. <https://doi.org/10.1145/3158105>
- [25] Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. 2019. Technical Appendix. <https://plv.mpi-sws.org/genmc>
- [26] Michalis Kokologiannakis and Konstantinos Sagonas. 2017. Stateless Model Checking of the Linux Kernel’s Hierarchical Read-copy-update (Tree RCU). In *SPIN 2017* (Santa Barbara, CA, USA). ACM, New York, NY, USA, 172–181. <https://doi.org/10.1145/3092282.3092287>
- [27] 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, New York, NY, USA, 618–632. <https://doi.org/10.1145/3062341.3062352>
- [28] Leslie Lamport. 1979. How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. *IEEE Trans. Computers* 28, 9 (Sept. 1979), 690–691. <https://doi.org/10.1109/TC.1979.1675439>
- [29] Leslie Lamport. 1987. A Fast Mutual Exclusion Algorithm. *ACM Trans. Comput. Syst.* 5, 1 (Jan. 1987), 1–11. <https://doi.org/10.1145/7351.7352>
- [30] L. Liang, P. E. McKenney, D. Kroening, and T. Melham. 2018. Verification of tree-based hierarchical read-copy update in the Linux kernel. In *DATE 2018*. 61–66. <https://doi.org/10.23919/DATE.2018.8341980>
- [31] Jeremy Manson, William Pugh, and Sarita V. Adve. 2005. The Java memory model. In *POPL 2005*. ACM, 378–391. <https://doi.org/10.1145/1040305.1040336>
- [32] Antoni Mazurkiewicz. 1987. Trace Theory. In *Petri nets: Applications and relationships to other models of concurrency (LNCS)*, Vol. 255. Springer, Berlin, Heidelberg, 279–324. https://doi.org/10.1007/3-540-17906-2_30
- [33] John M. Mellor-Crummey and Michael L. Scott. 1991. Algorithms for Scalable Synchronization on Shared-memory Multiprocessors. *ACM Trans. Comput. Syst.* 9, 1 (Feb. 1991), 21–65. <https://doi.org/10.1145/103727.103729>
- [34] Madanlal Musuvathi, Shaz Qadeer, Thomas Ball, Gérard Basler, P. Ramanayagam Arumuga Nainar, and Iulian Neamtii. 2008. Finding and Reproducing Heisenbugs in Concurrent Programs. In *OSDI 2008*. USENIX Association, 267–280.
- [35] Brian Norris and Brian Demsky. 2013. CDSChecker: Checking concurrent data structures written with C/C++ atomics. In *OOPSLA 2013*. ACM, 131–150. <https://doi.org/10.1145/2509136.2509514>
- [36] Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A Better x86 Memory Model: x86-TSO. In *TPHOLs 2009* (Munich, Germany). Springer, 391–407. https://doi.org/10.1007/978-3-642-03359-9_27
- [37] Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2018. Simplifying ARM concurrency: Multicopy-atomic axiomatic and operational models for ARMv8. *Proc. ACM Program. Lang.* 2, POPL (2018), 19:1–19:29. <https://doi.org/10.1145/3158107>

- [38] Azalea Raad, Marko Doko, Lovro Rožić, Ori Lahav, and Viktor Vafeiadis. 2019. On library correctness under weak memory consistency: Specifying and verifying concurrent libraries under declarative consistency models. *Proc. ACM Program. Lang.* 3, POPL (2019), 68:1–68:31. <https://doi.org/10.1145/3290381>
- [39] César Rodríguez, Marcelo Sousa, Subodh Sharma, and Daniel Kroening. 2015. Unfolding-based Partial Order Reduction. In *CONCUR 2015 (LIPIcs)*, Vol. 42. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 456–469. <https://doi.org/10.4230/LIPIcs.CONCUR.2015.456>
- [40] SPARC International Inc. 1994. *The SPARC architecture manual (version 9)*. Prentice-Hall.
- [41] Viktor Vafeiadis and Chinmay Narayan. 2013. Relaxed Separation Logic: A program logic for C11 concurrency. In *OOPSLA 2013*. ACM, New York, NY, USA, 867–884. <https://doi.org/10.1145/2509136.2509532>
- [42] Naling Zhang, Markus Kusano, and Chao Wang. 2015. Dynamic partial order reduction for relaxed memory models. In *PLDI 2015*. ACM, New York, NY, USA, 250–259. <https://doi.org/10.1145/2737924.2737956>

A Memory Model Instantiations in Our Framework

A.1 The SC Memory Model

An execution G is *SC-consistent* iff there exists an order $\text{mo} \triangleq \bigcup_{x \in \text{Loc}} \text{mo}_x$ such that:

- each $\text{mo}_x \subseteq G.W_x \times G.W_x$ is a strict total order on $G.W_x$;
- $G.\text{rf}$ is a total function and for all $\langle w, r \rangle \in G.\text{rf}$, $\text{val}_w(w) = \text{val}_r(r)$, where $\text{val}_w(w)$ denotes the value written by w and $\text{val}_r(r)$ denotes the value read by r ;
- $(G.\text{po} \cup G.\text{rf} \cup \text{mo} \cup (G.\text{rf}^{-1}; \text{mo}))^+$ is irreflexive.

A.2 The TSO Memory Model

Definition A.1 (TSO). An execution G is *TSO-consistent* iff there exists an order $\text{tso} \subseteq G.E \times G.E$ such that:

- tso is total on $G.W$
- $G.\text{po} \setminus (G.W \times G.R) \subseteq \text{tso}$
- $G.\text{rf} \subseteq \text{tso} \cup G.\text{po}$
- $\forall \langle w, r \rangle \in G.\text{rf}. \forall w' \in G.W \cap G.R. \langle w', r \rangle \in \text{tso} \cup G.\text{po} \wedge \text{loc}(w') = \text{loc}(r) \Rightarrow \langle w, w' \rangle \notin \text{tso}$
- $G.\text{rf}$ is a total function and for all $\langle w, r \rangle \in G.\text{rf}$, $\text{val}_w(w) = \text{val}_r(r)$, where $\text{val}_w(w)$ denotes the value written by w and $\text{val}_r(r)$ denotes the value read by r .

A.3 The RC11 Memory Model

First, let us define the *happens-before* (hb) order for RC11. Intuitively, hb records when an event is globally perceived as occurring before another one. To this end, several derived relations are needed:

$$\begin{aligned} G.\text{rseq} &\triangleq \bigcup_{x \in \text{Loc}} [W_x]; G.\text{po}^?; [W_x^{\exists r1x}]; (G.\text{rf}; [R \cap W])^* && (\text{release sequence}) \\ G.\text{sw} &\triangleq [E^{\exists \text{rel}}]; ([F]; G.\text{po})^?; G.\text{rseq}; G.\text{rf}; [R^{\exists r1x}]; (G.\text{po}; [F])^?; [E^{\exists \text{acq}}] && (\text{synchronizes with}) \\ G.\text{hb} &\triangleq (G.\text{po} \cup G.\text{sw})^+ && (\text{happens-before}) \end{aligned}$$

Happens-before is defined in terms of two more basic definitions:

1. The *release sequence* (rseq) relation: the release sequence of a write contains the write itself, all later writes to the same location in the same thread, as well as all RMWs that recursively read from such writes.
2. The *synchronizes with* (sw) relation: we say that a release event a *synchronizes with* an acquire event b , whenever b (or, in case b is a fence, some po-prior read) reads from the release sequence of a (or, in case a is a fence, of some po-later write).

Then, we say that an event a *happens-before* an event b if there is a path from a to b consisting of po and sw edges.

With this definition for hb , an RC11-consistent execution must satisfy the requirements described in the following subsections.

A.3.1 Completeness. The first constraint is very simple: every read should read a value from a write with the same value. Accordingly, an execution G is said to be *complete* if $G.\text{rf}$ is a total function and for all $\langle w, r \rangle \in G.\text{rf}$, $\text{val}_w(w) = \text{val}_r(r)$, where $\text{val}_w(w)$ denotes the value written by w and $\text{val}_r(r)$ denotes the value read by r .

A.3.2 Coherence. RC11 requires that there exist a total order called *modification order* (mo) that is a disjoint union of relations $\{\text{mo}_x\}_{x \in \text{Loc}}$, such that each mo_x is a strict partial order on $E \cap W_x$.

Then, coherence (a.k.a. SC-per-location) requires two things:

1. For every particular location, all threads agree on the order of accesses to that location, and
2. this order should be consistent with the *happens-before* order (hb).

To order accesses to a given location, the model requires that for every location x , $G.\text{mo}$ *totally* orders the writes to x , and defines an extension of mo , which is a partial order on *all* accesses to x :

$$G.\text{eco} \triangleq (G.\text{mo} \cup G.\text{rf} \cup G.\text{rf}^{-1}; G.\text{mo})^+ \quad (\text{extended coherence order})$$

Here, writes are ordered using mo , while reads are placed after the writes they read from, but before writes that are mo -later than the writes they read from.

Then, the coherence condition simply requires that $G.\text{eco}; G.\text{hb}$ is irreflexive.

A.3.3 Atomicity. There must be no two RMW events reading from the same write. Alternatively, there must be no event in modification order between the write from which an RMW u reads-from and the u itself.

Formally, we can demand that $G.\text{rf}; [R \cap W]; G.\text{mo}^{-1}; G.\text{mo}^{-1}$ is irreflexive.

A.3.4 Global SC Constraint. SC accesses and fences are subject to a global constraint, which, roughly speaking, requires threads to agree on their order. In fact, due to the interaction with other access modes, this is notably the most involved part of RC11, which addresses flaws of the original C/C++ memory model.

RC11 requires the acyclicity of a relation called *partial SC order*, denoted psc , which is, in turn, defined using additional helper notations:

$$\begin{aligned}
 G.\text{po}|_{\neq \text{loc}} &\triangleq \{\langle a, b \rangle \in G.\text{po} \mid \text{loc}(a) \neq \text{loc}(b)\} & G.\text{hb}|_{\text{loc}} &\triangleq \{\langle a, b \rangle \in G.\text{hb} \mid \text{loc}(a) = \text{loc}(b)\} \\
 G.\text{scb} &\triangleq G.\text{po} \cup G.\text{po}|_{\neq \text{loc}}; G.\text{hb}; G.\text{po}|_{\neq \text{loc}} \cup G.\text{hb}|_{\text{loc}} \cup G.\text{mo} \cup G.\text{rf}^{-1}; G.\text{mo} & (SC\text{-before}) \\
 G.\text{psc} &\triangleq ([E^{\text{sc}}] \cup [F^{\text{sc}}]; G.\text{hb}^?); G.\text{scb}; ([E^{\text{sc}}] \cup G.\text{hb}^?; [F^{\text{sc}}]) \cup \\
 & [F^{\text{sc}}]; (G.\text{hb} \cup G.\text{hb}; G.\text{eco}; G.\text{hb}); [F^{\text{sc}}] & (partial\ SC\ order)
 \end{aligned}$$

A.3.5 No porf cycles. In order to rule out “out-of-thin-air” behaviours, RC11 requires that porf be acyclic.

B The Writes-Before Relation

Given the RC11 memory model in §A.3, the writes-Before (wb) relation is an alternative coherence definition that ensures the existence of modification orders, without actually providing them. For a memory location x , wb_x can be defined as follows:

$$\text{wb}_x \triangleq [G.W_x]; \left((G.\text{rf}^*; G.\text{hb}; (G.\text{rf}^{-1})^?; G.\text{rf}^*) \setminus (G.\text{rf}^{-1})^* \right); [G.W_x]$$

C Tracking the Modification Order

```

1: procedure VERIFY( $P$ )
2:    $\langle G, T, U, S \rangle \leftarrow \langle G_0, \emptyset, \emptyset, \emptyset \rangle$ 
3:   VISITONE( $P, G, T, U, S$ )
4:   while  $el \leftarrow \text{RemoveMax}(S)$  do
5:     if  $el = \langle r, w, G_w \rangle$  then
6:        $\langle E_1, r, E_2 \rangle \leftarrow \text{split}(G.E, r)$ 
7:        $G \leftarrow \langle E_1 \uplus [r] \uplus G_w.E, G.\text{rf}|_{E_1} \cup G_w.\text{rf}, G.\text{mo}|_{E_1} \cup G_w.\text{mo} \rangle$ 
8:        $G.\text{rf}[r] \leftarrow w$ 
9:        $T \leftarrow T \cap (E_1 \uplus [r])$ 
10:       $U \leftarrow U \setminus \{U[e'] \mid e' \in E_2\}$ 
11:      if  $w \neq \perp$  then CALCREVISITS( $G, T, U, S, r$ )
12:    if  $el = \langle w, M \rangle$  then
13:       $\langle E_1, w, E_2 \rangle \leftarrow \text{split}(G.E, w)$ 
14:       $G \leftarrow \langle E_1 \uplus [r], G.\text{rf}|_{E_1}, G.\text{mo}|_{E_1} \rangle$ 
15:       $G.\text{mo}[\text{loc}(w)] \leftarrow M \times \{w\} \cup \{w\} \times (E_1 \cap W_{\text{loc}(w)} \setminus M)$ 
16:       $T \leftarrow T \cap (E_1 \uplus [w])$ 
17:       $U \leftarrow U \setminus \{U[e'] \mid e' \in E_2\}$ 
18:      CALCREVISITS( $G, T, U, S, w$ )
19:    VISITONE( $P, G, T, U, S$ )

1: procedure VISITONE( $G, T, U, S$ )
2:   while  $\text{cons}(G) \wedge a \leftarrow \text{next}_P(G)$  do
3:     if  $a \in \text{error}$  then exit("erroneous program")
4:      $G \leftarrow \text{Add}(G, a)$ 
5:     if  $a \in R$  then
6:        $W \leftarrow G.E \cap W_{\text{loc}(a)}$ 
7:        $B \leftarrow \{\langle \perp, \emptyset \rangle \mid \text{cons}(G.\text{rf}[a \mapsto \perp])\}$ 
8:       choose some  $w_0 \in W$ 
9:        $G.\text{rf}[r] \leftarrow w_0$ 
10:       $T \leftarrow T \cup \{r\}$ 
11:       $S[a] \leftarrow S[a] \cup \{\langle w, \emptyset \rangle \mid w \in W \setminus \{w_0\}\} \cup B$ 
12:       $U[a] \leftarrow U[a] \cup B$ 
13:    if  $a \in W$  then
14:       $G.\text{mo} \leftarrow G.\text{mo} \cup (G.E \cap W_{\text{loc}(a)} \times \{a\})$ 
15:       $S[a] \leftarrow S[a] \cup \{M \mid M \subseteq G.E \cap W_{\text{loc}(a)} \wedge \text{rng}([M]; G.\text{mo}) \subseteq M\}$ 
16:    CALCREVISITS( $G, T, U, S, a$ )

1: procedure CALCREVISITS( $G, T, U, S, a$ )
2:    $p_a \leftarrow \text{dom}(G.\text{porf}^?; [a])$ 
3:   for  $r \in T \cap R_{\text{loc}(a)} \setminus p_a$  do
4:      $\langle E_1, r, E_2 \rangle \leftarrow \text{split}(G.E, r);$ 
5:      $G_a \leftarrow G|_{E_2 \cap p_a}$ 
6:     if  $a \notin W$  then  $a \leftarrow \perp$ 
7:     if  $\langle a, G_a \rangle \notin U[r] \wedge \langle a, \emptyset \rangle \notin U[r]$  then
8:        $S[r] \leftarrow S[r] \cup \{\langle a, G_a \rangle\}$ 
9:        $U[r] \leftarrow U[r] \cup \{\langle a, G_a \rangle\}$ 

```

D Completeness and Optimality Proofs

In what follows, we write $G.\text{pred}$ for the strict total order induced by the $G.E$ sequence. That is, $\langle a, b \rangle \in G.\text{pred} \iff a <_{G.E} b$, where $a <_{G.E} b$ denotes that a appears before b in $G.E$.

D.1 Completeness Proof

Proposition D.1 (Read-extensibility). *Given execution G such that $\text{cons}(G)$ holds, and event $r = \langle i, |G.E_i| + 1, l \rangle$ with $\text{typ}(r) = R$, there exists $w \in G.W_{\text{loc}(r)}$ such that $\text{cons}(\text{Add}(G, r).\text{rf}[r \mapsto w])$.*

Proposition D.2 (Write-extensibility). *Given execution G and event w , with $\text{typ}(w) = W$ and $\text{rng}([w]; \text{porf}) = \emptyset$, if $\text{cons}(G \setminus \{w\})$ holds, then so does $\text{cons}(G)$.*

Proposition D.3 (Update-extensibility). *Given execution G such that $\text{cons}(G)$ holds, and event u with $\text{typ}(u) = \text{RMW}$ and $\text{rng}([u]; G.\text{po}) = \emptyset$, then the following also hold:*

- (a) *Given event r such that:*
 - $\text{typ}(r) = R$
 - $\langle u, r \rangle \in G.\text{rf}$
 - $\text{rng}([r]; G.\text{po}) = \emptyset$*then $\exists w \in G.E \setminus \{u\}$ such that $\text{cons}(G.\text{rf}[r \mapsto w])$.*
- (b) *Given events u', w such that:*
 - $\text{typ}(u') = \text{RMW}$ and $\text{typ}(w) \in \{W, \text{RMW}\}$
 - $\langle u, u' \rangle, \langle w, u \rangle \in G.\text{rf}$
 - $\text{rng}([u']; G.\text{porf}) = \emptyset$*then $\text{cons}(G \setminus \{u\}.\text{rf}[u' \mapsto w])$.*

Definition D.4 (Prefix construction). Given an execution $G_f \supseteq G_0$, an event e is added to G_0 in porf-order , written $G_0 \xrightarrow{e} G_1$, if $e \in G_f \setminus G_0$, $\text{dom}(G_f.\text{porf}; [e]) \subseteq G_0$, and $G_1 \setminus \{e\} = G_0$.

A graph G is constructed in porf-order if $\emptyset \rightarrow^* G$, where $\rightarrow^*$ denotes the reflexive-transitive closure of $\rightarrow$ (i.e., G is constructed via a succession of porf-steps).

Definition D.5 (Algorithmic construction). Given an execution G_0 , an event e is added to G_0 in *algorithmic order*, written $G_0 \xrightarrow{e} G_1$, if there exists T, S, U, T', S', U' such that Algorithm 2 calls VISITREAD (resp. VISITWRITE) from a configuration $\langle G_0, T, S, U \rangle$, and a configuration $\langle G_1, T', S', U' \rangle$ is produced either immediately after the call, or when REVISITREADS pops one of the entries pushed to S by VISITREAD (resp. VISITWRITE). Note that, in either case, it will be $G_1 \setminus \{e\} = G_0$.

Given a graph G_0 , an event e is added in algorithmic order *with no revisits*, written $G_0 \xrightarrow{e}_{nr} G_1$, if $G_0 \xrightarrow{e} G_1$ without revisiting any events along the way.

A graph G is constructed in *algorithmic order* if $\emptyset \Rightarrow^* G$, where $\Rightarrow^*$ denotes the reflexive-transitive closure of $\Rightarrow$ (i.e., there is a run of Algorithm 2 that produces a configuration $\langle G, -, -, - \rangle$).

Lastly, given an algorithmic construction $G_0 \Rightarrow^* G_n$, a graph G_i is part of this construction, written $G_i \in [G_0 \Rightarrow^* G_n]$, if it is $G_0 \Rightarrow^* G_i \Rightarrow^* G_n$.

Theorem D.6 (Completeness). *For any $G_f \in \llbracket P \rrbracket$ such that $\text{cons}(G_f)$ it is $\emptyset \Rightarrow^* G_f$.*

ASSUME: 1. P is a deterministic program with finite traces

2. $G_f \in \llbracket P \rrbracket$

3. $\text{cons}(G_f)$

PROVE: $\emptyset \Rightarrow^* G_f$

(1)1. $\emptyset \rightarrow^* G_f$

PROOF: Can be shown trivially by induction on the length of G_f .

(1)2. Q.E.D.

PROOF: From Lemma D.7, by taking $G = G_f = G^{\text{ext}}$.

□

Lemma D.7 (Prefix extension). *If $\emptyset \rightarrow^* G \Rightarrow^* G^{\text{ext}}$ with $\text{cons}(G)$ and $G^{\text{ext}} \in \llbracket P \rrbracket$ for some program P , then $\emptyset \Rightarrow^* G^{\text{ext}}$.*

PROOF SKETCH: Proof by induction: we assume that we have a fixed prefix construction G from which the algorithm can reach G^{ext} , and prove that the algorithm can reach the same extension running from the initial configuration.

For the proof of the inductive step, given a G such that $\emptyset \rightarrow^* G \xrightarrow{e} G'$, and a graph G'^{ext} such that $G' \Rightarrow^* G'^{ext}$, we show that it is also $G \Rightarrow^* G'^{ext}$, and then use the inductive hypothesis to prove that $\emptyset \Rightarrow^* G'^{ext}$.

ASSUME: 1. P is a deterministic program with finite traces

2. $\emptyset \rightarrow^* G$ with $\text{cons}(G)$

3. $G \Rightarrow^* G^{ext}$ with $G^{ext} \in \llbracket P \rrbracket$

PROVE: $\emptyset \Rightarrow^* G^{ext}$

$\langle 1 \rangle 1$. CASE: $G = \emptyset$

PROOF: Trivially holds from assumption 3 and case assumption $\langle 1 \rangle$.

$\langle 1 \rangle 2$. CASE: 1. $\emptyset \rightarrow^* G \xrightarrow{e} G'$

2. IH: For any G^{ext} such that $G \Rightarrow^* G^{ext}$, it is $\emptyset \Rightarrow^* G^{ext}$

PROOF: From D.8 by taking $G_0 = G$, $G_f = G^{ext}$, $G = G$ and $G' = G'$ we get that $G \Rightarrow^* G_f$. Then, we get the desired result from the IH.

$\langle 1 \rangle 3$. Q.E.D.

PROOF: Steps $\langle 1 \rangle 1$, $\langle 1 \rangle 2$, assumptions 1-3, and mathematical induction. □

Proposition D.8 (Commutativity). *If $G_0 \Rightarrow^* G \xrightarrow{e} G' \Rightarrow^* G_f$ for some G_0, G_f , with $\text{cons}(G_0), \text{cons}(G_f)$, and $G_f \in \llbracket P \rrbracket$, it is $G_0 \Rightarrow^* G_f$.*

PROOF SKETCH: We show that running the algorithm from G will also lead to G_f , thus proving that $G_0 \Rightarrow^* G_f$. For that, we do proof by induction on the length of the algorithmic construction $G' \Rightarrow^* G_f$.

For the proof of the inductive step, we show that the algorithmic steps taken when starting from G either commute with e , and then use the inductive hypothesis for a production of a smaller length, or that they will result into a graph that is also produced by the algorithm when starting from G' .

ASSUME: 1. P is a deterministic program with finite traces

2. $G_0 \Rightarrow^* G$

3. $G \xrightarrow{e} G'$

4. $G' \Rightarrow^* G_f$

PROVE: $G_0 \Rightarrow^* G_f$

$\langle 1 \rangle 1$. CASE: $G' = G_f$

PROOF: Algorithm 2 can simply take the step; it is $G \xRightarrow{e} G'$. From case hypothesis $\langle 1 \rangle$ and assumption $\langle 0 \rangle 4$ we get that $G_0 \Rightarrow^* G_f$.

$\langle 1 \rangle 2$. CASE: 1. $G' \xRightarrow{a} G_2 \Rightarrow^* G_f$

2. IH: For any G_3 such that $G_0 \Rightarrow^* G_3 \xrightarrow{e} G'_3 \Rightarrow^* G_f$, it is $G_0 \Rightarrow^* G_f$

$\langle 2 \rangle 1$. CASE: $\text{next}_P(G) = e$

PROOF: It is $G \xRightarrow{e} G'$ from graph extensibility. Thus, from $\langle 1 \rangle 1$ we get that $G_0 \Rightarrow^* G \xRightarrow{e} G' \xRightarrow{a} G_2 \Rightarrow^* G_f$.

$\langle 2 \rangle 2$. CASE: $\text{next}_P(G) = a$ with $a \neq e$

$\langle 3 \rangle 1$. CASE: $\langle e, a \rangle \notin G_2.\text{rf}$ and a performs no revisiting

$\langle 4 \rangle 1$. There exists G'' such that $G \xRightarrow{a} G'' \xrightarrow{e} G_2$

PROOF: Graph extensibility guarantees that a can be added in a consistent manner in G' and, in addition, in the case where $\text{typ}(a) \in \{R, \text{RMW}\}$, Def. 3.4 with assumption $\langle 3 \rangle$ guarantees that $\exists w \in G'', G_2$ such that $\langle w, a \rangle \in G''.\text{rf}$ and $\langle w, a \rangle \in G_2.\text{rf}$ (i.e., $G'' = G_2 \setminus \{e\}$). From case assumptions $\langle 1 \rangle 1$ and $\langle 3 \rangle$ we get that it is $G'' \xrightarrow{e} G_2$, i.e., the two transitions commute: both $G \xRightarrow{a} G'' \xrightarrow{e} G_2$ and $G \xrightarrow{e} G' \xRightarrow{a} G_2$ hold.

$\langle 4 \rangle 2$. Q.E.D.

PROOF: From $\langle 4 \rangle 1$ and IH (assumption $\langle 1 \rangle 2$).

$\langle 3 \rangle 2$. CASE: $\langle e, a \rangle \in G_2.\text{rf}$

$\langle 4 \rangle 1$. There is a run of Algorithm 2 such that $G \xRightarrow{a}_{nr} G'' \Rightarrow_{nr}^* G_1 \xrightarrow{e} G_2$

PROOF: From Prop. D.9, case assumptions $\langle 2 \rangle$, $\langle 1 \rangle 1-2$, and $\langle 3 \rangle$.

$\langle 4 \rangle 2$. Q.E.D.

From $\langle 4 \rangle 1$, assumption $\langle 0 \rangle 2$ and case assumption $\langle 1 \rangle 1$.

$\langle 3 \rangle 3$. CASE: a revisits some event $r_0 \neq e$ and $\langle e, a \rangle \notin G_2.\text{rf}$

⟨4⟩1. There exists G'_2 such that $G \xRightarrow{a}_{nr} G'_2$

PROOF: From graph extensibility (Prop. D.2). (Intuitively, this is the graph in which Algorithm 3 will revisit all consistent reads.)

⟨4⟩2. $e \in \text{rng}([r_0]; G'_2.\text{pred})$

From assumptions ⟨0⟩:3 and ⟨1⟩:1.

⟨4⟩3. ASSUME: $e \in \text{dom}(G'_2.\text{porf}^?; [a])$

PROVE: FALSE

PROOF: From case assumption ⟨3⟩ it is $\text{typ}(a) = W$. Then, case assumption ⟨4⟩ implies that it must be $a \in \text{rng}([e]; G'_2.\text{porf}^?; G'_2.\text{po})$. However, this in turn means that $\text{rng}([e]; G'_2.\text{porf}^?) \neq \emptyset$, which contradicts assumption ⟨1⟩:1.

⟨4⟩4. $e \notin G_2$

PROOF: When a revisit takes place for all e' such that $e' \in \text{rng}([r_0]; G'_2.\text{pred})$ and $e' \notin \text{dom}(G'_2.\text{porf}^?; [a])$ it is $e' \notin G_2$ (Algorithm 3 and §4.3). Since for e it also is $e \in \text{rng}([r_0]; G'_2.\text{pred})$ and $e \notin \text{dom}(G'_2.\text{porf}^?; [a])$ (steps ⟨4⟩2 and ⟨4⟩3), it is $e \notin G_2$.

⟨4⟩5. Q.E.D.

PROOF: From step ⟨4⟩4 we get that it is $G' \xRightarrow{a} G_2$. Then, from case assumption ⟨1⟩:1 and IH (assumption ⟨1⟩:2), we get the desired proof.

⟨3⟩4. CASE: a revisits e

⟨4⟩1. There exists G'' such that $G' \xRightarrow{a} G'' \xrightarrow{e} G_2$

PROOF: Note that, since a revisits e , it must be $\text{typ}(a) \in \{W, \text{RMW}\}$. However, from prefix construction (assumption ⟨0⟩:3) we get that $\text{typ}(a) = W$. Thus, from Prop. D.2, we get that $G' \xRightarrow{a} G''$. Then, from case assumptions ⟨1⟩:1 and ⟨3⟩, we get that the two transitions commute, i.e., both $G' \xRightarrow{a} G'' \xrightarrow{e} G_2$ and $G' \xrightarrow{e} G \xRightarrow{a} G_2$ hold.

⟨4⟩2. Q.E.D.

PROOF: From step ⟨4⟩1 and IH (assumption ⟨1⟩:2).

⟨3⟩5. Q.E.D.

PROOF: Cases ⟨3⟩1–⟨3⟩4 are exhaustive.

⟨2⟩3. Q.E.D.

PROOF: Cases ⟨2⟩1 and ⟨2⟩2 are exhaustive.

⟨1⟩3. Q.E.D.

PROOF: Steps ⟨1⟩1, ⟨1⟩2, and mathematical induction.

□

Proposition D.9 (Revisitable extension). *If $\emptyset \rightarrow^* G \xrightarrow{e} G' \xRightarrow{a} G_2 \Rightarrow^* G_f$, with $a <_G e$ and $\langle e, a \rangle \in G_2.\text{rf}$, then there is a run of Algorithm 2 from G such that $G \xRightarrow{a} G'' \Rightarrow_{nr}^* G_1 \xRightarrow{e} G_2$.*

PROOF SKETCH: We basically want to prove that when event e is added by the algorithm starting from G , event a will be considered as an option to revisit. Thus, we perform a case analysis on the type of e and show that there is a run where Algorithm 3 will not discard a from the set of revisitable options.

ASSUME: 1. P is a deterministic program with finite traces

2. $\emptyset \rightarrow^* G \xrightarrow{e} G' \xRightarrow{a} G_2 \Rightarrow^* G_f$

3. $a <_G e$

4. $\langle e, a \rangle \in G_2.\text{rf}$

PROVE: There exists a run of Algorithm 2 from G such that $G \xRightarrow{a}_{nr} G'' \Rightarrow_{nr}^* G_1 \xRightarrow{e} G_2$

⟨1⟩1. It is $\text{next}_P(G) = a$

PROOF: From assumptions ⟨0⟩:2 and ⟨0⟩:3.

⟨1⟩2. There is at least one run $G \xRightarrow{a}_{nr} G'' \Rightarrow_{nr}^* G_1$ with $\text{next}_P(G'_1) = e$

PROOF: Graph extensibility guarantees that Algorithm 2 will be able to add all events a' with $a' <_G e$, resulting to a graph G_0 , with no revisits performed. Since Algorithm 2 performs no revisiting, and all events e' such that $e' \in \text{dom}(G_2.\text{porf}^?; [e])$ are already in G (prefix construction; assumption ⟨0⟩:2), it will be $\text{next}_P(G_0) = e$.

⟨1⟩3. CASE: $\text{typ}(e) = W$

⟨2⟩1. Choose graphs G'_1 and G_p such that

1. $G \xRightarrow{a}_{nr} G_1 \Rightarrow_{nr}^* G'_1$

$$2. G'_1 \xRightarrow{e}_{nr} G_p$$

PROOF: Step $\langle 1 \rangle 2$ and graph extensibility (Prop. D.2) guarantee that we can select appropriate G'_1 and G_p .

$$\langle 2 \rangle 2. \text{ It is } a \in T \text{ when } G'_1 \text{ is reached in the production } G \xRightarrow{a}_{nr} G_1 \Rightarrow_{nr}^* G'_1$$

PROOF: From step $\langle 2 \rangle 1$ we know that no revisiting is performed, and only if an event gets revisited it is removed from T .

$$\langle 2 \rangle 3. \text{ It is } \text{cons}(G_2)$$

PROOF: By assumption $\langle 0 \rangle 2$.

$$\langle 2 \rangle 4. \text{ ASSUME: } \langle a, e \rangle \in G_p.\text{porf}$$

PROVE: FALSE

$$\langle 3 \rangle 1. \text{ It is } e \in \text{rng}([a]; G_p.\text{porf}^?; G_p.\text{po})$$

PROOF: This implies that revisiting is performed, which contradicts case assumption $\langle 2 \rangle 1$.

$$\langle 3 \rangle 2. \text{ It is } e \in \text{rng}([a]; G_p.\text{porf}^?; G_p.\text{rf})$$

This implies that $e \in \text{rng}([a]; G_p.\text{rf})$, which contradicts case assumption $\langle 1 \rangle$.

$$\langle 3 \rangle 3. \text{ Q.E.D.}$$

PROOF: From assumption $\langle 2 \rangle$ we get that cases $\langle 3 \rangle 1$ and $\langle 3 \rangle 2$ are exhaustive.

$$\langle 2 \rangle 5. \text{ Q.E.D.}$$

It will be $G'_1 \xRightarrow{e} G_2$, because a will be considered as an event to revisit by Algorithm 3, and will not be discarded (steps $\langle 2 \rangle 2$, $\langle 2 \rangle 3$, $\langle 2 \rangle 4$).

$$\langle 1 \rangle 4. \text{ CASE: } \text{typ}(e) = \text{RMW}$$

$$\langle 2 \rangle 1. \text{ Let } w \in G.E \text{ such that } \langle w, e \rangle \in G'.\text{rf}$$

PROOF: Such a w exists in G due to the prefix construction (assumption $\langle 0 \rangle 2$).

$$\langle 2 \rangle 2. \text{ There exists } m \in \mathbb{N} \text{ and } G_0, G_3, G_4 \text{ such that}$$

$$1. G_0 \Rightarrow_{nr}^m G_3$$

$$2. \text{tid}(\text{next}_P(G_3)) = \text{tid}(e)$$

$$3. \text{ One of the following holds}$$

$$\text{a. } G_0 = G' \text{ and } \nexists r \in G_3.\text{rf} \text{ such that } \langle e, r \rangle \in G_3.\text{rf}$$

$$\text{b. } G_0 = G' \setminus \{e\} \text{ and } e \notin G_3.E \text{ and } \exists u' \in G_3.E \text{ such that } \langle w, u' \rangle \in G_3.\text{rf}$$

PROOF: From assumption $\langle 0 \rangle 2$, case assumption $\langle 1 \rangle$, step $\langle 2 \rangle 1$, and by choosing an appropriate n (which is always possible due to program finiteness), Prop. D.10 yields the required result.

$$\langle 2 \rangle 3. \text{ Choose } G'_1 \text{ such that } G'_1 = G_3 \setminus \{e\}, G \Rightarrow_{nr}^m G'_1, \text{ and one of the following holds:}$$

$$1. \nexists r \in G'_1.E \text{ such that } \langle e, r \rangle \in G'_1.\text{rf}$$

$$2. \exists u' \in G'_1.E \text{ such that } \langle w, u' \rangle \in G'_1.\text{rf}$$

$$\langle 3 \rangle 1. \text{ Choose } G'_1 \text{ such that } G'_1 = G_3 \setminus \{e\}$$

PROOF: From Def. 3.4 and $\langle 2 \rangle 2.3$.

$$\langle 3 \rangle 2. \text{ CASE: } G_0 = G' \text{ and } \nexists r \in G_3.E \text{ such that } \langle e, r \rangle \in G_3.\text{rf}$$

$$\langle 4 \rangle 1. \text{ It is } G_0 \setminus \{e\} \Rightarrow_{nr}^m G'_1 \text{ and } \nexists r \in G'_1.E \text{ such that } \langle e, r \rangle \in G'_1.\text{rf}$$

PROOF: From step $\langle 2 \rangle 2.1$ and case assumption $\langle 3 \rangle$ we know that we can use Lemma D.11.

$$\langle 4 \rangle 2. \text{ Q.E.D.}$$

PROOF: From step $\langle 4 \rangle 1$, along with assumption $\langle 0 \rangle 2$ and case assumption $\langle 3 \rangle$.

$$\langle 3 \rangle 3. \text{ CASE: } G_0 = G' \setminus \{e\} \text{ and } e \notin G_3.E \text{ and } \exists u' \in G_3.E \text{ such that } \langle w, u' \rangle \in G_3.\text{rf}$$

PROOF: By choosing $G'_1 = G_3$, since $e \notin G_3$ and no revisiting is performed (assumptions $\langle 3 \rangle$ and $\langle 2 \rangle 2.1$), and also $G = G_0$ (assumptions $\langle 0 \rangle 2$ and $\langle 3 \rangle$), it trivially is $G \Rightarrow_{nr}^m G'_1$.

$$\langle 3 \rangle 4. \text{ Q.E.D.}$$

PROOF: From $\langle 2 \rangle 2.3$ we get that cases $\langle 3 \rangle 2$ and $\langle 3 \rangle 3$ are exhaustive.

$$\langle 2 \rangle 4. \text{ There is at least one } G_p \text{ such that } G'_1 \xRightarrow{e}_{nr} G_p$$

PROOF: From Prop. D.1 and steps $\langle 2 \rangle 3$ and $\langle 2 \rangle 2.2$.

$$\langle 2 \rangle 5. \text{ It is } a \in T \text{ when } G'_1 \text{ is reached in the production } G \xRightarrow{a}_{nr} G_1 \Rightarrow_{nr}^* G'_1$$

PROOF: From $\langle 2 \rangle 3$ and $\langle 1 \rangle 1$. Note that, if there are more than one alternatives for the first step of the production

$$G \xRightarrow{a}_{nr} G_1 \Rightarrow_{nr}^{m-1} G'_1, \text{ it will be } a \in T \text{ in one of them.}$$

$$\langle 2 \rangle 6. \text{ It is } \text{cons}(G_2)$$

PROOF: From assumption $\langle 0 \rangle 2$.

$$\langle 2 \rangle 7. \text{ We can choose } G_p \text{ such that either } \langle a, e \rangle \notin G_p.\text{porf} \text{ or } w \in W_{\text{excl}} \text{ when } e \text{ is added}$$

$$\langle 3 \rangle 1. \text{ CASE: } \nexists r \in G'_1.\text{rf} \text{ such that } \langle e, r \rangle \in G'_1.\text{rf}$$

$\langle 4 \rangle 1$. There exists G_p such that $\langle w, e \rangle \in G_p.\text{rf}$

PROOF: From steps $\langle 2 \rangle 2$ and $\langle 2 \rangle 4$, as well as case assumption $\langle 3 \rangle$, we can choose $G_p = G_3$.

$\langle 4 \rangle 2$. Q.E.D.

PROOF: From case assumption $\langle 3 \rangle$, and steps $\langle 2 \rangle 4$ and $\langle 4 \rangle 1$, it cannot be $\langle a, e \rangle \in G_p.\text{porf}$.

$\langle 3 \rangle 2$. CASE: $e \notin G'_1.E$ and $\exists u' \in G'_1.E$ such that $\langle w, u' \rangle \in G'_1.\text{rf}$

PROOF: From case assumption $\langle 3 \rangle$ we get that it will be $w \in W_{\text{excl}}$ for any G_p .

$\langle 3 \rangle 3$. Q.E.D.

PROOF: From $\langle 2 \rangle 3$ we get that cases $\langle 3 \rangle 1$ and $\langle 3 \rangle 2$ are exhaustive.

$\langle 2 \rangle 8$. Q.E.D.

$\langle 1 \rangle 5$. Q.E.D.

PROOF: From $\langle 0 \rangle 4$ we conclude that cases $\langle 1 \rangle 3$ and $\langle 1 \rangle 4$ are exhaustive.

□

Proposition D.10 (Non-sbrf extension). *If $\emptyset \rightarrow^* G \xrightarrow{e} G_1$, with $\text{typ}(e) = \text{RMW}$ and $\langle w, e \rangle \in G_1.\text{rf}$ for some $w \in G_1.E$, then $\forall n \in \mathbb{N}$ such that $G_1 \Rightarrow_{nr}^n G'_1$ it is $\exists m, G_0, G_2, G_3$ such that*

1. $G_0 \Rightarrow_{nr}^m G_2$

2. One of the following holds

a. $m = n$

b. $m < n$ and exists $a \in E$ such that $G_2 \xRightarrow{a}_{nr} G_3 \wedge \text{tid}(a) = \text{tid}(e)$

3. One of the following holds

a. $G_0 = G_1$ and $\nexists r \in G_2.E$ such that $\langle e, r \rangle \in G_2.\text{rf}$

b. $G_0 = G_1 \setminus \{e\}$ and $e \notin G_2.E$ and $\exists u' \in G_2.E$ such that $\langle w, u' \rangle \in G_2.\text{rf}$

PROOF SKETCH: We prove the argument inductively and use the graph extensibility properties to guarantee that the claims hold at each step.

ASSUME: 1. P is a deterministic and finite program

2. $\emptyset \rightarrow^* G \xrightarrow{e} G_1$

3. $\text{typ}(e) = \text{RMW}$

4. $\langle w, e \rangle \in G_1.\text{rf}$ for some $w \in G_1.E$

5. There exists $n \in \mathbb{N}$ such that $G_1 \Rightarrow_{nr}^n G'_1$

PROVE: There exist m'', G''_0, G''_2, G''_3 such that

1. $G''_0 \Rightarrow_{nr}^m G''_2$

2. One of the following holds

a. $m'' = n$

b. $m'' < n$ and exists $a'' \in E$ such that $G_2 \xRightarrow{a''}_{nr} G_3 \wedge \text{tid}(a'') = \text{tid}(e)$

3. One of the following holds

a. $G''_0 = G_1$ and $\nexists r \in G''_2.E$ such that $\langle e, r \rangle \in G''_2.\text{rf}$

b. $G''_0 = G_1 \setminus \{e\}$ and $e \notin G''_2.E$ and $\exists u' \in G''_2.E$ such that $\langle w, u' \rangle \in G''_2.\text{rf}$

$\langle 1 \rangle 1$. CASE: $n = 0$

PROOF: Obviously it is $G_1 = G'_1$. By choosing $m'' = 0$ and $G''_0 = G_1$, it is $m'' = n$, and $\nexists r \in G_1.E$ such that $\langle e, r \rangle \in G_1.\text{rf}$ (assumption $\langle 0 \rangle 2$).

$\langle 1 \rangle 2$. CASE: 1. $n = k + 1$

2. There exist m', G'_0, G'_2, G'_3 such that

a. $G'_0 \Rightarrow_{nr}^{m'} G'_2$

b. One of the following holds

i. $m' = k$

ii. $m' < k$ and exists $a' \in E$ such that $G'_2 \xRightarrow{a'}_{nr} G'_3 \wedge \text{tid}(a') = \text{tid}(e)$

c. One of the following holds

i. $G'_0 = G_1$ and $\nexists r \in G'_2.E$ such that $\langle e, r \rangle \in G'_2.\text{rf}$

ii. $G'_0 = G_1 \setminus \{e\}$ and $e \notin G'_2.E$ and $\exists u' \in G'_2.E$ such that $\langle w, u' \rangle \in G'_2.\text{rf}$

$\langle 2 \rangle 1$. CASE: $m' = k$

$\langle 3 \rangle 1$. Choose G'_3 such that $G'_2 \xRightarrow{e'}_{nr} G'_3$

PROOF: From assumptions $\langle 0 \rangle 5$, $\langle 1 \rangle 1$, and graph extensibility we get that Algorithm 2 can indeed take a step.

⟨3⟩2. CASE: $\text{tid}(e') = \text{tid}(e)$

PROOF: By choosing $m'' = k$, $G_0'' = G_0'$, $G_2'' = G_2'$, $G_3'' = G_3'$ we get that

1. $m'' < k + 1 = n$ and for $a'' = e'$ it is $G_2'' \xRightarrow{a''}_{nr} G_3'' \wedge \text{tid}(a'') = \text{tid}(e)$ (assumptions ⟨1⟩:1, ⟨2⟩, and ⟨3⟩)
2. One of the following holds (from assumption ⟨1⟩:2c)
 - a. $G_0'' = G_1$ and $\nexists r \in G_2''.E$ such that $\langle e, r \rangle \in G_2''.rf$
 - b. $G_0'' = G_1 \setminus \{e\}$ and $e \notin G_2''.E$ and $\exists u' \in G_2''.E$ such that $\langle w, u' \rangle \in G_2''.rf$

⟨3⟩3. CASE: $\text{tid}(e') \neq \text{tid}(e)$

⟨4⟩1. CASE: $G_0' = G_1$ and $\nexists r \in G_2'.E$ such that $\langle e, r \rangle \in G_2'.rf$

⟨5⟩1. CASE: $\text{typ}(e') = W$

PROOF: By choosing $m'' = k + 1$, $G_0'' = G_0'$, $G_2'' = G_3'$ we get

1. $G_0'' \xRightarrow{m''}_{nr} G_2''$ (assumptions ⟨1⟩:2a, ⟨3⟩, and step ⟨3⟩1)
2. $m'' = n$ (assumption ⟨1⟩:1)
3. $G_0'' = G_1$ and $\nexists r \in G_2''.E$ such that $\langle e, r \rangle \in G_2''.rf$ (step ⟨3⟩1, $e' \notin \text{rng}(G_2''.rf)$ because of ⟨5⟩, and case assumptions ⟨3⟩ and ⟨4⟩)

⟨5⟩2. CASE: $\text{typ}(e') = R$

PROOF: By choosing $m'' = k + 1$, $G_0'' = G_0'$, $G_2'' = G_3'$ we get

1. $G_0'' \xRightarrow{m''}_{nr} G_2''$ (assumptions ⟨1⟩:2a, ⟨3⟩, and step ⟨3⟩1)
2. $m'' = n$ (assumption ⟨1⟩:1)
3. $G_0'' = G_1$ and $\nexists r \in G_2''.E$ such that $\langle e, r \rangle \in G_2''.rf$ (from step ⟨3⟩1, case assumption ⟨5⟩ and Prop. D.3–Item (a), as well as case assumptions ⟨3⟩ and ⟨4⟩)

⟨5⟩3. CASE: $\text{typ}(e') = \text{RMW}$

⟨6⟩1. Exists $w' \in G_3'.E$ such that $\langle w', e' \rangle \in G_3'.rf$

PROOF: From ⟨3⟩1, case assumption ⟨5⟩, and graph extensibility.

⟨6⟩2. CASE: $w' \neq e$

PROOF: By choosing $m'' = k + 1$, $G_0'' = G_0'$, $G_2'' = G_3'$ we get

1. $G_0'' \xRightarrow{m''}_{nr} G_2''$ (assumptions ⟨1⟩:2a, ⟨3⟩, and step ⟨3⟩1)
2. $m'' = n$ (assumption ⟨1⟩:1)
3. $G_0'' = G_1$ and $\nexists r \in G_2''.E$ such that $\langle e, r \rangle \in G_2''.rf$ (from steps ⟨3⟩1 and ⟨6⟩2, as well as case assumptions ⟨3⟩ and ⟨4⟩)

⟨6⟩3. CASE: $w' = e$

⟨7⟩1. Let $G_2'' \triangleq G_3' \setminus \{e\}.rf[e' \mapsto w]$ such that $G_0' \xRightarrow{m''}_{nr} G_2''$ and $\exists u' \in G_2''.E$ such that $\langle w, u' \rangle \in G_2''.rf$

PROOF: From ⟨1⟩:2a we know that $G_0' \xRightarrow{m''}_{nr} G_2'$. Then, from graph extensibility (Prop. D.3–Item (b)) with case assumptions ⟨6⟩, ⟨5⟩, ⟨3⟩, and assumptions ⟨0⟩:3 and ⟨0⟩:4 we get the desired result, with $u' = e'$.

⟨7⟩2. Q.E.D.

PROOF: By choosing $m'' = k + 1$, $G_0'' = G_0' \setminus \{e\}$ and G_2'' from ⟨7⟩1 we get

1. $G_0'' \xRightarrow{m''}_{nr} G_2''$ (from Lemma D.11 with step ⟨7⟩1, and case assumptions ⟨4⟩ and ⟨1⟩:2a)
2. $m'' = n$ (assumption ⟨1⟩:1)
3. $G_0'' = G_1 \setminus \{e\}$ and $e \notin G_2''.E$ and $\exists u' \in G_2''.E$ such that $\langle w, u' \rangle \in G_2''.rf$ (from case assumption ⟨4⟩ and step ⟨7⟩1)

⟨6⟩4. Q.E.D.

PROOF: Cases ⟨6⟩2 and ⟨6⟩3 are exhaustive.

⟨5⟩4. Q.E.D.

PROOF: Cases ⟨5⟩1–⟨5⟩3 are exhaustive.

⟨4⟩2. CASE: $G_0' = G_1 \setminus \{e\}$ and $e \notin G_2'.E$ and $\exists u' \in G_2'.E$ such that $\langle w, u' \rangle \in G_2'.rf$

PROOF: By choosing $m'' = k + 1$, $G_0'' = G_0'$, $G_2'' = G_3'$ we get

1. $G_0'' \xRightarrow{m''}_{nr} G_2''$ (case assumptions ⟨1⟩:2a, ⟨3⟩, and step ⟨3⟩1)
2. $m'' = n$ (assumption ⟨1⟩:1)
3. $G_0'' = G_1 \setminus \{e\}$ and $e \notin G_2''.E$ and $\exists u' \in G_2''.E$ such that $\langle w, u' \rangle \in G_2''.rf$ (from graph extensibility, step ⟨3⟩1, and case assumptions ⟨3⟩ and ⟨4⟩)

⟨4⟩3. Q.E.D.

PROOF: From ⟨1⟩:2c we get that cases ⟨4⟩1 and ⟨4⟩2 are exhaustive.

⟨3⟩4. Q.E.D.

PROOF: Cases ⟨3⟩2 and ⟨3⟩3 are exhaustive.

$\langle 2 \rangle 2$. CASE: $m' < k$ and exists $a' \in E$ such that $G'_2 \xRightarrow{a'}_{nr} G'_3 \wedge \text{tid}(a') = \text{tid}(e)$

PROOF: By choosing $m'' = m'$, $G''_0 = G'_0$, $G''_2 = G'_2$, $G''_3 = G'_3$, and $a'' = a'$, we get that

1. $m'' < k + 1 = n$ and exists $a'' \in E$ such that $G''_2 \xRightarrow{a''}_{nr} G''_3 \wedge \text{tid}(a'') = \text{tid}(e)$ (assumptions $\langle 1 \rangle : 1$ and $\langle 2 \rangle$)

2. One of the following holds (from assumption $\langle 1 \rangle : 2c$)

a. $G''_0 = G_1$ and $\nexists r \in G''_2.E$ such that $\langle e, r \rangle \in G''_2.\text{rf}$

b. $G''_0 = G_1 \setminus \{e\}$ and $e \notin G''_2.E$ and $\exists u' \in G''_2.E$ such that $\langle w, u' \rangle \in G''_2.\text{rf}$

$\langle 2 \rangle 3$. Q.E.D.

PROOF: Cases $\langle 2 \rangle 1$ and $\langle 2 \rangle 2$ are exhaustive (assumption $\langle 1 \rangle : 2b$).

$\langle 1 \rangle 3$. Q.E.D.

PROOF: Steps $\langle 1 \rangle 1$, $\langle 1 \rangle 2$ and mathematical induction. □

Lemma D.11 (Equivalent production). *If there are $n \in \mathbb{N}$, $G_0, G_n, e \in E$ such that $\text{cons}(G_0), G_0 \Rightarrow_{nr}^n G_n$ and for all $G_i \in [G_0 \Rightarrow_{nr}^n G_n]$ it is $\text{next}_P(G_i) \neq e$ and either $\text{rng}([e]; G_i.\text{porf}) = \emptyset$ or $e \notin G_i.E$, then there is an equivalent production $G_0 \setminus \{e\} \Rightarrow_{nr}^n G_n \setminus \{e\}$.*

PROOF SKETCH: Simple proof by induction.

ASSUME: 1. P is a deterministic program with finite traces

2. $n \in \mathbb{N}$, $e \in E$, and G_0 such that $\text{cons}(G_0)$

3. $G_0 \Rightarrow_{nr}^n G_n$ such that, for all $G_i \in [G_0 \Rightarrow_{nr}^n G_n]$, it is

a. $\text{next}_P(G_i) \neq e$

b. $\text{rng}([e]; G_i.\text{porf}) = \emptyset$ or $e \notin G_i.E$

PROVE: $G_0 \setminus \{e\} \Rightarrow_{nr}^n G_n \setminus \{e\}$ and for all $G'_i \in [G_0 \setminus \{e\} \Rightarrow_{nr}^n G_n \setminus \{e\}]$ it is $e \notin G'_i$

$\langle 1 \rangle 1$. CASE: $n = 0$

PROOF: It obviously is $G_0 = G_n$ (assumption 3). From prefix-closedness (Def. 3.4) and assumption 2, it also is $\text{cons}(G_0 \setminus \{e\})$.

$\langle 1 \rangle 2$. CASE: 1. $n = k + 1$

2. IH: Given a G'_k such that $G_0 \Rightarrow_{nr}^k G'_k$, where for all $G'_i \in [G_0 \Rightarrow_{nr}^k G'_k]$ it is

a. $\text{next}_P(G'_i) \neq e$

b. $\text{rng}([e]; G'_i.\text{porf}) = \emptyset$ or $e \notin G'_i.E$

then it also is $G_0 \setminus \{e\} \Rightarrow_{nr}^k G'_k \setminus \{e\}$

$\langle 2 \rangle 1$. Let G_k, G_{k+1} such that $G_0 \Rightarrow_{nr}^k G_k \xRightarrow{e'}_{nr} G_{k+1}$ and $\forall G_i \in [G_0 \Rightarrow_{nr}^k G_k \xRightarrow{e'}_{nr} G_{k+1}]$ it is

1. $\text{next}_P(G_i) \neq e$

2. $\text{rng}([e]; G_i.\text{porf}) = \emptyset$ or $e \notin G_i.E$

PROOF: From assumptions $\langle 0 \rangle : 3$ and $\langle 1 \rangle : 1$.

$\langle 2 \rangle 2$. CASE: $e \notin G_{k+1}$

PROOF: It is $G_{k+1} \setminus \{e\} = G_{k+1}$ (case assumption $\langle 2 \rangle$), and since from the IH we get that $G_0 \setminus \{e\} \Rightarrow_{nr}^k G_k \setminus \{e\}$ (assumption

$\langle 1 \rangle : 2$) with $\text{cons}(G_k \setminus \{e\})$ (courtesy of Algorithm 2), it also is $G_0 \setminus \{e\} \Rightarrow_{nr}^k G_k \setminus \{e\} \xRightarrow{e'}_{nr} G_{k+1} \setminus \{e\}$ (from step $\langle 2 \rangle 1$).

$\langle 2 \rangle 3$. CASE: $\text{rng}([e]; G_{k+1}.\text{porf}) = \emptyset$

$\langle 3 \rangle 1$. It is $e' \neq e$

PROOF: From $\langle 2 \rangle 1$.

$\langle 3 \rangle 2$. Q.E.D.

PROOF: From prefix-closedness (Def. 3.4) we get that it is $\text{cons}(G_{k+1} \setminus \{e\})$, and from the IH (assumption $\langle 1 \rangle : 2$), we get that $G_0 \setminus \{e\} \Rightarrow_{nr}^k G_k \setminus \{e\}$ with $\text{cons}(G_k \setminus \{e\})$ (algorithmic construction). Thus, from step $\langle 3 \rangle 1$, it can also be $G_0 \setminus \{e\} \Rightarrow_{nr}^{k+1} G_{k+1} \setminus \{e\}$.

$\langle 2 \rangle 4$. Q.E.D.

PROOF: From assumption $\langle 0 \rangle : 3b$ and step $\langle 2 \rangle 1$ we get that cases $\langle 2 \rangle 2$ and $\langle 2 \rangle 3$ are exhaustive.

$\langle 1 \rangle 3$. Q.E.D.

PROOF: Steps $\langle 1 \rangle 1$, $\langle 1 \rangle 2$ and mathematical induction. □

Algorithm 4 Explore one program execution

```

1: procedure VISITONE( $G, T, U, S, A$ )
2:   while  $\text{cons}(G) \wedge a \leftarrow \text{next}_P(G)$  do
3:     if  $a \in \text{error}$  then exit("erroneous program")
4:      $G \leftarrow \text{Add}(G, a)$ 
5:     if  $a \in R$  then
6:        $W \leftarrow G.E \cap W_{\text{loc}}(a)$ 
7:        $B \leftarrow \{ \langle \perp, \emptyset \rangle \mid \text{cons}(G.\text{rf}[a \mapsto \perp]) \}$ 
8:       choose some  $w_0 \in W$ 
9:        $G.\text{rf}[r] \leftarrow w_0$ 
10:       $T \leftarrow T \cup \{r\}$ 
11:       $S[a] \leftarrow S[a] \cup \{ \langle w, \emptyset \rangle \mid w \in W \setminus \{w_0\} \} \cup B$ 
12:       $U[a] \leftarrow U[a] \cup B$ 
13:       $\text{CALCREVISITS}(G, T, U, S, a)$ 
14:       $A \leftarrow A \cup \{G\}$ 

```

D.2 Optimality Proof

Definition D.12 (Compatibility). Two executions G_1 and G_2 are *compatible* over a program P , written $G_1 \#_P G_2$, if:

$$\begin{aligned}
G_1 \#_P G_2 \stackrel{\text{def}}{\iff} & \exists G'_1, G'_2, E_1, E_2, r_1, r_2, s_1, s_2. \begin{aligned} & G_1 \Rightarrow_{nr}^* G'_1 \wedge \text{next}_P(G'_1) = \perp \wedge E_1 = G'_1.E \\ & \wedge G_2 \Rightarrow_{nr}^* G'_2 \wedge \text{next}_P(G'_2) = \perp \wedge E_2 = G'_2.E \\ & \wedge G_1.E \cup E_1 = G_2.E \cup E_2 \\ & \wedge G_1.\text{rf} \cup r_1 = G_2.\text{rf} \cup r_2 \\ & \wedge G_1.\text{po} \cup s_1 = G_2.\text{po} \cup s_2 \\ & \wedge s_1 \subseteq (G_1.E \cup E_1) \times E_1 \wedge s_2 \subseteq (G_2.E \cup E_2) \times E_2 \end{aligned}
\end{aligned}$$

Two executions are *orthogonal* over a program P , written $G \perp_P G'$, if they are incompatible: $G \perp_P G' \stackrel{\text{def}}{\iff} \neg(G \#_P G')$.

Lemma D.13. *Compatibility relation is an equivalence (reflexive, symmetric and transitive) relation.*

For all G_1, G_2, G_3 , if $G_1 \leq_P G_2$ and $G_1 \perp_P G_3$, then $G_2 \perp_P G_3$, where:

$$\begin{aligned}
G \leq_P G' \stackrel{\text{def}}{\iff} & \exists G'', E, r, s. \begin{aligned} & G \Rightarrow_{nr}^* G'' \wedge \text{next}_P(G'') = \perp \wedge E = G''.E \\ & \wedge G.E \cup E = G'.E \wedge G.\text{rf} \cup r = G'.\text{rf} \wedge G.\text{po} \cup s = G'.\text{po} \\ & \wedge s \subseteq G'.E \times E \end{aligned}
\end{aligned}$$

In what follows, when the choice of program P is clear from the context, we drop the P subscript and simply write $G_1 \perp G_2$ for $G_1 \perp_P G_2$; and write $G_1 \leq G_2$ for $G_1 \leq_P G_2$.

Definition D.14 (Instrumented configuration). An instrumented configuration is a tuple $\langle G, T, U, S, A \rangle$, where G, T, U and S are defined as before, and A is a set of consistent executions.

Let us *instrument* our configurations such that an instrumented configuration additionally carries a set A of all the executions generated so far. Similarly, we instrument the implementation in Algorithm 1, Algorithm 2 and Algorithm 3 such that they operate on these instrumented configurations. We thus extend Algorithm 2 as in Algorithm 4 by adding the **highlighted** line, recording the generated execution in A .

Definition D.15 (Instrumented construction). Given a program P , an instrumented configuration $\langle G, T, U, S, A \rangle$ is ordered before $\langle G', T', U', S', A' \rangle$ in the *instrumented order*, written $\langle G, T, U, S, A \rangle \Rightarrow_P \langle G', T', U', S', A' \rangle$, if a step of the instrumented algorithm for P transforms $\langle G, T, U, S, A \rangle$ to $\langle G', T', U', S', A' \rangle$.

When the choice of P is clear from the context, we drop the P subscript and simply write $\Rightarrow$ for $\Rightarrow_P$.

Lemma D.16. For all $\langle G, T, U, S, A \rangle$ and $\langle G', T', U', S', A' \rangle$,

$$\text{inv}(\langle G, T, U, S, A \rangle) \wedge \langle G, T, U, S, A \rangle \Rightarrow \langle G', T', U', S', A' \rangle \Rightarrow \text{inv}(\langle G', T', U', S', A' \rangle)$$

where

$$\begin{aligned}
 \text{inv}(\langle G, T, U, S, A \rangle) \triangleq & \\
 & \forall r. S[r] \cup U[r] \neq \emptyset \Rightarrow r \in G.E \\
 & \wedge \forall (w, r) \in G.\text{rf}. r \in T \wedge (r, w) \in G.\text{pred} \Rightarrow \\
 & \quad \exists P_w. P_w = \text{sbrf}(G, r, w) \wedge (w, P_w) \in U[r] \wedge \text{new}(G, r, w, P_w) \leq G \\
 & \wedge \forall r_1, r_2, w_1. (w_1, r_1) \in G.\text{rf} \wedge (r_1, r_2) \in G.\text{pred} \wedge (r_2 \in T \vee S[r_2] \neq \emptyset) \Rightarrow \\
 & \quad (G.\text{porf}^?; [w_1]) \times \{r_2\} \subseteq G.\text{pred} \\
 & \wedge \forall G_1, G_2 \in A \cup \{G\}. G_1 \neq G_2 \Rightarrow G_1 \perp G_2 \\
 & \wedge \forall G_a \in A. \forall r. \forall (w, P_w) \in S[r]. \text{new}(G, r, w, P_w) \perp G_a \\
 & \wedge \forall G_a \in A. \forall r \in T \cap G.E. \forall w. (w, r) \in G_a.\text{rf} \Rightarrow \\
 & \quad (w, r) \in G.\text{pred} \vee \text{cut}(G, r) \perp G_a \vee \exists P_w. (w, P_w) \in U[r] \wedge \text{new}(G, r, w, P_w) \leq G_a \\
 & \wedge \forall r. (w, P_w) \in S[r]. G.\text{rf}^{-1}(r) \neq w \vee \text{sbrf}(G, r, w) \perp P_w
 \end{aligned}$$

with

$$\begin{aligned}
 \text{new}(G, r, w, P_w) &\triangleq \text{Combine}(G', P_w).\text{rf}[r \mapsto w] \text{ with } G' = \text{Remove}(G, \text{rng}([r]; G.\text{pred})) \\
 \text{sbrf}(G, w, r) &\triangleq G|_{\text{dom}(G.\text{porf}^?; [w]) \setminus \text{dom}(G.\text{pred}^?; [r])} \\
 \text{cut}(G, r) &\triangleq \text{Remove}(G, \text{rng}([r]; G.\text{pred}^?))
 \end{aligned}$$

Proof. Pick arbitrary $\langle G, T, U, S, A \rangle, \langle G', T', U', S', A' \rangle$ such that $\text{inv}(\langle G, T, U, S, A \rangle)$ holds and $\langle G, T, U, S, A \rangle \Rightarrow \langle G', T', U', S', A' \rangle$. That is, one step of the instrumented algorithm transforms $\langle G, T, U, S, A \rangle$ to $\langle G', T', U', S', A' \rangle$. There are then three cases to consider: 1) $\text{next}_P(G) \in R$; or 2) $\text{next}_P(G) \in W$; or 3) $\text{next}_P(G) = \perp$.

Case 1.

Let $r = \text{next}_P(G) \in R$. We then know $\langle G', T', U', S', A' \rangle = \text{VISITREAD}(G_r, T, U, S, A, r)$ with $G_r = \text{Add}(G, r)$. Let $W = W_{\text{cons}} \cup W_{\text{excl}}$ with W_{cons} and W_{excl} as defined in the algorithm. From the instrumented algorithm we then know there exists $w_0 \in W$ such that $G' = G_r.\text{rf}[r \mapsto w_0]$, $T' = T \cup \{r\}$, $A' = A$, $U' = U$, and $S'[r'] = S[r']$ for all $r' \neq r$ and $S'[r] = \{(w, \emptyset) \mid w \in W \setminus \{w_0\}\}$. We are then required to show:

$$\forall r'. S'[r'] \cup U'[r'] \neq \emptyset \Rightarrow r' \in G'.E \quad (1)$$

$$\begin{aligned}
 & \forall (w, r') \in G'.\text{rf}. r' \in T' \wedge (r', w) \in G'.\text{pred} \Rightarrow \\
 & \quad \exists P_w. P_w = \text{sbrf}(G', r', w) \wedge (w, P_w) \in U'[r'] \wedge \text{new}(G', r', w, P_w) \leq G'
 \end{aligned} \quad (2)$$

$$\begin{aligned}
 & \forall r_1, r_2, w_1. (w_1, r_1) \in G'.\text{rf} \wedge (r_1, r_2) \in G'.\text{pred} \wedge (r_2 \in T' \vee S'[r_2] \neq \emptyset) \Rightarrow \\
 & \quad (G'.\text{porf}^?; [w_1]) \times \{r_2\} \subseteq G'.\text{pred}
 \end{aligned} \quad (3)$$

$$\forall G_1, G_2 \in A \cup \{G'\}. G_1 \neq G_2 \Rightarrow G_1 \perp G_2 \quad (4)$$

$$\forall G_a \in A. \forall r'. \forall (w, P_w) \in S'[r']. \text{new}(G', r', w, P_w) \perp G_a \quad (5)$$

$$\begin{aligned}
 & \forall G_a \in A. \forall r' \in T' \cap G'.E. \forall w. (w, r') \in G_a.\text{rf} \Rightarrow \\
 & \quad (w, r') \in G'.\text{pred} \vee \text{cut}(G', r') \perp G_a \vee \exists P_w. (w, P_w) \in U'[r'] \wedge \text{new}(G', r', w, P_w) \leq G_a
 \end{aligned} \quad (6)$$

$$\forall r'. (w, P_w) \in S'[r']. G'.\text{rf}^{-1}(r') \neq w \vee \text{sbrf}(G', r', w) \perp P_w \quad (7)$$

Part (1) follows immediately from the definitions of G', S', U' and $\text{inv}(\langle G, T, U, S, A \rangle)$.

Part (2) follows from the definitions of G', U' , the $\text{inv}(\langle G, T, U, S, A \rangle)$, that $(w_0, r) \in G'.\text{pred}$, that $G \leq G'$, and that for all w, P_w and $r' \neq r$, $\text{new}(G, r', w, P_w) = \text{new}(G', r', w, P_w)$.

Part (3) follows from the definitions of G', S' , $\text{inv}(\langle G, T, U, S, A \rangle)$ and the fact that r is the maximal element in $G'.\text{pred}$.

For (4), as $\text{inv}(\langle G, T, U, S, A \rangle)$ holds, it suffices to show $\forall G_1 \in A. G' \perp G_1$. Pick an arbitrary $G_1 \in A$; from the definition of G' we then have $G \leq G'$. As such, since from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $G \perp G_1$, from Lemma D.13 we have $G' \perp G_1$ as required.

For (5), pick arbitrary $G_a \in A$, r' and $(w, P_w) \in S'[r']$. There are then two cases to consider: i) $r' \neq r$; ii) $r' = r$. In case (i), we know that $(w, P_w) \in S[r']$ and thus since r is the maximal (in pred order) in G' , we know that $\text{new}(G, r', w, P_w) = \text{new}(G', r', w, P_w)$ and thus from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $\text{new}(G', r', w, P_w) \perp G_a$, as required. In case (ii), we have $P_w = \emptyset$. From the definition of G' we have $G \leq \text{new}(G', r', w, P_w)$. As such, since from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $G \perp G_a$, from Lemma D.13 we have $\text{new}(G', r', w, P_w) \perp G_a$, as required.

For (6), pick arbitrary $G_a \in A$, $r' \in T' \cap G'.E$ and w such that $(w, r') \in G_a.\text{rf}$. There are then two cases to consider: i) $r' \neq r$; ii) $r' = r$. In case (i), we know $r' \in T \cap G.E$. As such, from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $(w, r') \in G.\text{pred} \vee \text{cut}(G, r') \perp G_a \vee \exists P_w. (w, P_w) \in U[r'] \wedge \text{new}(G, r', w, P_w) \leq G_a$. That is, since $U' = U$, from the definition of G' we

have $(w, r') \in G'.\text{pred} \vee \text{cut}(G, r') \perp G_a \vee \exists P_w. (w, P_w) \in U'[r'] \wedge \text{new}(G, r', w, P_w) \leq G_a$. Moreover, since r is the maximal (in pred order) in G' , we know that $\text{cut}(G, r') = \text{cut}(G', r')$ and $\text{new}(G, r', w, P_w) = \text{new}(G', r', w, P_w)$. We thus have $(w, r') \in G'.\text{pred} \vee \text{cut}(G', r') \perp G_a \vee \exists P_w. (w, P_w) \in U'[r'] \wedge \text{new}(G', r', w, P_w) \leq G_a$, as required. In case (ii), from the definition of G' we have $G = \text{cut}(G', r')$. As such, since from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $G \perp G_a$, we have $\text{cut}(G', r') \perp G_a$, as required.

For (7), pick arbitrary w, P_w, r' such that $(w, P_w) \in S'[r']$. There are two cases to consider: i) $r' \neq r$; ii) $r' = r$. In case (i), we know $(w, P_w) \in S[r]$. From the definition of G' we thus have: $G.\text{rf}^{-1}(r') = G'.\text{rf}^{-1}(r')$ and $\text{sbrf}(G, r', w) = \text{sbrf}(G', r', w)$ and thus the desired result follows immediately from $\text{inv}(\langle G, T, U, S, A \rangle)$. In case (ii), from the definition of G' we know that $w \neq w_0$, as required.

Case 2.

Let $w = \text{next}_P(G) \in W$. We then know $\langle G', T', U', S', A' \rangle = \text{VISITWRITE}(G_w, T, U, S, A, r)$ with $G_w = \text{Add}(G, w)$. Let set R be as defined in the algorithm. From the instrumented algorithm we then know $G' = G_w, T' = T, A' = A$, and for all r , when $r \notin R$ we have $S'[r] = S[r]$ and $U'[r] = U[r]$; and when $r \in R$ we have $S'[r] = S[r] \cup \langle w, P_w \rangle$ and $U'[r] = U[r] \cup \langle w, P_w \rangle$ with $P_w = \text{sbrf}(G, w, r)$. We are then required to show:

$$\forall r'. S'[r'] \cup U'[r'] \neq \emptyset \Rightarrow r' \in G'.E \quad (8)$$

$$\begin{aligned} \forall (w', r) \in G'.\text{rf}. r \in T' \wedge (r, w') \in G'.\text{pred} \Rightarrow \\ \exists P_w. P_w = \text{sbrf}(G', r, w') \wedge (w', P_w) \in U'[r] \wedge \text{new}(G', r, w', P_w) \leq G' \end{aligned} \quad (9)$$

$$\begin{aligned} \forall r_1, r_2, w_1. (w_1, r_1) \in G'.\text{rf} \wedge (r_1, r_2) \in G'.\text{pred} \wedge (r_2 \in T' \vee S'[r_2] \neq \emptyset) \Rightarrow \\ (G'.\text{porf}^2; [w_1]) \times \{r_2\} \subseteq G'.\text{pred} \end{aligned} \quad (10)$$

$$\forall G_1, G_2 \in A \cup \{G'\}. G_1 \neq G_2 \Rightarrow G_1 \perp G_2 \quad (11)$$

$$\forall G_a \in A. \forall r. \forall (w', P_w) \in S'[r]. \text{new}(G', r, w', P_w) \perp G_a \quad (12)$$

$$\begin{aligned} \forall G_a \in A. \forall r \in T' \cap G'.E. \forall w'. (w', r) \in G_a.\text{rf} \Rightarrow \\ (w', r) \in G'.\text{pred} \vee \text{cut}(G', r) \perp G_a \vee \exists P_w. (w', P_w) \in U'[r] \wedge \text{new}(G', r, w', P_w) \leq G_a \end{aligned} \quad (13)$$

$$\forall r, (w', P_w) \in S'[r]. G'.\text{rf}^{-1}(r) \neq w' \vee \text{sbrf}(G', r, w') \perp P_w \quad (14)$$

Part (8) follows immediately from the definitions of G', S', U' and $\text{inv}(\langle G, T, U, S, A \rangle)$.

Part (9) follows from the definitions of G', U' , the $\text{inv}(\langle G, T, U, S, A \rangle)$, the fact that $\neg \exists r. (w, r) \in G'.\text{rf}$, that $G \leq G'$, and the fact that and that for all r, P_w and $w' \neq w$, $\text{new}(G, r, w', P_w) = \text{new}(G', r, w', P_w)$.

Part (10) follows from the definitions of G', S' , $\text{inv}(\langle G, T, U, S, A \rangle)$, the fact that for all r', w', P'_w if $(w', P'_w) \in S'[r'] \setminus S[r']$ then $r' \in T$, and the fact that for all r_1 we have $(w, r_1) \notin G'.\text{rf}$.

For (11), as $\text{inv}(\langle G, T, U, S, A \rangle)$ holds, it suffices to show $\forall G_a \in A. G' \perp G_a$. Pick an arbitrary $G_a \in A$; from the definition of G' we then have $G \leq G'$. As such, since from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $G \perp G_a$, from Lemma D.13 we have $G' \perp G_a$ as required.

For (12), pick arbitrary $G_a \in A, r$ and $(w', P_w) \in S'[r]$. There are two cases to consider: i) $(w', P_w) \in S[r]$; or ii) $(w', P_w) \in S'[r] \setminus S[r]$ and thus $r \in T \cap G.E, w' = w, P_w = \text{sbrf}(G', r, w)$ and $(w, P_w) \notin U[r]$. In case (i), since w is maximal (in pred order) in G' , we know $\text{new}(G, r, w', P_w) = \text{new}(G', r, w', P_w)$ and thus from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $\text{new}(G', r, w', P_w) \perp G_a$, as required.

In case (ii), let $G'' = \text{new}(G', r, w, P_w)$. We know that either $r \notin G_a.E$ or $r \in G_a.E$. In the former case, since G_a is a completed graph and thus $\text{next}_P(G_a) = \perp$, by definition we simply have $G_a \perp G''$, as required. In the latter case, pick w'' such that $(w'', r) \in G_a.\text{rf}$. Since $r \in T \cap G.E$, from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $(w'', r) \in G.\text{pred} \vee \text{cut}(G, r) \perp G_a \vee \exists P''_w. (w'', P''_w) \in U[r] \wedge \text{new}(G, r, w'', P''_w) \leq G_a$. In case of the first disjunct, since $w \notin G.E$ and $w'' \in G.E$, we have $w'' \neq w$, i.e., $G_a.\text{rf}^{-1}(r) \neq G''.\text{rf}^{-1}(r)$ and thus $G'' \perp G_a$, as required. In case of the second disjunct, from the definition of G'' we have $\text{cut}(G, r) \leq G''$. As such, from Lemma D.13 we have $G'' \perp G_a$, as required. In case of the last disjunct, since $(w, P_w) \notin U[r]$ and $(w'', P''_w) \in U[r]$, we know that either $w'' \neq w$ or $w'' = w \wedge P''_w \neq P_w$. In the former case we then have $G_a.\text{rf}^{-1}(r) \neq G''.\text{rf}^{-1}(r)$ and thus $G'' \perp G_a$, as required. In the latter case, it is straightforward to demonstrate that as w is the porf maximal event in both P_w and P''_w , since $P''_w \neq P_w$, we also have $P''_w \perp P_w$. As such, we have $\text{new}(G, r, w, P''_w) \perp \text{new}(G, r, w, P_w)$. Moreover, since $w'' = w$ and $\text{new}(G, r, w'', P''_w) \leq G_a$, we have $\text{new}(G, r, w, P''_w) \leq G_a$. Consequently, from Lemma D.13 we have $\text{new}(G, r, w, P_w) \perp G_a$. On the other hand, from the definition of G' we have $\text{new}(G, r, w, P_w) = \text{new}(G', r, w, P_w)$. As such, we have $\text{new}(G, r, w, P_w) \perp G_a$, as required.

For (13), pick arbitrary $G_a \in A$, $r \in T' \cap G'.E$ and w' with $(w', r) \in G_a.\text{rf}$. From $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $(w', r) \in G.\text{pred} \vee \text{cut}(G, r) \perp G_a \vee \exists P_w. (w', P_w) \in U[r] \wedge \text{new}(G, r, w', P_w) \leq G_a$. From the definition of G' and since w is maximal (in $G'.\text{pred}$ order), we know that $\text{cut}(G, r) = \text{cut}(G', r)$ and $\text{new}(G, r, w', P_w) = \text{new}(G', r, w', P_w)$. As such, since $U[r] \subseteq U'[r]$, from the definition of G' we have $(w', r) \in G'.\text{pred} \vee \text{cut}(G', r) \perp G_a \vee \exists P_w. (w', P_w) \in U'[r] \wedge \text{new}(G, r, w', P_w) \leq G_a$, as required.

For (14), pick arbitrary w', P_w, r such that $(w', P_w) \in S'[r]$. There are two cases to consider: i) $(w', P_w) \in S[r]$; ii) $(w', P_w) \in S'[r] \setminus S[r]$ and thus $w' = w$. In case (i), since w is maximal (in $G'.\text{pred}$ order), we have $\text{sbrf}(G, r, w') \leq \text{sbrf}(G', r, w')$. As such, the desired result follows from $\text{inv}(\langle G, T, U, S, A \rangle)$ and Lemma D.13. In case (ii), since $w \notin G$, we know that $G.\text{rf}^{-1}(r) \neq w$. As such, from the definition of G' we have $G'.\text{rf}^{-1}(r) = G.\text{rf}^{-1}(r) \neq w$, as required.

Case 3.

Let $\perp = \text{next}_P(G)$. We then know $\langle G', T', U', S', A' \rangle = \text{REVISITREADS}(G, T, U, S, A)$. That is, there exists w, P_w, r such that $r \in G.E$; $\langle w, P_w \rangle \in S[r]$; for all r' if $(r, r') \in G.\text{pred}$ then $S[r'] = \emptyset$; $G' = \text{new}(G, r, w, P_w)$; $T' = T \setminus P_w.E$; $A' = A \cup \{G\}$; $S'[r] = S[r] \setminus \{\langle w, P_w \rangle\}$; for all $r' \neq r$ we have $S'[r'] = S[r']$; for all r' such that $(r, r') \in G.\text{pred}$ we have $U'[r'] = \emptyset$; and for all r' such that $(r', r) \in G.\text{pred}^2$ we have $U'[r'] = U[r']$. We are then required to show:

$$\forall r'. S'[r'] \cup U'[r'] \neq \emptyset \Rightarrow r' \in G'.E \quad (15)$$

$$\begin{aligned} \forall (w', r') \in G'.\text{rf}. r' \in T' \wedge (r', w') \in G'.\text{pred} \Rightarrow \\ \exists P'_w. P'_w = \text{sbrf}(G', r', w') \wedge (w', P'_w) \in U'[r'] \wedge \text{new}(G', r', w', P'_w) \leq G' \end{aligned} \quad (16)$$

$$\begin{aligned} \forall r_1, r_2, w_1. (w_1, r_1) \in G'.\text{rf} \wedge (r_1, r_2) \in G'.\text{pred} \wedge (r_2 \in T' \vee S'[r_2] \neq \emptyset) \Rightarrow \\ (G'.\text{porf}^2; [w_1]) \times \{r_2\} \subseteq G'.\text{pred} \end{aligned} \quad (17)$$

$$\forall G_1, G_2 \in A \cup \{G\} \cup \{G'\}. G_1 \neq G_2 \Rightarrow G_1 \perp G_2 \quad (18)$$

$$\forall G_a \in A \cup \{G\}. \forall r'. \forall (w', P'_w) \in S'[r']. \text{new}(G', r', w', P'_w) \perp G_a \quad (19)$$

$$\begin{aligned} \forall G_a \in A \cup \{G\}. \forall r' \in T' \cap G'.E. \forall w'. (w', r') \in G_a.\text{rf} \Rightarrow \\ (w', r') \in G'.\text{pred} \vee \text{cut}(G', r') \perp G_a \vee \exists P'_w. (w', P'_w) \in U'[r'] \wedge \text{new}(G', r', w', P'_w) \leq G_a \end{aligned} \quad (20)$$

$$\forall r', (w', P'_w) \in S'[r']. G'.\text{rf}^{-1}(r') \neq w' \vee \text{sbrf}(G', r', w') \perp P'_w \quad (21)$$

Part (15) follows immediately from the definitions of $G', S', U', \text{inv}(\langle G, T, U, S, A \rangle)$ and the maximality of r in S .

For part (16), pick an arbitrary $(w', r') \in G'.\text{rf}$ such that $r' \in T'$ and $(r', w') \in G'.\text{pred}$. Since $r' \in T'$ we know $r' \notin P_w$ and thus $(r', r) \in G'.\text{pred}^2$, and thus from the definition of G' : $(r', r) \in G.\text{pred}^2$. There are now two cases to consider: 1) $r' = r$; 2) $r' \neq r$. In case (1), from the definition of G' we have $\text{sbrf}(G', r, w) = P_w$ and that $G' = \text{new}(G', r, w, P_w)$. Moreover, since $(w, P_w) \in U[r]$, from the definition of U' we have $(w, P_w) \in U'[r]$. We thus know there exists $w' = w$ and $P'_w = P_w$ such that $\text{sbrf}(G', r, w') = P'_w$, $(w', P'_w) \in U'[r]$ and $G' = \text{new}(G', r, w', P'_w) \leq G'$, as required.

In case (2), since $(r', r) \in G.\text{pred}^2$ and $r' \neq r$, we have $(r', r) \in G.\text{pred}$. As $(w', r') \in G'.\text{rf}$, from the definition of G' we know $(w', r') \in G.\text{rf}$. Moreover, since $S[r] \neq \emptyset$ ($(w, P_w) \in S[r]$), from $\text{inv}(\langle G, T, U, S, A \rangle)$ we know $(G.\text{porf}^2; [w]) \times \{r\} \subseteq G.\text{pred}$ and thus $(G'.\text{porf}^2; [w]) \times \{r\} \subseteq G'.\text{pred}$. That is, $(w', r) \in G.\text{pred}$ and thus $(w', r) \in G'.\text{pred}$, and $\text{sbrf}(G, r', w') = \text{sbrf}(G', r', w')$. As $(w', r) \in G.\text{pred}$ and $(r', w') \in G'.\text{pred}$, from the definition of G' we know $(r', w') \in G.\text{pred}$. On the other hand, since $r' \in T'$, from the definition of T we know $r \in T$. As such, since $(w', r') \in G.\text{rf}$, $r \in T$ and $(r', w') \in G.\text{pred}$, from $\text{inv}(\langle G, T, U, S, A \rangle)$ we know there exists P'_w such that $P'_w = \text{sbrf}(G, r', w')$, $(w', P'_w) \in U[r'] \wedge \text{new}(G, r', w', P'_w) \leq G$. On the other hand, from the definition of G' and U' and since $(r', r) \in G.\text{pred}$, we know $\text{new}(G, r', w', P'_w) = \text{new}(G', r', w', P'_w)$ and $U'[r'] = U[r']$. As $\text{sbrf}(G, r', w') = \text{sbrf}(G', r', w')$, we know there exists P'_w such that $P'_w = \text{sbrf}(G', r', w')$, $(w', P'_w) \in U'[r']$. Moreover, since $P'_w = \text{sbrf}(G', r', w')$, and $(w', r') \in G'.\text{rf}$, from the definitions of $\text{new}(\cdot, \cdot, \cdot, \cdot)$ and $\text{sbrf}(\cdot, \cdot, \cdot)$ we have $\text{new}(G', r', w', P'_w) \leq G'$. That is, there exists P'_w such that $P'_w = \text{sbrf}(G', r', w')$, $(w', P'_w) \in U'[r'] \wedge \text{new}(G', r', w', P'_w) \leq G'$, as required.

For part (17), pick arbitrary w_1, r_1, r_2 such that $(w_1, r_1) \in G'.\text{rf}$, $(r_1, r_2) \in G'.\text{pred}$ and $r_2 \in T' \vee S'[r_2] \neq \emptyset$. We first demonstrate that $(r_2, r) \in G.\text{pred}^2$. Given the disjunction $r_2 \in T' \vee S'[r_2] \neq \emptyset$, if the first disjunct holds, since $r_2 \in G'.E$ and from the definition of T' we know $T' \cap P_w.E = \emptyset$, from the definition of G' we have $(r_2, r) \in G.\text{pred}^2$. If however the second disjunct holds, from the definition of S' and since r is picked to be maximal in S , we have $(r_2, r) \in G.\text{pred}^2$.

Since $(r_1, r_2) \in G'.\text{pred}$, $(r_2, r) \in G.\text{pred}^2$ and thus $(r_2, r) \in G'.\text{pred}^2$, from the definition of G' we know $(r_1, r_2) \in G'.\text{pred}$ and thus $(r_1, r_2) \in G.\text{pred}$. Moreover, since $(w_1, r_1) \in G'.\text{rf}$, from the definition of G' we know $(w_1, r_1) \in G.\text{rf}$. Consequently, since $r_2 \in T' \vee S'[r_2] \neq \emptyset$, $T' \subseteq T$ and $S'[r_2] \subseteq S[r_2]$, from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $(G.\text{porf}^2; [w_1]) \times \{r_2\} \subseteq G.\text{pred}$. As such, since $(r_2, r) \in G.\text{pred}^2$, from the definition of G' we have $(G.\text{porf}^2; [w_1]) \times \{r_2\} \subseteq G'.\text{pred}$. Moreover, from

the definition of G' and since $(G.\text{porf}^?; [w_1]) \times \{r_2\} \subseteq G'.\text{pred}$, and $(r_2, r) \in G.\text{pred}^?$ (thus $(r_2, r) \in G'.\text{pred}^?$), we know $(G.\text{porf}^?; [w_1]) = (G'.\text{porf}^?; [w_1])$. That is, we have $(G'.\text{porf}^?; [w_1]) \times \{r_2\} \subseteq G'.\text{pred}$, as required.

For (18), as $\text{inv}(\langle G, T, U, S, A \rangle)$ holds, it suffices to show: i) $G' \perp G$ and ii) $G' \perp G_a$ for all $G_a \in A$. Part (i) follows from the definition of G' and the last conjunct of $\text{inv}(\langle G, T, U, S, A \rangle)$. Part (ii) follows from the definition of G' and the fifth conjunct of $\text{inv}(\langle G, T, U, S, A \rangle)$.

For (19), pick an arbitrary $G_a \in A \cup \{G\}$, r' and $(w', P'_w) \in S'[r']$. From the definition of S' we then have $(w', P'_w) \in S[r']$. Moreover, as r is the maximal entry in S , we know that $(r', r) \in G'.\text{pred}^?$, and thus from the definition of G' we have $\text{new}(G', r', w', P'_w) = \text{new}(G, r', w', P'_w)$. There are now two cases to consider: 1) $G_a \in A$; or 2) $G_a = G$.

In case (1), from the definition of G' we have $\text{new}(G', r', w', P'_w) = \text{new}(G, r', w', P'_w)$. from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $\text{new}(G, r', w', P'_w) \perp G_a$, and since $\text{new}(G', r', w', P'_w) = \text{new}(G, r', w', P'_w)$, we have $\text{new}(G', r', w', P'_w) \perp G_a$, as required.

In case (2), from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $G.\text{rf}^{-1}(r') \neq w' \vee \text{sbrf}(G, r', w') \perp P'_w$, and thus $\text{new}(G, r', w', P'_w) \perp G$. As such, since $\text{new}(G', r', w', P'_w) = \text{new}(G, r', w', P'_w)$, we have $\text{new}(G', r', w', P'_w) \perp G$, as required.

For (20), pick an arbitrary $G_a \in A \cup \{G\}$, $r' \in T' \cap G'.E$ and w' such that $(w', r') \in G_a.\text{rf}$. From the definition of T' and G' we then have $r' \in T \cap G.E$ and $r' \notin P_w$, and since r is the maximal entry in S , we know that $(r', r) \in G'.\text{pred}^?$. There are now two cases to consider: 1) $G_a \in A$; or 2) $G_a = G$.

In case (1), from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $(w', r') \in G.\text{pred} \vee \text{cut}(G, r') \perp G_a \vee \exists P'_w. (w', P'_w) \in U[r'] \wedge \text{new}(G, r', w', P'_w) \leq G_a$. From the definition of G' we have $\text{cut}(G', r') = \text{cut}(G, r')$ and $\text{new}(G', r', w', P'_w) = \text{new}(G, r', w', P'_w)$. Moreover, from the definition of U' and since $(r', r) \in G'.\text{pred}^?$, we have $U'[r'] = U[r']$. As such, from the definitions of G' we have $(w', r') \in G'.\text{pred} \vee \text{cut}(G', r') \perp G_a \vee \exists P'_w. (w', P'_w) \in U'[r'] \wedge \text{new}(G', r', w', P'_w) \leq G_a$, as required.

In case (2), there are two cases to consider: either $(w', r') \in G.\text{pred}$; or $(w', r') \notin G.\text{pred}$. In the former case from the definition of G' we have $(w', r') \in G'.\text{pred}$, as required. In the latter case, from $\text{inv}(\langle G, T, U, S, A \rangle)$ we know there exists P'_w such that $P'_w = \text{sbrf}(G, r', w')$, $(w', P'_w) \in U[r']$ and $\text{new}(G, r', w', P'_w) \leq G$. From the definition of U' and since $(r', r) \in G'.\text{pred}^?$, we have $U'[r'] = U[r']$. Moreover, from the definition of G' we have $\text{new}(G', r', w', P'_w) = \text{new}(G, r', w', P'_w) \leq G$. As such, we know there exists P'_w such that $(w', P'_w) \in U'[r']$ and $\text{new}(G', r', w', P'_w) \leq G$, as required.

For (21), pick arbitrary w', P'_w, r' such that $(w', P'_w) \in S'[r']$. There are now two cases to consider: 1) $r' = r$; or 2) $r' \neq r$ and thus since r is chosen to be pred -maximal in G , we have $(r', r) \in G.\text{pred}$. In case (1), since $(w, P_w), (w', P'_w) \in S[r]$, and $S[r]$ is a set, we know that $w \neq w' \vee P_w \neq P'_w$. On the other hand, from the definition of G' we have $\text{sbrf}(G', r, w) = P_w$. It is then straightforward to demonstrate that $P'_w \neq \text{sbrf}(G', r, w) \Rightarrow P'_w \perp \text{sbrf}(G', r, w)$. As such, we have $w \neq w' \vee \text{sbrf}(G', r, w) \perp P'_w$, as required.

In case (2), from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $G.\text{rf}^{-1}(r') \neq w' \vee \text{sbrf}(G, r', w') \perp P'_w$. If the first disjunct holds $(G.\text{rf}^{-1}(r') \neq w')$, since $(r', r) \in G.\text{pred}$, from the definition of G' we have $G.\text{rf}^{-1}(r') = G'.\text{rf}^{-1}(r') \neq w'$, as required.

On the other hand, if the first disjunct does not hold and instead the second disjunct holds, we then have $G.\text{rf}^{-1}(r') = w' \wedge \text{sbrf}(G, r', w') \perp P'_w$. Since $(w', r') \in G.\text{rf}$, $(r', r) \in G.\text{pred}$ and $S[r] \neq \emptyset$ ($(w, P_w) \in S[r]$), from $\text{inv}(\langle G, T, U, S, A \rangle)$ we have $G.\text{porf}^?; [w'] \times \{r\} \subseteq G.\text{pred}$. As such, from the definition of G' we know $w' \in G'$ and $\text{sbrf}(G, r', w') = \text{sbrf}(G', r', w')$. Consequently, since $\text{sbrf}(G, r', w') \perp P'_w$, we have $\text{sbrf}(G', r', w') \perp P'_w$, as required. $\square$

Theorem D.17 (Optimality). *For all G, T, U, A , if $\langle G_0, \emptyset, \emptyset, \emptyset, \emptyset \rangle \Rightarrow^* \langle G, T, U, \emptyset, A \rangle$ then:*

$$\forall G_1, G_2 \in A \cup \{G\}. G_1 \neq G_2 \Rightarrow G_1 \perp G_2$$

Proof. From the definition of $\text{inv}(\cdot)$, the $\text{inv}(G_0, \emptyset, \emptyset, \emptyset, \emptyset)$ holds vacuously. The desired result then follows from Lemma D.16 by straightforward induction on the number of steps in $\Rightarrow^*$. $\square$