



Effective Lock Handling in Stateless Model Checking

MICHALIS KOKOLOGIANNAKIS, MPI-SWS, Germany

AZALEA RAAD, MPI-SWS, Germany

VIKTOR VAFEIADIS, MPI-SWS, Germany

Stateless Model Checking (SMC) is a verification technique for concurrent programs that checks for safety violations by exploring all possible thread interleavings. SMC is usually coupled with Partial Order Reduction (POR), which exploits the independence of instructions to avoid redundant explorations when an equivalent one has already been considered. While effective POR techniques have been developed for many different memory models, they are only able to exploit independence at the *instruction level*, which makes them unsuitable for programs with coarse-grained synchronization mechanisms such as *locks*.

We present a lock-aware POR algorithm, *LAPOR*, that exploits independence at both *instruction* and *critical section levels*. This enables *LAPOR* to explore exponentially fewer interleavings than the state-of-the-art techniques for programs that use locks conservatively. Our algorithm is sound, complete, and optimal, and can be used for verifying programs under several different memory models. We implement *LAPOR* in a tool and show that it can be exponentially faster than the state-of-the-art model checkers.

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

Additional Key Words and Phrases: Model checking, mutual exclusion locks, weak memory models

ACM Reference Format:

Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. 2019. Effective Lock Handling in Stateless Model Checking. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 173 (October 2019), 26 pages. <https://doi.org/10.1145/3360599>

1 INTRODUCTION

Coarse-grained locking (CGL) is the most widespread form of synchronization in concurrent shared-memory programs. It is a fairly straightforward way of making a sequential program “thread-safe” with programming languages, such as Java, even providing syntactic support for synchronizing all the methods operating on a shared object (regardless of whether they access the same fields of the object or not). CGL often performs better than finer-grained schemes in cases of low contention or when “hardware lock elision” is available (e.g., on modern x86 processors). Lock elision enables concurrent execution of lock-protected code regions that do not conflict (e.g., access only disjoint memory locations), by optimistically executing them and rolling back in case a conflict is detected.

Intuitively, CGL simplifies reasoning about concurrent programs because it reduces the number of interleavings that a programmer or an automated verifier has to consider. Somewhat surprisingly, however, CGL has the opposite effect on verification using the state-of-the-art *stateless model checking* (SMC) tools (e.g., [Abdulla et al. 2014, 2018; Aronis et al. 2018; Kokologiannakis et al. 2017,

Authors’ addresses: Michalis Kokologiannakis, MPI-SWS, Saarland Informatics Campus, Germany, michalis@mpi-sws.org; Azalea Raad, MPI-SWS, Saarland Informatics Campus, Germany, azalea@mpi-sws.org; Viktor Vafeiadis, MPI-SWS, Saarland Informatics Campus, Germany, viktor@mpi-sws.org.



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

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

2475-1421/2019/10-ART173

<https://doi.org/10.1145/3360599>

2019]). More specifically, CGL increases the verification cost significantly because it inhibits the most important optimization of SMC, which makes the approach scalable, namely *(dynamic) partial order reduction* (POR) [Abdulla et al. 2014; Flanagan and Godefroid 2005].

SMC with POR explores all possible program executions (i.e., thread interleavings under sequential consistency) up to a fixed equivalence relation, where two interleavings are deemed equivalent if one can be obtained from the other by commuting pairs of independent actions (e.g., accesses to different locations). Since accesses of different threads are most often independent, POR leads to a huge reduction in the space that needs to be explored, which in turn makes SMC with POR scalable.

In programs with coarse-grained locks, however, existing POR algorithms achieve very little reduction in the state space because they treat locking commands as normal instructions. Since acquisitions of the same lock cannot be commuted (because it would affect which thread acquires the lock), existing SMC/POR tools consider all the possible ways in which critical sections of the same lock could be interleaved. This leads to a lot of redundant exploration because the relative order of critical sections is often irrelevant—it does not affect the program outcome.

To make SMC scalable for coarse-grained locking programs, we develop the *Lock-Aware Partial Order Reduction* (LAPOR) algorithm, which treats lock acquisitions in each program execution as independent unless their respective order can be deduced from the memory accesses of the execution. As such, although our algorithm works at the instruction level, it also exploits independence at the critical section level, thus achieving exponential reductions in the number of explored executions compared to the state-of-the-art model checking tools.

In more detail, this paper makes the following contributions.

- (§2) We present an intuitive account of LAPOR through a series of examples, and show how the lock acquisition order can be inferred on the fly when necessary, without exploring all possible orderings in advance.
- (§3) We present the formalism underpinning LAPOR. Our algorithm is parametric both in the choice of the memory model and in when two concurrent writes are deemed conflicting.
- (§4) We describe LAPOR in detail, and show that it is sound (produces no false positives), complete (explores all possible behaviours), and optimal (explores each behaviour exactly once), and that it can be applied to any memory model that satisfies a few basic assumptions.
- (§5) We implement LAPOR over the state-of-the-art GENMC model checker [Kokologiannakis et al. 2019] and evaluate its performance. We show that LAPOR is significantly faster than the state of the art for programs with coarse-grained locks, while incurring only a small overhead for programs with fine-grained locks, and negligible overhead for programs without locks.

2 OVERVIEW

2.1 Stateless Model Checking and Dynamic Partial Order Reduction

Stateless Model Checking (SMC) is an effective concurrency verification technique based on the observation that one can verify a concurrent program by exploring all possible interleavings of its threads¹. However, as exploring all interleavings quickly becomes non-scalable even for small programs, *partial order reduction* (POR) techniques are often employed alongside SMC. Given a suitable notion of *instruction independence*, POR techniques deem two interleavings equivalent if one can be obtained from the other by swapping adjacent, independent execution steps. POR techniques then explore (at least) one interleaving from each equivalence class—in the context of SC, these equivalence classes are known as Mazurkiewicz traces [Mazurkiewicz 1987].

¹Although the approach presented in this paper is parametric within the choice of the memory model (see §3), in this section we assume a setting of sequential consistency (SC) [Lampert 1979] for ease of presentation.

Amongst POR techniques, Dynamic Partial Order Reduction (DPOR) [Flanagan and Godefroid 2005] stands out. Unlike earlier POR techniques, DPOR does not rely on static over-approximations to infer whether events are independent. It instead keeps track of the conflicting events along the exploration path, and uses them to guide further explorations.

To make this concrete, consider the following program with N concurrent reads.²

$$a_1 := x \parallel \cdots \parallel a_N := x \quad (N\text{-READS})$$

While this program has $N!$ interleavings, almost all (D)POR techniques will explore only one interleaving by detecting that the read operations are independent (they return the same value irrespective of their ordering) and hence all interleavings are deemed equivalent.

Each equivalence class of interleavings can be represented as an *execution graph*, such as the one shown in Fig. 1 for the N -READS program. The nodes of an execution graph denote execution events which correspond to program instructions, while its edges describe different relations on events. More specifically, the solid black edges represent the *program-order* (po) relation, which totally orders the events in each thread and orders the initialization events before all other events. The dashed edges represent the *reads-from* (rf) relation, which relates each write event to the read events that read from it and obtain its value. rf is total and functional on its second argument: every reads reads from precisely one write. In the graph of Fig. 1, all read events read from the initialization event. Note that the reads are not related to one another as they are independent.




Fig. 1. Execution graph.

The objective of SMC is to explore all execution graphs of a given program in a systematic way, while keeping only one graph in memory at any given time. In the remainder of this article we use the terms ‘execution graphs’, ‘executions’, ‘graphs’ and ‘equivalence classes’ interchangeably.

2.2 Locks in Stateless Model Checking

Although the notion of instruction independence is well-studied for memory accesses (reads and writes), exploiting independence is much more difficult in the presence of locks. To see why, consider a variant of the N -READS program where each read is wrapped in a critical section using locks:

$$\begin{array}{l} lock(l); \\ a_1 := x; \\ unlock(l) \end{array} \parallel \cdots \parallel \begin{array}{l} lock(l); \\ a_N := x; \\ unlock(l) \end{array} \quad (N\text{-READS-LOCK})$$

For such locking programs, existing (D)POR techniques take an either *fine-grained* or *coarse-grained* approach to locks.

Fine-grained Approaches. Fine-grained approaches consider interleavings at the instruction level, treating $lock()$ and $unlock()$ operations as individual instructions. In the case of N -READS-LOCK, the lock operations are considered as dependent because they access the same lock and cannot be reordered. This yields $N!$ distinct equivalence classes capturing the order in which the lock is acquired. That is, existing (D)POR techniques explore $N!$ interleavings (one per execution) instead of one as in the case of N -READS.

More concretely, for $N = 2$, assume without loss of generality that a (D)POR algorithm first obtains execution **A** in Fig. 2, where the left thread acquires the lock first. The algorithm subsequently detects that the two threads compete to obtain the same lock; it thus reverses the order in which the lock is acquired, yielding execution **B**. As such, although adding locks to N -READS does not alter the program behaviour, it results in additional *unnecessary* executions.

²In our examples, we use a, b, c for local variables, and x, y, z for global (shared) variables, which are initialized with 0.




Fig. 2. Executions of the N -READS-LOCK program for $N = 2$ under a classic POR algorithm.

The reason for this is that traditional (D)POR algorithms treat $lock()$ and $unlock()$ operations as memory accesses (reads and writes). More specifically, $lock()$ is assigned read semantics, $unlock()$ is assigned write semantics, and rf is required to be one-to-one on lock/unlock events. This characterization of locks captures the order of lock acquisition elegantly: when the lock operation of a critical section cs_2 reads from the unlock operation in critical section cs_1 , then cs_1 is ordered before cs_2 . That is, the rf relation induces a total order on critical sections and thus determines the lock acquisition order. As lock operations participate in rf , SMC algorithms always consider the lock operations as dependent. As we have seen, this can then lead to an exponential increase in the number of executions, hindering performance needlessly.


This can be very surprising to programmers using SMC to test their code because it contradicts the intuition that the behaviour of concurrent programs is *monotonic* with respect to synchronization: adding more synchronization (e.g., locks) to a program should not introduce additional behaviours. As such, when inserting locks in the N -READS program to obtain N -READS-LOCK, the possible behaviours of N -READS-LOCK must be included in those of N -READS. As SMC techniques are used to explore the possible behaviours of a program by considering all its executions, under the monotonicity intuition one then does not expect N -READS-LOCK to have more executions than N -READS. However, while recent SMC techniques generate only one execution for N -READS (ignoring the read order), they generate $N!$ executions for the N -READS-LOCK variant (considering all lock orderings), thus breaking the monotonicity property.

Coarse-grained Approaches. In coarse-grained approaches, (D)POR techniques treat critical sections as *atomic* blocks, and only consider interleavings at the level of critical sections. As such, for the N -READS-LOCK above, these techniques correctly identify the $N!$ possible interleavings of the critical sections as equivalent, thus generating one execution. Such atomic treatment of critical sections may improve performance significantly in *conservative* settings where all accesses to shared variables are enclosed within critical sections (e.g., N -READS-LOCK). However, in *mixed* settings where shared variables may be accessed both within and without critical sections, such atomic treatment is *unsound* as it fails to explore all executions.

To see this, consider the conservative program below (left) and its mixed variant (right):

$\begin{array}{l} lock(l); \\ x := 1; \\ x := 2; \\ unlock(l) \end{array} \parallel \begin{array}{l} lock(l); \\ a := x; \\ unlock(l) \end{array} \quad (\text{WW+R-CONS})$	$\begin{array}{l} lock(l); \\ x := 1; \\ x := 2; \\ unlock(l) \end{array} \parallel \begin{array}{l} a := x \end{array} \quad (\text{WW+R-MIXED})$
---	--

The **WW+R-CONS** program has two possible executions: one where the left critical section is executed first and thus $a = 2$, another where the right critical section is executed first and thus $a = 0$. The **WW+R-MIXED** program also has a third possible execution. Since the right thread does not acquire the lock l , the critical section in the left thread need not execute atomically with respect to the right thread, and so the read may interleave between the two writes yielding the outcome $a = 1$.

Fig. 3. Full exploration for the N -READS-LOCK program.

When treating critical sections atomically, (D)POR techniques cannot distinguish between $\mathbf{ww+r-cons}$ and its mixed variant $\mathbf{ww+r-mixed}$, failing to explore the valid execution in which $a = 1$. As such, atomic treatment of critical sections is unsound in the mixed setting.

2.3 LAPOR: Keeping Locks Unordered

LAPOR follows the fine-grained approach of considering interleavings at the level of instructions. It therefore records each individual instruction as an event in an execution. In order to avoid redundant executions due to locks, however, unlike classical (D)POR approaches, LAPOR does not assign read/write semantics to locks; nor does it record the total lock acquisition order. Instead, LAPOR keeps the lock operations *unordered*, and infers the lock order on the fly, only *when necessary*.

Let us demonstrate LAPOR via the N -READS-LOCK example. As shown in Fig. 3, starting from an empty execution (containing only the initialization events), LAPOR inspects the program and adds execution events one at a time, thus arriving at a full execution where the locks are unordered. Since there is only one reads-from option (i.e., the initialization event) for the reads in each thread, and since locks do not participate in \mathbf{rf} , LAPOR explores *only one* execution rather than $N!$ executions.



2.4 LAPOR: Inferring Lock Orderings On the Fly

In the N -READS-LOCK program above, the order of lock acquisition does not matter as the locks do not alter the behaviour of the program. This, however, is not always the case because locks may induce mutual exclusion constraints on the execution. In such cases, LAPOR *infers* the necessary lock orderings on the fly. To demonstrate this, consider the example below:

$$\begin{array}{l|l}
 \text{lock}(l); & \text{lock}(l); \\
 x := 42; & a := x; \\
 y := 42; & b := y; \\
 \text{unlock}(l) & \text{unlock}(l)
 \end{array} \quad (\mathbf{ww+rr-LOCK})$$

Here, the order of lock acquisition matters because the values read from x and y depend on the order in which the locks are acquired: when the left thread acquires the lock first, the right thread reads 42 for x and y ; when the right thread acquires the lock first, the right thread reads 0 (the initialization value) for x and y . Let us describe how LAPOR infers the lock ordering in $\mathbf{ww+rr-LOCK}$.

As before, starting from the empty graph (containing initialization events only), LAPOR incrementally adds all events of the first thread, as well as the lock event of the second thread, arriving at execution ① of Fig. 4. LAPOR then proceeds to add the event associated with the $a := x$ read. Inspecting the execution so far, LAPOR notes that this read may read either 0 or 42, and thus both possibilities must be explored, as depicted in executions ② and ③ of Fig. 4. However, rather than exploring both executions at once, LAPOR proceeds with one of the executions (e.g., ②), and pushes the alternative execution into a *work set* to be explored at a later time.

Fig. 4. Key exploration step for **WW+RR-LOCK**.Fig. 5. Key exploration step for **WW+RR-LOCK** (continued).

Although reading 0 and 42 are both consistent options, reading either value induces a certain lock ordering, as discussed above. That is, the order of lock acquisition matters because the two critical sections are *not independent*. We refer to two such critical sections as *conflicting*. For example, in **WW+RR-LOCK**, the two critical sections are read-write conflicted because the right thread reads from location x , and the left thread also writes at that memory location. As we describe shortly, LAPOR detects conflicting critical sections, infers their induced lock orderings, and tracks them as an additional relation in the execution, namely the *locks-before* relation: **lb**.

In the case where $a := x$ reads 0 (in ②), LAPOR proceeds as follows. First, it adds a *reads-before* (a.k.a. from-read) edge (**rb**) from $a := x$ to $x := 42$, as depicted in ②. Intuitively, **rb** edges relate each read r (e.g., $a := x$) to all the writes (e.g., $x := 42$) that overwrite the value read by r . As such, the presence of an **rb** edge between the critical sections of distinct threads indicates a conflict: both critical sections access the same location (e.g., x) and one is overwriting the value read by the other.

Next, LAPOR determines the induced lock ordering as prescribed by **rb**: when there is an **rb** edge from critical section cs_1 to cs_2 , then cs_1 must have acquired the lock before cs_2 , and thus there is an **lb** edge from cs_1 to cs_2 . This is illustrated in ② (Fig. 4) by the **lb** edge from the critical section in the right thread to that in the left.

Analogously, the presence of an **rf** edge between the critical sections of distinct threads indicates a conflict: both critical sections access the same location and one reads a value written by the other. As such, the two critical sections in ③ (Fig. 4) are deemed conflicted and LAPOR adds an **lb** edge from the critical section in the left thread to that in the right.

Continuing from ②, LAPOR next adds the $b := y$ read, which may once again read either 0 or 42, as depicted in graphs ④ and ⑤ of Fig. 5). Note that if locks did not enforce mutual exclusion, both read options would be valid. However, since locks do enforce mutual exclusion, reading 0 is

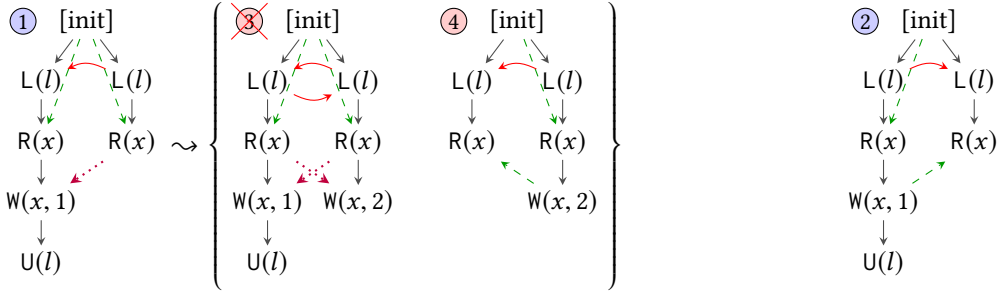


Fig. 6. Key exploration step for 2RW-LOCK.

the only valid option as reading 42 violates mutual exclusion due to the lock ordering imposed by 1b. Let us now discuss how LAPOR detects this mutual exclusion violation.

In the case where $b := y$ reads 42, i.e., there is an **rf** edge from $y := 42$ to $b := y$ as shown in 5, LAPOR detects a conflict between the two critical sections, and thus adds an **lb** edge from the critical section in the left thread to that in the right. This then leads to an **lb** cycle between the two critical sections, which means that the resulting graph is inconsistent, because a critical section cannot be simultaneously ordered both before and after another one. Thus reading 42 is discarded as invalid. Subsequently, the execution in 4 is completed by adding the unlock event of the right thread, thus concluding the first run of LAPOR.

Lastly, LAPOR completes execution 3 in Fig. 4 following analogous steps: it adds the $b := y$ read, excludes 0 as a possible reads-from value for $b := y$ (otherwise we obtain an **lb** cycle resulting in an inconsistent execution), and arrives at the execution where both $a := x$ and $b := y$ read 42.

2.5 LAPOR: Excluding Inconsistent Executions

In §2.4 we showed that when adding a read r to the graph, some reads-from options for r may lead to inconsistent executions. Nevertheless, a read can always be safely added to the graph as there is at least one reads-from option that preserves consistency. However, as we demonstrate here, this is not the case when adding a write to the graph; i.e., the mere addition of a write to the graph may render it inconsistent. Consider the example below:

$$\begin{array}{l}
 \text{lock}(l); \\
 a := x; \\
 \text{if } a = 0 \text{ then } x := 1; \\
 \text{unlock}(l)
 \end{array}
 \parallel
 \begin{array}{l}
 \text{lock}(l); \\
 b := x; \\
 x := 2; \\
 \text{unlock}(l)
 \end{array}
 \quad (2\text{RW-LOCK})$$

Starting from the empty graph with initialization events, LAPOR first adds $\text{lock}(l)$ and $a := x$ in the left thread. The $a := x$ read can only read from the initialization event since the graph contains no other writes. This makes the condition of the **if** statement true, enabling the addition of $x := 1$. LAPOR then adds $\text{unlock}(l)$ in the left thread, as well as $\text{lock}(l)$ in the second thread. It next adds the $b := x$ read, detecting that there are two available reads-from options (either from the initialization event, or the $x := 1$ write in the first thread), as depicted in executions 1 and 2 of Fig. 6.

Observe that as depicted in 1 (resp. 2), the choice of reads-from option for $b := x$ results in an **rb** (resp. **rf**) edge between $x := 1$ and $b := x$, which in turn induces an order on the two critical sections, as discussed earlier. As before, LAPOR continues with one of the executions (e.g., 1), and the other execution (e.g., 2) is added to a work set for later exploration. Continuing from 1, LAPOR next adds the $x := 2$ write, which brings up two issues that require additional care.

First, the addition of the $x := 2$ write results in an **rb** edge from $a := x$ to $x := 2$, which in turn induces an **lb** edge from the left critical section to the right, as depicted in ③. This then leads to an **lb** cycle, rendering the execution inconsistent; as such, graph ③ is not explored further.

Second, as discussed above, previously when $a := x$ was added, there was only one read-from option: the initialization event. However, note that in an interleaving of **2RW-LOCK** where the right thread is executed before the left, it is possible for $a := x$ to read from $x := 2$ in the right thread. As such, we must ensure that this alternative interleaving is also explored. To do this, when adding a write w (e.g., $x := 2$), regardless of whether adding w makes the graph inconsistent, we check if w may *revisit* an existing read r in the graph (e.g., $a := x$), in that r may read from the newly added w , leading to an additional execution (e.g., ④ in Fig. 6).

When this is the case, following Abdulla et al. [2018] and Kokologiannakis et al. [2019], we only keep the following events in this additional execution: (1) r itself (e.g., $a := x$), (2) the events that were added *before* r (e.g., $lock(l)$ in the left thread), (3) w itself (e.g., $x := 2$), and (4) the events that led to w (e.g., $lock(l)$ and $b := x$ in the right thread). In other words, we remove all events that were added *after* r but did not lead to w (e.g., the $x := 1$ and $unlock(l)$ in the left thread). This is because the new value read by r (from w) may affect the control flow, and lead to different events in the graph. For example, in this new execution where $x := 2$ revisits $a := x$ (i.e., $a := x$ reads from $x := 2$), the condition of the **if** statement is no longer true, and thus $x := 1$ is no longer part of the execution, and must be removed from the graph, as depicted in ④.

Once again, the new execution ④ is added to the work set for later exploration, and thus the work set contains ② and ④ at this point. As the current execution ③ was dropped due to inconsistency, LAPOR proceeds by exploring the remaining executions in the work set.

Exploring ② is straightforward and is simply completed by adding the remaining events of the right thread, yielding an execution where $a = 0$ and $b = 1$. Analogously, continuing from ④, the $x := 1$ write is no longer added (as the conditional fails), and the exploration is completed by adding the $unlock(l)$ event of the left thread, yielding an execution where $a = 2$ and $b = 0$.

Remark 1. Note that when adding $b := x$ reading 0 (from the initialization event) in graph ③, LAPOR cannot know in advance that this leads to an inconsistent execution at the next step (after adding $x := 2$). This is because LAPOR does not know of the next events that are added to the graph in due course. In particular, if $b := x$ were the only event in the critical section of the right thread, reading 0 would be a valid option and would not lead to an inconsistent execution. However, were we to disallow $b := x$ reading 0 (fearing that it may lead to inconsistency), we would inadvertently miss the *consistent* execution in which both $a := x$ and $b := x$ read 0 from the initialization event.

2.6 LAPOR: Parametric Treatment of Write-Write Conflicts

In the previous sections we showed how LAPOR uses **rb** and **rf** to detect *read-write* conflicts, i.e., when one critical section writes to a location read by the other. We next describe how LAPOR deals with *write-write* conflicts, i.e., when two critical sections write to the same location without reading from it. Consider the **w+w+LOCK** example below (left), whose two critical sections are conflicting because they both access x and at least one of them is a write.

$$\begin{array}{l} lock(l); \\ x := 1; \\ unlock(l) \end{array} \parallel \begin{array}{l} lock(l); \\ x := 2; \\ unlock(l) \end{array} \quad (w+w+LOCK) \qquad x := 1 \parallel x := 2 \quad (w+w)$$

The **w+w+LOCK** program has two interleavings: one where the left thread acquires the lock first, another where the right thread acquires the lock first. The order in which the two critical sections execute determines the final value of x , but does not affect any observation made by the program (since the program does not read x). As such, it suffices for a model checker to explore only one of




Fig. 7. Executions of the $w+w+LOCK$ program via mo (left) and wb (right)

the two interleavings. LAPOR can achieve this, but it depends on how exactly it is instantiated to order concurrent writes. As we will shortly see, LAPOR is parametric in its treatment of write-write conflicts and will explore one or two interleavings depending on the underlying DPOR algorithm used for ordering normal memory accesses.

To explain this, let us look at how existing DPOR algorithms handle the $w+w$ example above (right), obtained from $w+w+LOCK$ by removing the lock operations.

Classical DPOR algorithms [Abdulla et al. 2014; Flanagan and Godefroid 2005] consider the two writes on x conflicting, and thus explore two interleavings: one where $x := 1$ executes first and another where $x := 2$ executes first. To distinguish between the two interleavings, execution graphs in classical DPOR algorithms (in addition to po and rf) record the ‘modification-order’, mo , which totally orders the writes on each location. Execution graphs recording mo are often called Mazurkiewicz traces [Mazurkiewicz 1987] or Shasha-Snir traces [Shasha and Snir 1988] in the literature. LAPOR can be instantiated with such execution graphs by also using mo to infer the induced lock ordering in the same way as with rf and rb . For the $w+w+LOCK$ program, this leads to two executions as depicted in Fig. 7 (left).

However, several more recent DPOR algorithms [Abdulla et al. 2018; Aronis et al. 2018; Chalupa et al. 2017; Kokologiannakis et al. 2017, 2019] do not consider all orderings on the writes of the same location *unless* this order is *observable*. As such, since the value of x is not read in $w+w$, the order of the writes cannot possibly be observable, and thus these techniques explore only one execution for $w+w$, ignoring the order in which the writes are executed.

To do this, several such techniques compute the ‘writes-before’ relation, $wb \subseteq mo$, which (similarly to lb) *infers* the order between writes. That is, wb orders writes on the same location only when necessary. For example, consider the execution of Fig. 8. In this execution, $W(x, 2)$ must *write before* $W(x, 3)$ as they are po -related. Moreover, $W(x, 1)$ must write before $W(x, 3)$: otherwise, $R(x)$ would read 1 because of coherence. However, observe that $W(x, 2)$ and $W(x, 1)$ are not ordered. Intuitively, so long as wb is respected for x , these can execute in either order as this order cannot be observed.




Fig. 8. The wb relation

LAPOR can be instantiated with such techniques by using wb instead of mo to induce the lock orderings between critical sections with write-write conflicts. For the $w+w+LOCK$ example above, the order between the two writes is unnecessary; thus the recorded wb relation is empty, leading to a single execution as in Fig. 7 (right).

This parametricity of LAPOR is a significant part of our contribution: we reduce the number of explorations exponentially by recording the lock ordering when necessary by (1) identifying read-write conflicted critical sections, and (2) taking advantage of the underlying (D)POR technique and utilizing its notion of conflicts to detect write-write-conflicted critical sections.

3 FORMAL MODEL

We present the formalism underpinning LAPOR.

Labels, Events and Executions. As discussed in §2, in the literature of declarative (a.k.a. axiomatic) concurrency models, the traces of a program are typically represented as a set of execution graphs, where the graph nodes denote execution *events*, and graph edges capture the sundry relations on events. Each event corresponds to the execution of an instruction. An event is a tuple of the form $\langle i, n, l \rangle$, where $i \in \text{Tid} \uplus \{0\}$ 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 $l \in \{\text{error}\} \uplus \text{Lab}$ 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) an *instruction* label $l \in \text{Lab}$. As the set of instructions is memory model (MM) specific, the set of event labels is also MM-specific. As such, we keep our definitions parametric in the choice of MM, and assume a set of *labels*, Lab , associated with the instructions of the underlying MM. For instance, under the SC memory model, the store instruction $x := 1$ is associated with the label $W(x, 1)$. To model instruction labels, we assume a set of memory locations, Loc , ranged over by x, y, z . We further assume a set of *read labels*, $\text{RLab} \subseteq \text{Lab}$, a set of *write labels*, $\text{WLab} \subseteq \text{Lab}$, a set of *lock labels*, $\text{LLab} \subseteq \text{Lab}$, and a set of *unlock labels*, $\text{ULab} \subseteq \text{Lab}$, associated with *read* (load), *write* (store), *lock* and *unlock* instructions, respectively. For instance, the SC label $W(x, 1)$ is a write label.

Definition 3.1 (Events). Given a set of labels Lab , an *event* $e \in \text{Event}$ is a tuple $\langle i, n, l \rangle$, where $i \in \text{Tid}$ is a thread identifier, $n \in \mathbb{N}$ is a sequence number, and $l \in \{\text{error}\} \uplus \text{Lab}$ is a label.

We typically use a, b and e to range over events. The functions tid , sn , and lab project the thread identifier, the sequence number, and the label of an event respectively. We assume a set of locations Loc , and assume that functions loc , val_r and val_w respectively project the location, the read value and the written value of a label, where applicable. For instance, $\text{loc}(l)=x$ and $\text{val}_w(l)=1$ for $l=W(x, 1)$. We lift the label functions loc , val_r and val_w to events, and write e.g., $\text{loc}(e)$ for $\text{loc}(\text{lab}(e))$, given an event e . We define the *read events* as $R \triangleq \{e \in \text{Event} \mid \text{lab}(e) \in \text{RLab}\}$; the write (W), lock (L) and unlock (U) events are defined analogously. We define *initialization events* as $\text{Event}_0 \triangleq \{e \in W \mid \text{tid}(e)=0\}$.

Definition 3.2 (Executions). An *execution* is a tuple $G=\langle E, rf \rangle$, where $E \subseteq \text{Event}$ is a sequence of *events*, and $rf \subseteq (E \cap W) \times (E \cap R)$ is the *reads-from* relation, which is total and functional on its second argument: every read event reads from exactly one write event.

The order of events in the sequence E denotes the order in which the events were added to G . Given an execution G , we write $G.E$ and $G.rf$ for its components, and write $G.R$ for $G.E \cap R$; similarly for $G.W$, $G.L$ and $G.U$. We write $G.E_0$ for $G.E \cap \text{Event}_0$; we write $G.E_i$ for $\{\langle i', -, - \rangle \in G.E \mid i = i'\}$; and write $G.po$ for the *program order* defined as follows:

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

Initialization events are *po*-before all other events, and events of the same thread are ordered according to their sequence number. Further, we define $\text{SetRF}(G, w, r)$ as graph obtained from changing the incoming *rf* edge of r to come from w . Formally, $\text{SetRF}(\langle E, rf \rangle, w, r) \triangleq \langle E, rf \setminus E \times \{r\} \cup \{\langle w, r \rangle\} \rangle$.

The Memory Model. LAPOR can be instantiated for any memory model M that satisfies some basic assumptions.




Fig. 9. Illustration of the cs_l and lb relations on an execution of $WW+RR$ -LOCK.

First, M should define a *consistency predicate* on executions, $cons_M(\cdot)$ saying which executions are allowed. The permitted behaviours of a given program P under M are described as the set of its *consistent* executions. For the correctness proofs, we require that $cons_M(\cdot)$ is *well-formed*, *prefix-closed*, and *lock-extensible* (see §4.2).

Second, we assume that M defines a *happens-before* relation, hb , that contains the po relation and may additionally contain some synchronization edges. For instance, in the case of C11, happens-before contains an edge from a release write w to an acquire read r when r reads from w . Consistency of an execution G must require that hb is a strict partial order (irreflexive and transitive).

The wo Relation. Recall from §2 that LAPOR infers the lock order between conflicting critical sections by computing the ‘locks-before’ relation, lb . As discussed in §2.6, LAPOR is parametric in its notion of write-write conflicts. To model this, we assume the existence of a ‘write order’ parameter, wo , describing when two writes on the same location are deemed conflicted. As discussed in §2.6, wo can then be instantiated by either (1) the ‘modification order’ relation mo , describing a strict *total* order on the writes of each location; or (2) the smaller ‘writes-before’ relation wb by Lahav and Vafeiadis [2015], described below.

Notation. Given a relation r and a set A , we write $r^?$, r^+ and r^* for the reflexive, transitive and reflexive-transitive closures of r , respectively. We write $dom(r)$ and $rng(r)$ for the domain and range of r , respectively. We write r^{-1} for the inverse of r ; $r|_A$ for $r \cap (A \times A)$; $[A]$ for the identity relation on A : $\{(a, a) \mid a \in A\}$; and $[a]$ for $\{(a, a)\}$. We write r_e to restrict r to its *external* subset, i.e., the r edges between different threads: $r_e \triangleq r \setminus (po \cup po^{-1})$. 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. When r is a strict partial order, we write $r|_{imm}$ for the *immediate* edges in r , i.e., $r \setminus r$; r . Given an event set E , we write E_x for $\{e \in E \mid loc(e)=x\}$. Finally, we write $\#$ for sequence concatenation.

The cs_l Relation. The ‘critical section’ relation, cs_l , relates events in the same critical section of lock l . To define it, we introduce an equivalence (reflexive, transitive and symmetric) relation cs_l on events contained in a critical section guarded by lock l . An event a belongs to a critical section guarded by lock l if it is po -after a $L(l)$ event e and there is no $U(l)$ event po -between e and a .

Definition 3.3. Given an execution G and a lock l , the *critical section relation on l* , $G.cs_l$, is defined as follows:

$$G.cs_l \triangleq \{(a, b) \mid \exists e \in G.L_l. e \xrightarrow{pocs_l^?} a \wedge e \xrightarrow{pocs_l^?} b\}$$

with $G.pocs_l \triangleq G.po \setminus (G.po^?; [G.U_l]; G.po)$.

For example, let us consider the complete execution ③ of the **ww+rr-lock** example of § 2.4 (Fig. 9). As can be seen (highlighted in yellow), the execution has two critical sections: cs_l relates all the events of the first thread to one another, and all the events of the second thread to one another.

The lb Relation. We move on to the **lb** relation, which we incorporate into the underlying model's **hb** relation. Recall from § 2 that an **rf** edge between two critical sections in different threads renders the two sections write-read conflicted, and thus induces an **lb** edge between the events of the two sections in the **rf** direction. Put formally, if a and b are events in a critical section guarded by l , and $(a, b) \in G.rf_e$, then for all a', b' such that $(a, a') \in cs_l$ and $(b, b') \in cs_l$ (i.e., a' and b' are in the same critical sections as a and b , respectively), we have $(a', b') \in lb$. It is sufficient to consider the inter-thread **rf** edges (rf_e): same-thread **rf** edges either belong to the same critical section or to po-ordered critical sections, in which case they are already ordered by **hb**.

Similarly, if a read of a critical section a reads from a write that is overwritten by a write in a critical section b of the same lock, then all events in a are **lb**-before the events in b . Formally, we define the 'reads-before' relation as: $G.rb \triangleq G.rf^{-1}; G.wo$, and require that $G.cs_l; G.rb_e; G.cs_l \subseteq G.lb$. Intuitively, **rb** relates each read r to the writes **wo**-after the write r reads from.

Analogously, in the case of write-write conflicted critical sections, the (parametric) **wo** relation induces **lb** edges in the **wo** direction. We thus define the **lb** relation as the smallest transitive relation that contains the orderings induced by the rf_e , rb_e and wo_e edges between critical sections.

Definition 3.4. Given an execution G , the *locks-before relation*, $G.lb$, is the smallest transitive relation that satisfies the following for all locations l :

$$G.cs_l; (G.rf_e \cup G.rb_e \cup G.wo_e); G.cs_l \subseteq G.lb$$

We adapt the definition of the underlying memory model, so as to include **lb** in the definition of **hb**. As a result, checking $cons_M(\cdot)$ ensures that **lb** is irreflexive.

Returning to the execution in Fig. 9, the **rf** edges induce an **lb** edge from the critical section of the first thread to that of the second thread. Technically, there is an **lb** edge from every event of thread 1 to every event of thread 2, but we only show one edge for brevity.

The wb Instantiation of wo. As discussed above, the **wo** parameter may be instantiated as either: (1) the 'modification order', $mo \triangleq \bigcup_{x \in Loc} mo_x$, where each mo_x is a strict total order on the writes of location x ; or (2) the 'writes-before' order, $wb \triangleq \bigcup_{x \in Loc} wb_x$, where each wb_x is a strict partial order on the writes of location x , satisfying certain conditions.

We next present the formal definition of wb_x in Def. 3.5, as defined by Lahav and Vafeiadis [2015]. In the original definition by Lahav and Vafeiadis [2015], the $G.hb$ relation corresponds to the 'happens-before' relation of the release-acquire fragment of the C11 memory model [Batty et al. 2011; Lahav et al. 2017].

Definition 3.5 (The wb relation). Given an execution G , its 'happens-before' relation $G.hb$ and a location x , if $G.hb$ is transitive and $G.po \cup G.lb \subseteq G.hb$, then the *writes-before relation on x* , $G.wb_x$, is defined as the smallest transitive relation that satisfies the following inequality:

$$[G.W_x]; \left((G.rf^*; G.hb; (G.rf^{-1})^?; G.rf^*) \setminus (G.rf^{-1})^* \right); [G.W_x] \subseteq G.wb_x$$

The *writes-before* relation is defined as: $wb \triangleq \bigcup_{x \in Loc} wb_x$.

Calculating hb, lb, and wb. Note that the definitions of **hb**, **lb**, and **wb** are mutually recursive. The definition of **wb** depends on **hb**, which includes **lb**, whose definition recursively depends on $wo \triangleq wb$. We therefore calculate the three relations together by fixpoint iteration. Starting with $lb = wb = \emptyset$ and **hb** as defined by the underlying memory model, we iteratively add to the **lb** and




Fig. 10. An example of the recursive calculation of **hb**, **lb**, and **wb**.

wb relations the LHS of the inequalities in Def. 3.4 and Def. 3.5 respectively, as well as any transitive edges, and extend **hb** to include **lb** until a fixpoint is reached.

Figure 10 shows an example of the fixpoint calculation. In the first iteration of the fixpoint, **wb** edges are added from the initializer events to all writes of the graph. (For brevity, Fig. 10 does not display these edges and instead shows the induced **rb** edge from $R(x)$ to $W(x, 42)$.) In the second iteration, the **rb** edge induces **lb** and **hb** edges from the events of the second thread to the events of the first thread. (Again for brevity, Fig. 10 shows only one such edge.) In the third iteration, a **wb** edge is created between the $W(y, 1)$ and the $W(y, 42)$, due to the previously created **lb/hb** edge. In the fourth (and last) iteration, no further edges can be added and so the fixpoint iteration finishes.

4 LAPOR: LOCK-AWARE PARTIAL ORDER REDUCTION

Our model checking algorithm for locks, LAPOR, can be built as an extension of any (D)POR algorithm, as long as it calculates its alternative exploration options at each step and not at the end of each execution. For concreteness, here we present a version of LAPOR built on top of the GENMC model checking algorithm by Kokologiannakis et al. [2019]. Basing LAPOR over GENMC enables our algorithm to be parametric with respect to the memory model and allows us to prove its correctness and optimality (see § 4.2).

Configurations. Recall from §2 that given a program P , LAPOR explores the executions of P one at a time, recording the alternative explorations encountered along the way in a *work set*. For instance, when adding a read event r to the current execution G , if r has several reads-from options in G , then one of these options is considered in G , and the remaining options are added to the work set for future exploration. As such, as in GENMC, LAPOR maintains a *configuration* of the form $\langle G, \Gamma \rangle$, where G denotes the current execution, and Γ denotes the work set. A work set Γ in GENMC is a tuple comprising several components in order to track alternative future explorations efficiently. However, as the representation details of work sets are not necessary for understanding LAPOR, here we represent it as an abstract component Γ , and refer the reader to Kokologiannakis et al. [2019] for more details.

The Main VERIFY Procedure. LAPOR begins exploring the executions of a program P by calling $\text{VERIFY}(P)$. This procedure creates an initial configuration comprising the initial graph G_0 (containing initialization events only) and an empty work set Γ . It then generates the executions of P one at a time by calling $\text{VISITONE}(P, G, \Gamma)$ on Line 4, which fully explores one execution extending G , and pushes alternative explorations encountered to the work set Γ . Once $\text{VISITONE}(P, G, \Gamma)$ returns a full execution, remaining executions are generated by exploring the options in Γ (Line 5).

Algorithm 1 Exploration algorithm

```

1: procedure VERIFY( $P$ )
2:    $\langle G, \Gamma \rangle \leftarrow \langle G_0, \emptyset \rangle$ 
3:   do
4:     VISITONE( $P, G, \Gamma$ )
5:   while  $\langle G, \Gamma \rangle \leftarrow \text{nextExp}(\Gamma)$ 

6: procedure VISITONE( $P, G, \Gamma$ )
7:   while  $\text{cons}_M(G)$  do
8:      $a \leftarrow \text{nextEvent}_P(G)$ 
9:     if  $a = \perp$  then
10:      output("Generated execution  $G$ .")
11:      return
12:     if  $a \in \text{error}$  then exit("Erroneous program.")
13:      $G.E \leftarrow G.E \# [a]$ 
14:     if  $a \in W$  then
15:        $\text{calcRevisits}(G, \Gamma, a)$ 
16:     if  $a \in R$  then
17:        $W \leftarrow G.E \cap W_{\text{loc}}(a)$ 
18:       choose some  $w_0 \in W$ 
19:        $G \leftarrow \text{SetRF}(G, w_0, a)$ 
20:        $\text{push}(\Gamma, \{\text{SetRF}(G, w, a) \mid w \in W \setminus \{w_0\}\})$ 

```

4.1 The VISITONE Procedure

The $\text{VISITONE}(P, G, \Gamma)$ procedure carries out the crux of the algorithm by exploring one execution at a time. The VISITONE procedure in Algorithm 1 is essentially that of GENMC, but has much more elaborate implementations of the $\text{cons}_M(\cdot)$ and the $\text{nextEvent}_P(\cdot)$ functions (see below).

The VISITONE procedure proceeds as follows. In each iteration of this loop, while the current graph G is M -consistent ($\text{cons}_M(G)$ holds), it is extended with its next event a , given by the $\text{nextEvent}_P(G)$ function described shortly. If $\text{nextEvent}_P(G)$ returns \perp , the current execution graph cannot be extended further because the program has either terminated or all the threads are blocked. In this case, VISITONE outputs the execution graph and terminates. If the next event a is an assertion violation, then an error is reported (Line 12), and the algorithm terminates. Otherwise, we add a to the graph by placing it at the end of the sequence of events (Line 13).


If the new event a is a write, we check whether it may revisit existing reads in the graph; i.e., if existing reads may read from a (see §2.5). This is done by calling $\text{calcRevisits}(G, \Gamma, a)$ on Line 15, which adds all possible revisit options to the work set Γ . For instance, in Fig. 6 of §2.5, upon adding the write event $W(x, 1)$ to ③, the algorithm detects that $W(x, 1)$ may revisit the read event $R(x)$ in the left thread as depicted in ④, and thus ④ is added to the work set. The $\text{calcRevisits}(G, \Gamma, a)$ procedure is that of GENMC and remains unchanged in our LAPOR extension. It is not necessary to understand the details of $\text{calcRevisits}(G, \Gamma, a)$ to understand LAPOR. We have thus omitted the $\text{calcRevisits}(G, \Gamma, a)$ procedure here and refer the reader to [Kokologiannakis et al. 2019].

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 (Line 17), choose a write w_0 for the current exploration (Line 18), set a to read from w_0 in G (Line 19), and push the remaining revisit options to Γ (Line 20).

Checking Consistency: The $\text{cons}_M(\cdot)$ Function. As mentioned in §3, $\text{cons}_M(G)$ includes a check that **hb** is irreflexive. Further recall that we have adapted the definition of **hb** to include **lb**, which in turn also depends on **wo** (which, in the case of $\text{wo} \triangleq \text{wb}$, recursively depends on **hb**). Therefore, as part of $\text{cons}_M(G)$, we calculate the three relations together via a fixpoint iteration as explained at the end of §3.

The nextEvent_P Function. As described in §2, LAPOR explores each execution of a given program incrementally by adding one event at a time, respecting the program-order (po) in each thread; i.e., the events of each thread are added in po order. To do this, LAPOR uses the nextEvent_P function of GENMC (Line 8) which returns the next event a to be added to G . More concretely, given a program P and an execution G of P , the $\text{nextEvent}_P(G)$ function in GENMC returns an event a of thread i in G such that: (1) a is the next event of i in $G.\text{po}$ (i.e., $a.\text{tid} = i$ and $a.\text{sn} = |G.E_i| + 1$); (2) i is not stuck (e.g., due to a failed **assume** statement); and (3) i has not finished execution. When no such event exists in G (i.e., all threads are stuck or finished), then $\text{nextEvent}_P(G)$ returns \perp .

In the presence of locks, defining nextEvent_P requires additional care. Recall the **WW+RR-LOCK** program in §2.4 and consider a run of LAPOR where the $\text{lock}(l)$ events and the $x := 42$ and $a := x$ events have been added to the execution, with $a := x$ reading from $x := 42$, as shown in ① below.



Note that as discussed in §2, the **rf** edge between $x := 42$ and $a := x$ induces an **lb** edge between the two critical sections. Given the conditions outlined above, the $\text{nextEvent}_P(G)$ may next choose to add either $y := 42$ or $b := y$. Let us assume that $\text{nextEvent}_P(G)$ chooses $b := y$. As depicted above in ②, since the execution thus far contains no other writes on y , the $b := y$ read has only one reads-from option: the initialization event. However, as shown in ③, reading from the initialization leads to an inconsistent execution: when $y := 42$ is eventually added to the graph, it results in an **rb** edge between $b := y$ and $y := 42$, which in turn induces an **lb** between the two critical sections, resulting in an **lb** cycle. As such, starting from execution ①, we fail to generate the execution in which $a := x$ and $b := y$ both read 42 from the first thread. In other words, the only way to extend ① consistently is for $b := y$ to read from $x := 42$, and thus $\text{nextEvent}_P(G)$ must add $x := 42$ before $b := y$. Intuitively, when two critical sections are **lb**-ordered, $cs_1 \xrightarrow{\text{lb}} cs_2$, the events in cs_1 may affect the reads-from options of those in cs_2 , and thus all events of cs_1 must be added before those of cs_2 .

More generally, we define $\text{nextEvent}_P(G)$ to prioritize the events of *open critical sections* (those without an *unlock* event), and to add such events in the **lb** order. That is, when $\text{nextEvent}_P(G)=a$, as well as the (1)-(3) conditions stipulated by GENMC outlined above, we require that: (4) either G contain no open critical sections, or a be an event in an **lb**-minimal open critical section:

$$G.\text{OCS}=\emptyset \vee a \in G.\text{OCS} \wedge \nexists b. b \in G.\text{OCS} \wedge (b, a) \in G.\text{lb}$$

$$\text{where } G.\text{OCS} \triangleq \bigcup_{x \in \text{Loc}} G.\text{OCS}_x \quad \text{with } G.\text{OCS}_x \triangleq \{e \mid [e]; G.\text{cs}_x \neq \emptyset \wedge [e]; G.\text{cs}_x; [G.U_x]=\emptyset\}$$

In our LAPOR development, we implement nextEvent_P to choose an event a in the left-most thread i (i.e., one with the smallest thread identifier) that satisfies conditions (1)-(4).

Note that if the executions of a program P are not well-formed in their acquisition of locks, then nextEvent_P may get stuck. For instance, consider a variant of the **WW+RR-LOCK** program in §2.4 without $\text{unlock}(l)$ in the left thread. Continuing from graph ① above, after adding the $y := 42$ event, the nextEvent_P function gets stuck: we cannot add the events of the right thread since the left critical section is open and is **lb**-before the right one.

To address this, when $\text{nextEvent}_P(G) = \perp$, LAPOR checks that the current execution G is *lock-well-formed*, in that *lock* and *unlock* events are paired. Put formally, let $G.\text{lpox} \triangleq (G.\text{po} \upharpoonright_{L_x \cup U_x})_{\text{imm}}$; a (full) execution G is lock-well-formed iff:

$$\forall x, a \in G.E. (a \in G.L_x \Rightarrow \exists u \in G.U_x. (a, u) \in G.\text{lpox}) \wedge (a \in G.U_x \Rightarrow \exists l \in G.L_x. (l, a) \in G.\text{lpox})$$

When a full execution G is not lock-well-formed, LAPOR reports an error.

4.2 LAPOR: Soundness, Completeness and Optimality

In this section, we establish the correctness of LAPOR by showing that its generated execution graphs satisfy three properties: *soundness*, *completeness* and *optimality*. Given a program P and a memory model m , we say that LAPOR generates execution G for P if running $\text{VERIFY}(P)$ outputs G (at Line 10 of VISITONE) among other executions. Soundness ensures that if LAPOR generates G for P , then G is consistent (i.e., $\text{cons}_m(G)$ holds); completeness ensures that if G is a consistent execution of P , then LAPOR generates G for P ; and optimality ensures that the executions of P generated by LAPOR are pairwise distinct.

The soundness proof is straightforward: LAPOR checks consistency after each step, dropping inconsistent executions (Line 7 of VISITONE); as such, it only outputs consistent executions. The optimality of LAPOR follows immediately from the optimality of its parent algorithm GENMC , as established in [Kokologiannakis et al. 2019]. That is, as LAPOR is implemented by extending GENMC 's consistency check and next event selection, and these extensions do not affect the optimality of the algorithm, the optimality of LAPOR follows straightforwardly from that of GENMC .

Establishing completeness is, however, more difficult. Even though GENMC is complete, this does not necessarily imply the completeness of its LAPOR extension. This is because in order to achieve completeness, GENMC requires the underlying memory model to be *well-formed*, *prefix-closed* and *extensible*. While the first two properties (described shortly) are preserved in the presence of locks, extensibility is not. In particular, extensibility requires that execution consistency be preserved under extension with events. That is, for all executions G and all events a , if G is consistent, then after adding a to G the execution remains consistent. As we discussed in §2.5, this is not necessarily the case in the presence of locks: the mere addition of a *write* event may render an execution inconsistent. In other words, the extensibility requirement of GENMC is too strong in the presence of locks. We thus weaken the extensibility requirement and show the LAPOR completeness under a weaker requirement: *lock-extensibility*. That is, if the underlying memory model is well-formed, prefix-closed and lock-extensible, then LAPOR is complete.

Well-formedness ensures that consistent executions do not contain $\text{po} \cup \text{rf}$ cycles: for all G , if G is consistent, then $G.\text{porf} \triangleq (G.\text{po} \cup G.\text{rf})^+$ is irreflexive. As argued by Kokologiannakis et al. [2019], this property precludes problematic executions exhibiting the ‘out of thin air’ behaviour. Moreover, well-formedness ensures that the execution consistency is independent of the order in which events are added: for all $G = \langle E, \text{rf} \rangle$, if G is consistent, then $\langle E', \text{rf} \rangle$ is also consistent for any permutation E' of E . This is because the order of events in $G.E$ constitutes auxiliary instrumentation used by the algorithm, and thus execution consistency must be agnostic to this order.

Prefix-closedness ensures that each execution G of a program can be generated by adding its event in any order that extends $G.\text{porf}$. As noted by Kokologiannakis et al. [2019], all well-known

memory models including SC [Lamport 1979], TSO [Owens et al. 2009], PSO [SPARC International Inc. 1994], and RC11 [Lahav et al. 2017], satisfy this property. This allows us to *fix* the order in which events are added via the nextEvent_P function (described above), and thus generate all program executions systematically. Put formally, the prefix-closedness states that 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 G to those events in E yields a consistent execution, i.e., $\langle E, G.\text{rf}|_E \rangle$ is consistent.

Recall from our discussion above (see the nextEvent_P function) that **lb**-ordered critical sections may affect the reads-from options of one another, thereby affecting consistency. That is, when the event addition order does not respect **lb**, consistency may be compromised. This is indeed the scenario that led to inconsistency in our **2RW-LOCK** example of § 2.5: even though the right critical section in execution ① of Fig. 6 is **lb**-ordered before the left one, its events are added after those of the left one, leading to inconsistency in ③. Lock-extensibility thus weakens the notion of extensibility to allow for such scenarios; i.e., we require extension to preserve consistency *only if* the construction of the execution thus far respects the **lb** order. That is, when extending an execution G with an event a (with $a.\text{sn} = |G.E_{a.\text{tid}}| + 1$), if G is consistent and either a is not in a critical section, or a is in an **lb**-maximal critical section, then lock-extensibility requires that the graph obtained by adding a to G be consistent. Note that when a is in a non- G .**lb**-maximal critical section cs , i.e., there exists cs' such that $cs \xrightarrow{\text{lb}} cs'$, then the cs' events in G were added before a , thereby violating the **lb** order; we thus do not require the resulting execution to be consistent. Put formally, lock-extensibility states that for all G , a and $G' = \text{add}(G, a)$, if G is consistent and $\exists l. a \in G'.\text{cs}_l \Rightarrow \nexists b. (a, b) \in G.\text{lb}$, then either (1) $a \notin G'.R$ and G' is consistent; or (2) there exists $w \in G'$ such that $G'.\text{rf}[a \mapsto w]$ is consistent.

THEOREM 4.1 (CORRECTNESS). *The LAPOR algorithm is sound and optimal. If the underlying memory model is well-formed, prefix-closed and lock-extensible, then LAPOR is complete.*

PROOF (SKETCH). Soundness follows immediately from the VISITONE algorithm, while optimality follows immediately from the optimality of GENMC in Kokologiannakis et al. [2019].

To show that LAPOR is complete, we must show that it generates all executions of a program P under a given memory model M that is well-formed, prefix-closed and lock-extensible. As discussed above and demonstrated in [Kokologiannakis et al. 2019], the well-formedness and prefix-closedness of M 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 LAPOR if we permute the order in which its events are added. That is, for each execution G obtained by adding events in the order $\sigma = e_1, e_2, \dots, e_n$, there exists some permutation σ' of σ , such that LAPOR adds events in the σ' order and arrives at G . In order to show that such a permutation exists, we may need to remove events and re-add them later (corresponding to the $\text{calcRevisits}(\cdot)$ step). As such, we must ensure that we preserve execution consistency after re-adding each event. This can always be done thanks to the lock-extensibility of M which ensures that LAPOR never gets stuck. More concretely, at each step when adding event a to G_1 and obtaining G_2 , either a is not in a critical section in which case consistency of G_2 is ensured by lock-extensibility, or a is in a critical section. In the latter case when a is in a critical section cs , we add a in a way that renders cs a *maximal* critical section in G_2 , and thus lock-extensibility ensures that G_2 is consistent. To ensure that cs is deemed a maximal critical section, we proceed as follows. If a is a read event, we pick its reads-from option to be a **wo**-maximal write in G_1 ; if a is a write event, we set a to be a **wo**-maximal event in G_2 ; and if a is a lock event, we set a to be an **lb**-maximal event in G_2 . Given the definition of **lb**, this then ensure that the critical section of a , namely cs , is a maximal critical section in G_2 .

The detailed completeness proof is given in the Appendix. \square

5 EVALUATION

To evaluate the performance of our approach, we have implemented LAPOR as a verification tool for C programs that use C11 atomics and the pthread library for concurrency. More specifically, our LAPOR implementation is an extension of the open-source GENMC model checker [Kokologiannakis et al. 2019], available at <https://github.com/MPI-SWS/genmc>. We note that since GENMC supports a DPOR algorithm that computes `wb` for ordering writes on the same memory location, we have instantiated LAPOR based on this DPOR with `wb`.

In the following, we compare LAPOR against the DPOR algorithm that GENMC employs. We avoid direct comparisons with other SMC tools for two reasons. First, since LAPOR is built over GENMC, the direct comparison between the two tools better demonstrates the pros and cons of the lock-aware approach, while including other tools in the tables would dilute the message. Second, GENMC has already been extensively evaluated against other SMC tools (see Kokologiannakis et al. [2019]) and demonstrated to outperform them in most cases. Although we do not report the measurements, we have confirmed that the same holds for the benchmarks of this paper. In particular, because of GENMC's optimality with respect to a coarse equivalence relation, none of the existing SMC tools can be faster except perhaps by a polynomial factor.

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

5.1 Benchmark Selection

To evaluate the benefit of our approach over the baseline GENMC, we have used a diverse set of benchmarks.

First (§5.2), in order to measure the overhead of calculating the `lb` relation, we have chosen a set of benchmarks where we expect that LAPOR will have no benefit over GENMC (i.e., the coarser-grained equivalence partitioning of LAPOR will coincide with that of GENMC). For this purpose, we took a standard set of benchmarks from the TACAS competition on Software Verification [SV-COMP 2019] that have been used extensively in the DPOR literature by many different tools [Abdulla et al. 2015, 2016, 2018; Alglave et al. 2013; Huang and Huang 2016]. These benchmarks use locks scarcely and only in a very fine-grained fashion, and so we do not expect any performance improvement by using LAPOR over GENMC.

More specifically, we took the benchmarks from two different concurrency categories: pthread and pthread-atomic. From these, we excluded those that do not contain locks (as they are not interesting for our purposes), as well as those that are not data-deterministic (because data-determinism is a prerequisite for stateless model checking). For the remaining 11 benchmarks, we modelled verifier directives that indicate certain sections should be executed atomically (e.g., the `__VERIFIER_atomic_begin()` and `__VERIFIER_atomic_end()` primitives) by having our tool acquire/release a designated lock accordingly. We note that these benchmarks are small C programs (50-100 LoC) with the exception of `scull`, a toy Linux-kernel driver (~455 LoC).

Next (§5.3), in order to evaluate the pros and cons of our approach, we shift our focus to more synthetic benchmarks. We used benchmarks similar to the ones presented in §2 (e.g., `N-READS`, `N-READS-LOCK`, `WW+RR-LOCK`), both with and without lock-induced synchronization, and we evaluate LAPOR in both versions.

These synthetic benchmarks show (1) that our tool can explore *exponentially fewer* executions than GENMC by exploiting independence at the critical-section level, and (2) that the introduction of additional synchronization in a program does *not* have a significant impact in LAPOR. They also

Table 1. Benchmarks from [SV-COMP 2019]. The LB column is the loop bound used.

	LB	Executions		Time (s)	
		<i>GENMC</i>	<i>LAPOR</i>	<i>GENMC</i>	<i>LAPOR</i>
fib_bench	5	16 632	5922	11.60	3.83
indexer	4	64	64	0.22	2.60
lazy	-	6	6	0.04	0.04
queue_ok	10	2	2	0.04	0.28
stack-1	5	252	252	0.16	0.61
stateful01-2	-	6	4	0.04	0.04
triangular	5	16 632	5922	11.03	2.03
gcd	3	12	12	0.04	0.06
read_write_lock	-	96	96	0.08	0.14
scull	-	1749	922	0.25	0.23
time_var_mutex	-	2	2	0.04	0.03

provide further measurements showing how the calculation of **lb** impacts the verification time both in benchmarks that have conflicting critical sections, and in benchmarks that do not.

Finally (§5.4), we provide a more realistic workload for evaluating the performance of LAPOR, by running it on a few coarse-grained and fine-grained implementations of concurrent data structures under two different workloads. The data structures we considered are: (1) concurrent sets implemented as a doubly-linked list, (2) concurrent binary search trees (BSTs), and (3) concurrent hashtables. The coarse-grained versions use a global lock to protect the entire data structure, while the fine-grained versions use hand-over-hand locking for the sets and the BSTs, and per-bucket locking for the hashtables. Our first workload includes only searching, and the second includes both searching and updating. For the implementations requiring a linked list, we ported the linked list implementation used in the Linux kernel.

As we will see, LAPOR achieves exponential reductions not only for synthetic benchmarks, but also for these more realistic workloads, and is able to exhaustively verify client programs of such data structures within a few seconds.

5.2 SV-COMP Benchmarks

We start with the standard benchmarks from SV-COMP [2019], whose results are summarized in Table 1. For each benchmark, the table reports both the number of executions explored by each tool and the time taken to generate them. The LB column contains the loop unrolling bound used where applicable.

As shown in Table 1, the performance of LAPOR is comparable to (and sometimes better than) that of GENMC in these benchmarks. In most cases, GENMC performs somewhat better than LAPOR, but there are several cases where LAPOR explores considerably fewer executions than GENMC, and is therefore much faster.

Let us first focus on benchmarks where GENMC is faster than LAPOR, and explain why this is the case (see the indexer and stack-1 benchmarks). In general, in cases where both tools explore the same number of executions, LAPOR has additional overhead compared to a traditional DPOR technique, as expected. This overhead is induced by the calculation of **lb**, which can be considerable or insignificant depending on the program under test.

As mentioned in §3, **lb** includes the parametric **wo** relation that orders writes, which in our implementation is instantiated as **wb**. However, **wb** includes **hb**, and **hb** in turn includes **lb**. This

means that in order to calculate **lb** (necessary for the consistency checks), we have to perform a fixpoint calculation. This calculation can be relatively expensive (at least $O(n^3)$ in the number of critical sections and/or conflicting write accesses), and is performed repeatedly. Our LAPOR implementation recalculates **lb** from scratch at each step, whereas GENMC checks consistency in a more incremental fashion. This is reflected in the performance of our tool in relevant test cases.

For example, consider the indexer program. This test has 13 threads concurrently modifying a concurrent array, and locks are used to ensure the mutual exclusion of compare-and-swap (CAS) operations. This, however, results in conflicting critical sections (where the lock order does matter) within for-loops that are executed multiple times for each thread. Due to the large number of threads and the number of critical sections in each thread, large executions are created, in which the calculation of **lb** requires a lot of time.

On the other hand, there are several benchmarks where LAPOR is considerably faster than GENMC (see the `fib_bench` and `triangular` benchmarks). As shown, LAPOR is able to leverage the independence between different critical sections, resulting in far fewer executions (explorations), which in turn reduces the verification time as well.

The reason behind this is the structure of the benchmarks. In particular, both benchmarks have two concurrent writer threads manipulating shared locations within a critical section, and a third concurrent “checker” thread ensuring that each location satisfies a particular invariant (predicate) at all times. What is crucial, however, is that the two writers write to *different* locations, and the checker reads each location *individually*, rather than reading all locations in one critical section. As such, LAPOR detects that each read of the checker need not be ordered with respect to the writer that does not write to this particular location, and hence explores significantly fewer executions.

In summary, the overhead LAPOR induces is expected, and does not have a considerable impact on the performance of the tool. We note that most of the [SV-COMP 2019] benchmarks are small programs that use locks for the manipulation of shared locations, even though plain atomic accesses would suffice. Nevertheless, despite the fact that there is little to gain in such small test cases where the cost of each execution is quite low, LAPOR outperforms GENMC in a number of cases. As we show in §5.4, in programs with significantly more complex critical sections, the benefits of using an approach such as LAPOR become even more prevalent.

5.3 Synthetic Benchmarks

We continue with a number of synthetic benchmarks with their results summarized in Table 2, for both non-locking (left) and locking (right) versions of the benchmarks.

As in §5.1, the first important observation here is that LAPOR can explore exponentially fewer executions than GENMC. Indeed, inspecting the locking version of these benchmarks, LAPOR is generally significantly faster than GENMC, as it explores exponentially fewer executions. In the only case where both tools explore the same number of executions (because the ordering between the critical sections matters), namely the `rw-lock` benchmark, LAPOR maintains comparable performance to GENMC, and is only a constant factor slower than GENMC.

The second important observation here is that the introduction of lock-induced synchronization can only reduce the number of executions explored by LAPOR. As shown in the two columns of §5.1, LAPOR explores the same number or fewer executions when going from the non-locking to the locking version of a benchmark. By contrast, when going from the non-locking to the locking version, GENMC explores exponentially more executions for benchmarks such as `nreads`, `nwrites` and `nww-rr`, while exploring fewer executions only for the `rw` benchmark. This is because in the `rw` benchmark each thread unconditionally performs a number of operations, and this is why the introduction of lock-induced synchronization in the locking version reduces the number of explored executions. We note that even the mere presence of empty critical sections can lead GENMC to

Table 2. Non-locking (left) and locking (right) versions of synthetic benchmarks

	Executions		Time (s)	
	<i>GENMC</i>	<i>LAPOR</i>	<i>GENMC</i>	<i>LAPOR</i>
nreads(6)	1	1	0.04	0.04
nreads(7)	1	1	0.04	0.04
nreads(8)	1	1	0.04	0.04
nwrites(6)	1	1	0.04	0.04
nwrites(7)	1	1	0.04	0.04
nwrites(8)	1	1	0.04	0.04
nww-rr(6)	36	36	0.04	0.04
nww-rr(7)	49	49	0.04	0.04
nww-rr(8)	64	64	0.04	0.04
rw(6)	16 807	16 807	0.69	0.82
rw(7)	262 144	262 144	16.09	17.78
rw(8)	4 782 969	4 782 969	1014.49	1083.55

	Executions		Time (s)	
	<i>GENMC</i>	<i>LAPOR</i>	<i>GENMC</i>	<i>LAPOR</i>
nreads-lock(6)	720	1	0.14	0.04
nreads-lock(7)	5040	1	1.17	0.04
nreads-lock(8)	40 320	1	18.95	0.04
nwrites-lock(6)	720	1	0.14	0.04
nwrites-lock(7)	5040	1	1.12	0.04
nwrites-lock(8)	40 320	1	18.08	0.04
nww-lock-rr(6)	720	6	0.21	0.04
nww-lock-rr(7)	5040	7	1.40	0.04
nww-lock-rr(8)	40 320	8	22.32	0.04
rw-lock(6)	720	720	0.14	0.37
rw-lock(7)	5040	5040	1.27	3.10
rw-lock(8)	40 320	40 320	21.96	38.00
complex-lock(2)	42	6	0.04	0.04
complex-lock(3)	11 214	90	2.81	0.11

Table 3. Coarse-grained (left) and fine-grained (right) data structure benchmarks (seekers workload)

	Executions		Time (s)	
	<i>GENMC</i>	<i>LAPOR</i>	<i>GENMC</i>	<i>LAPOR</i>
cset(5)	120	1	0.30	0.11
cset(6)	720	1	1.57	0.15
cset(7)	5040	1	11.46	0.18
cset(8)	40 320	1	134.02	0.25
cbst(5)	120	1	0.12	0.08
cbst(6)	720	1	0.55	0.11
cbst(7)	5040	1	4.12	0.13
cbst(8)	40 320	1	56.82	0.16
chtable(5)	113 400	1	81.90	0.05
chtable(6)	–	1	–	0.07
chtable(7)	–	1	–	0.09
chtable(8)	–	1	–	0.11

	Executions		Time (s)	
	<i>GENMC</i>	<i>LAPOR</i>	<i>GENMC</i>	<i>LAPOR</i>
fset(5)	120	1	0.70	0.13
fset(6)	720	1	3.97	0.17
fset(7)	5040	1	29.48	0.22
fset(8)	40 320	1	296.35	0.29
fbst(5)	120	1	0.16	0.09
fbst(6)	720	1	0.90	0.11
fbst(7)	5040	1	6.76	0.14
fbst(8)	40 320	1	81.34	0.17
fhtable(5)	2	1	0.05	0.05
fhtable(6)	4	1	0.05	0.05
fhtable(7)	8	1	0.05	0.06
fhtable(8)	16	1	0.07	0.07

explore exponentially more executions as the number of threads increases (this is the case for the complex-locks benchmark), while such critical sections do not have any effect in LAPOR.

Finally, it is worth noting that the calculation of **lb** does not seem to impose a significant overhead on LAPOR. In the non-locking versions of these benchmarks, both tools perform similarly, even in the case of rw, where the number of explored executions is rather large.

5.4 Data Structure Benchmarks

We conclude our evaluation with data structure benchmarks. The results for a reading workload are shown in Table 3; the results for a mixed (reading and writing) workload are shown in Table 4. Entries with a dash “–” denote that the tool did not finish after six hours (the timeout limit).

First, for the reading workload (see Table 3), LAPOR explores exponentially fewer executions than GENMC for both version of these benchmarks with coarse-grained and fined-grained locking schemes. Since there are no concurrent writes on these data structures, LAPOR does not order any of the critical sections, both for the coarse-grained and the fine-grained versions. By contrast, GENMC does keep track of the lock ordering, which has a significant impact even in the fine-grained version. More specifically, even though a hand-over-hand locking scheme is used for the concurrent

Table 4. Coarse-grained (left) and fine-grained (right) data structure benchmarks (mixed workload)

	Executions		Time (s)	
	<i>GENMC</i>	<i>LAPOR</i>	<i>GENMC</i>	<i>LAPOR</i>
cset(5)	120	4	0.20	0.50
cset(6)	720	6	1.05	1.65
cset(7)	5040	6	7.83	2.52
cset(8)	40 320	18	96.20	4.95
cbst(5)	120	5	0.12	0.81
cbst(6)	720	5	0.77	2.45
cbst(7)	5040	5	5.78	3.32
cbst(8)	40 320	5	76.40	4.36
chtable(5)	113 400	2	103.59	0.11
chtable(6)	–	4	–	0.34
chtable(7)	–	8	–	0.80
chtable(8)	–	16	–	2.96

	Executions		Time (s)	
	<i>GENMC</i>	<i>LAPOR</i>	<i>GENMC</i>	<i>LAPOR</i>
fset(5)	352	4	0.87	0.51
fset(6)	4590	6	12.14	1.94
fset(7)	38 536	6	121.71	3.10
fset(8)	340 348	18	2274.48	5.99
fbst(5)	120	5	0.20	1.14
fbst(6)	720	5	1.22	3.57
fbst(7)	5040	5	9.00	4.78
fbst(8)	40 320	5	102.35	5.47
fhtable(5)	2	2	0.05	0.10
fhtable(6)	4	4	0.05	0.34
fhtable(7)	8	8	0.06	0.81
fhtable(8)	16	16	0.08	2.95

set and the concurrent BST, which would, for example, reduce lock contention if this program was running on a real CPU, *GENMC* still needs to serialize the lock acquisitions of the root node in these data structures in order to guarantee full coverage. Thus, *GENMC* explores the same number of executions for the coarse-grained and the fine-grained version of most of the test cases.

The concurrent hashtable is the only program where *GENMC* does better when moving from a coarse-grained locking scheme to a fine-grained one. In this benchmark, each thread searches for two items, and collisions occur for 5 or more threads. *GENMC* is unable to cope with the large state space of the coarse-grained version of the program, timeouts for 6 threads or more, and it is very slow even for 5 threads. In the fine-grained version of the same benchmark *GENMC* does much better. Nevertheless the number of executions it explores increases exponentially with the number of threads, while *LAPOR* explores only 1 execution for both versions.

Next, we move on to the mixed workload including both searching and inserting elements (see Table 4). Similarly for this workload, in most cases *LAPOR* explores exponentially fewer executions and is significantly faster than *GENMC*. Despite the increased complexity of calculating **1b** (due to concurrent modifications taking place), *LAPOR* maintains a very good performance and greatly outperforms *GENMC*, even though it is a bit slower compared to itself on the previous workload. This is expected because the results of the searches in the data structure now depend on the concurrent additions and hence *LAPOR* explores more than one execution per benchmark.

Finally, we conclude with two observations about two benchmarks under this workload, namely the concurrent set and the concurrent hashtable. For the concurrent set, observe that the number of executions explored increases when going from the coarse-grained to the fine-grained version. This is because the set is implemented as a sorted doubly-linked list, and when an updater reaches the end of the list attempting to insert an element, it must also acquire the lock on the list head, thus contending with both updaters and seekers attempting to traverse the list. For the concurrent hashtable, due to the collisions and concurrent modifications, tracking the lock order is deemed necessary; hence *GENMC* and *LAPOR* explore the same number of executions. As expected, *LAPOR* is slower than *GENMC* in this case, although *LAPOR* is still competitive in terms of performance.

6 RELATED WORK

Stateless Model Checking. Following the success of tools such as *VERISOFT* [Godefroid 1997] and *CHESS* [Musuvathi et al. 2008] paving the way for Stateless Model Checking, the literature

of DPOR algorithms grew rapidly [Abdulla et al. 2014; Flanagan and Godefroid 2005; Kokologiannakis et al. 2017]. Initially, DPOR algorithms partitioned the state space based on Mazurkiewicz traces [Mazurkiewicz 1987] under sequential consistency [Abdulla et al. 2014; Flanagan and Godefroid 2005], or based on Shasha-Snir traces [Shasha and Snir 1988]³ under weak memory models [Abdulla et al. 2015, 2016; Kokologiannakis et al. 2017; Norris and Demsky 2013; Zhang et al. 2015]. Subsequently, algorithms that use coarser equivalence partitioning were proposed first for SC [Aronis et al. 2018; Chalupa et al. 2017], and later for Release-Acquire [Abdulla et al. 2018]. Recently, GENMC [Kokologiannakis et al. 2019] proposed a memory model-agnostic DPOR algorithm that exploits such coarse partitioning. However, despite their effectiveness, these algorithms all utilize independence only at the instruction level, and not at the critical section level.

To our knowledge, the only DPOR algorithm that can utilize independence at the critical section level is CDPOR (Constrained Dynamic Partial Order Reduction) by Albert et al. [2018]. More specifically, CDPOR exploits *conditional independence* between critical sections: two critical sections may be considered independent, depending on the current execution state. To check for such independence, CDPOR uses checks based on pre-generated constraints, which may result in (potentially expensive) state-equivalence checks. For example, consider the following program:

$$\begin{array}{l} \text{lock}(l); \\ a_0 := x; \\ \text{unlock}(l) \end{array} \parallel \begin{array}{l} \text{lock}(l); \\ x := v_1; \\ \text{unlock}(l) \end{array} \parallel \cdots \parallel \begin{array}{l} \text{lock}(l); \\ x := v_n; \\ \text{unlock}(l) \end{array} \quad (\text{R+N-WRITES})$$

If $v_1 = v_2 = \cdots = v_n$, then CDPOR detects that reading from any of the N writes leads to an equivalent state, and hence explores 2 executions (one reading the initial value, and one reading from one of the writes). On the other hand, LAPOR explores $N + 1$ executions (one for each possible value the read can read), as it does not take conditional independence into account. However, if v_1, v_2, \dots, v_n are pairwise distinct (and non-zero), CDPOR explores $(N + 1)!$ executions (as the N writes all lead to different states), while LAPOR still explores $N + 1$ executions.

In other words, CDPOR introduces a semantic pruning in order to exploit independence at the critical section level, while LAPOR *extends* the underlying DPOR (which is a form of syntactic pruning) to the critical section level. Indeed, the two techniques are complementary and could be combined, potentially yielding even better reductions.

Unfoldings. Rodríguez et al. [2015] propose a combination of unfoldings and DPOR to reduce the number of explored executions under SC only. Although exponential reductions (compared to Mazurkiewicz traces) can be achieved in this case as well, this work is also another form of semantic pruning, and is orthogonal to the underlying DPOR framework.

Maximal Causality Reduction. Finally, Maximal Causality Reduction (MCR) [Huang 2015; Huang and Huang 2016] is another technique for stateless model checking which is different from (D)POR. MCR uses a maximal causal model to partition the state space into equivalence classes, which can lead to a coarser equivalence partitioning than the one obtained by **wb**.

Intuitively, given one execution of the program, MCR generates a set of constraints that allow a particular read in the execution to read a different value (rather than read from a different place), and the feasibility of these constraints is checked with an SMT solver. These constraints also include Φ_{lock} , which totally orders the critical sections for each lock location. It may well be possible to replace Φ_{lock} with a constraint similar to **lb** acyclicity that uses MCR's data validity constraints to deduce necessary lock orderings.

³Shasha-Snir traces are the natural extension of Mazurkiewicz traces for weak-memory consistency.

Race Detection. The idea of keeping critical sections unordered has also been applied in the context of race detection [Kini et al. 2017; Roemer et al. 2018; Smaragdakis et al. 2012]. Given a trace of an execution of a concurrent program (under SC), these techniques attempt to discern whether there exists a pair of unordered, conflicting events, which constitutes a race. However, instead of using `hb` for ordering events, critical sections are only ordered when necessary (via similar relations to `lb`), enabling the detection of more races from one given trace.

While these works are typically sound (i.e., no false positives), they are not complete, as they only guarantee that a race does not exist in the given trace, and in all traces that can be inferred from the given one (by taking all linear extensions of the partial lock ordering). By contrast, our approach is both sound and complete, although verification is in general more expensive.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we presented LAPOR, a new stateless model checking approach for handling programs with locks. We demonstrated how LAPOR reduces the number of explorations exponentially by recording the lock ordering only when necessary. LAPOR achieves this by identifying read-write conflicted critical sections, and exploiting the underlying (D)POR technique and utilizing its notion of conflicts to detect write-write-conflicted critical sections. In comparison to classical DPOR approaches, LAPOR explores far fewer executions for programs that use locks conservatively, i.e., programs that acquire locks more often than strictly necessary. This, however, means that the time spent per execution is longer, as the consistency condition checked is more complex.

More interestingly, LAPOR can be used to model check *transactional* programs under the *serializability* consistency model [Papadimitriou 1979]. As with locks, serializability requires that concurrent transactions (critical sections) appear to execute one after the other in a total sequential order. Each serializable transaction is of the form `[C]`, where `C` denotes a block of code to be executed atomically, tantamount to a critical section delimited by locks. More concretely, a transactional program P under serializability can be converted into an equivalent locking program by replacing each transaction `[C]` in P , with `lock(l); C; unlock(l)` instead, where l denotes a global lock acquired by all transactions. As such, LAPOR can model check transactional programs under serializability by treating each transaction as a critical section.

We believe that the idea of not recording the order between lock events can be generalized to other synchronization idioms, such as reader-writer locks and barriers, and can decrease the verification time for other classes of programs. More generally, the same idea could perhaps be applied to programs with reads whose returned value does not affect the thread's subsequent memory accesses. For such read events, it should be possible to avoid recording the incoming `rf`-edges, and hence drastically reduce the number of execution graphs explored. In the future, we would like to explore these ideas further.

As an additional direction of future work, we would like to investigate combining LAPOR with CDPOR (constrained dynamic partial order reduction) [Albert et al. 2018]. As discussed in §6, the semantic pruning of CDPOR is orthogonal to the syntactic pruning of LAPOR, and the combination of the two techniques may further reduce the number of explorations.

ACKNOWLEDGMENTS

We thank the OOPSLA'19 reviewers for their valuable 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

- 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>
- 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
- 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
- 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>
- Elvira Albert, Miguel Gómez-Zamalloa, Miguel Isabel, and Albert Rubio. 2018. Constrained Dynamic Partial Order Reduction. In *CAV 2018 (LNCS)*, Vol. 10982. Springer, Berlin, Heidelberg, 392–410. https://doi.org/10.1007/978-3-319-96142-2_24
- 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
- Stavros Aronis, Bengt Jonsson, Magnus Lång, and Konstantinos Sagonas. 2018. Optimal Dynamic Partial Order Reduction with Observers. In *TACAS 2018 (LNCS)*, Vol. 10806. Springer, Berlin, Heidelberg, 229–248. https://doi.org/10.1007/978-3-319-89963-3_14
- Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ Concurrency. In *POPL 2011*. ACM, New York, NY, USA, 55–66. <https://doi.org/10.1145/1926385.1926394>
- 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>
- 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>
- Patrice Godefroid. 1997. Model Checking for Programming Languages using VeriSoft. In *POPL 1997*. ACM, New York, NY, USA, 174–186. <https://doi.org/10.1145/263699.263717>
- 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>
- Shiyong 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>
- Dileep Kini, Umang Mathur, and Mahesh Viswanathan. 2017. Dynamic Race Prediction in Linear Time. In *PLDI 2017*. ACM, New York, NY, USA, 157–170. <https://doi.org/10.1145/3062341.3062374>
- 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>
- Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. 2019. Model Checking for Weakly Consistent Libraries. In *PLDI 2019*. ACM, New York, NY, USA, 15. <https://doi.org/10.1145/3314221.3314609>
- Ori Lahav and Viktor Vafeiadis. 2015. Owicki-Gries Reasoning for Weak Memory Models. In *ICALP 2015 (LNCS)*, Vol. 9135. Springer, Berlin, Heidelberg, 311–323. https://doi.org/10.1007/978-3-662-47666-6_25
- Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing Sequential Consistency in C/C++11. In *PLDI 2017*. ACM, New York, NY, USA, 618–632. <https://doi.org/10.1145/3062341.3062352>
- Leslie Lamport. 1979. How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. *IEEE Trans. Computers* 28, 9 (Sept. 1979), 690–691. <https://doi.org/10.1109/TC.1979.1675439>
- 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
- Madanlal Musuvathi, Shaz Qadeer, Thomas Ball, Gérard Basler, Piramanayagam Arumuga Nainar, and Iulian Neamtiu. 2008. Finding and Reproducing Heisenbugs in Concurrent Programs. In *OSDI 2008*. USENIX Association, Berkeley, CA, USA, 267–280.
- Brian Norris and Brian Densky. 2013. CDSChecker: Checking concurrent data structures written with C/C++ atomics. In *OOPSLA 2013*. ACM, New York, NY, USA, 131–150. <https://doi.org/10.1145/2509136.2509514>
- Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A Better x86 Memory Model: x86-TSO. In *TPHOLs 2009*. Springer, 391–407. https://doi.org/10.1007/978-3-642-03359-9_27
- Christos H. Papadimitriou. 1979. The Serializability of Concurrent Database Updates. *J. ACM* 26, 4 (Oct. 1979), 631–653. <https://doi.org/10.1145/322154.322158>

- 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>
- Jake Roemer, Kaan Genç, and Michael D. Bond. 2018. High-coverage, Unbounded Sound Predictive Race Detection. In *PLDI 2018*. ACM, New York, NY, USA, 374–389. <https://doi.org/10.1145/3192366.3192385>
- Dennis Shasha and Marc Snir. 1988. Efficient and Correct Execution of Parallel Programs That Share Memory. *ACM Trans. Program. Lang. Syst.* 10, 2 (April 1988), 282–312. <https://doi.org/10.1145/42190.42277>
- Yannis Smaragdakis, Jacob Evans, Caitlin Sadowski, Jaeheon Yi, and Cormac Flanagan. 2012. Sound Predictive Race Detection in Polynomial Time. In *POPL 2012*. ACM, New York, NY, USA, 387–400. <https://doi.org/10.1145/2103656.2103702>
- SPARC International Inc. 1994. *The SPARC architecture manual (version 9)*. Prentice-Hall.
- SV-COMP. 2019. Competition on Software Verification (SV-COMP). <https://sv-comp.sosy-lab.org/2019/> [Online; accessed 27-March-2019].
- 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>