Optimal Prediction of Synchronization-Preserving Races

UMANG MATHUR, University of Illinois at Urbana-Champaign, USA

ANDREAS PAVLOGIANNIS, Aarhus University, Denmark

MAHESH VISWANATHAN, University of Illinois at Urbana-Champaign, USA

Concurrent programs are notoriously hard to write correctly, as scheduling nondeterminism introduces subtle errors that are both hard to detect and to reproduce. The most common concurrency errors are (*data*) *races*, which occur when memory-conflicting actions are executed concurrently. Consequently, considerable effort has been made towards developing efficient techniques for race detection. The most common approach is *dynamic race prediction*: given an observed, race-free trace σ of a concurrent program, the task is to decide whether events of σ can be correctly reordered to a trace σ^* that witnesses a race hidden in σ .

In this work we introduce the notion of *sync(hronization)-preserving races*. A sync-preserving race occurs in σ when there is a witness σ^* in which synchronization operations (e.g., acquisition and release of locks) appear in the same order as in σ . This is a broad definition that *strictly subsumes* the famous notion of happens-before races. Our main results are as follows. First, we develop a sound and complete algorithm for predicting sync-preserving races. For moderate values of parameters like the number of threads, the algorithm runs in $\tilde{O}(N)$ time and space, where N is the length of the trace σ . Second, we show that the problem has a $\Omega(N/\log^2 N)$ space lower bound, and thus our algorithm is essentially *time and space optimal*. Third, we show that predicting races with *even just a single* reversal of two sync operations is NP-complete and even $W[1]$ -hard when parameterized by the number of threads. Thus, sync-preservation characterizes *exactly* the tractability boundary of race prediction, and our algorithm is nearly *optimal* for the tractable side. Our experiments show that our algorithm is fast in practice, while sync-preservation characterizes races often missed by state-of-the-art methods.

CCS Concepts: • **Software and its engineering** → **Software verification and validation**; • **Theory of computation** → *Theory and algorithms for application domains*; *Program analysis*.

Additional Key Words and Phrases: concurrency, dynamic analysis, race detection, complexity

ACM Reference Format:

Umang Mathur, Andreas Pavlogiannis, and Mahesh Viswanathan. 2021. Optimal Prediction of Synchronization-Preserving Races. *Proc. ACM Program. Lang.* 5, POPL, Article 36 (January 2021), 29 pages. <https://doi.org/10.1145/3434317>

1 INTRODUCTION

The verification of concurrent programs is one of the main challenges in formal methods. Concurrency adds a dimension of non-determinism to program behavior which stems from inter-process communication. Accounting for such non-determinism during program development is a challenging mental task, making concurrent programming significantly error-prone. At the same time, bugs

Authors' addresses: Umang Mathur, Department of Computer Science, University of Illinois at Urbana-Champaign, USA, umathur3@illinois.edu; Andreas Pavlogiannis, Department of Computer Science, Aarhus University, Denmark, pavlogiannis@cs.au.dk; Mahesh Viswanathan, Department of Computer Science, University of Illinois at Urbana-Champaign, USA, vmahesh@illinois.edu.

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

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/1-ART36

<https://doi.org/10.1145/3434317>

due to concurrency are very hard to reproduce manually, and automated techniques for doing so are crucial in enhancing the productivity of software developers.

Data races are the most common form of concurrency errors. A data race (sometimes just called a race) occurs when a thread of a multi-threaded program accesses a shared memory location while another thread is modifying it without proper synchronization. The presence of a data race is often symptomatic of a serious bug in the program [Lu et al. 2008]; races have caused data corruption and compilation errors [Boehm 2011; Kasikci et al. 2013; Narayanasamy et al. 2007], and significant system errors [Boehm 2012; Zhivich and Cunningham 2009] in the past. Therefore, considerable research has focused on detecting and preventing races in multi-threaded programs.

One of the most popular approaches to race prediction is via dynamic analysis [Bond et al. 2010; Flanagan and Freund 2009; Pozniansky and Schuster 2003]. Unlike static analysis, dynamic race prediction is performed at runtime. Such techniques determine if an observed execution provides evidence for the existence of a *possibly alternate* program execution that can concurrently perform conflicting data accesses¹. The underlying principle is that a race is present but “hidden” in a large number of different program executions; hence techniques that uncover such hidden races can accelerate the process of debugging concurrent programs significantly. The popularity of dynamic race prediction techniques further stems (i) from their scalability to large production software, and (ii) from their ability to produce only sound error reports.

The most popular dynamic race prediction techniques are based on Lamport’s happens-before partial order [Lamport 1978]. These techniques scan the input trace, determine happens-before orderings on-the-fly, and report a race on a pair of conflicting data accesses if they are unordered by happens-before. This approach is sound, in that the presence of unordered conflicting data accesses ensures the existence of an execution with a race. While happens-before based analysis fails to predict races in various cases [Smaragdakis et al. 2012], its wide deployment is based on the fact that the algorithm is fast, single pass, and runs in linear time. The principle that forms the basis of its efficiency is the following. When reasoning about alternate executions, happens-before analysis does not consider any execution in which the order of synchronization primitives is reversed from that in the observed execution. We call such alternate executions *sync(hronization)-preserving executions*. Other, more powerful race prediction techniques [Genç et al. 2019; Huang et al. 2014; Huang and Rajagopalan 2016; Pavlogiannis 2019; Roemer et al. 2018; Smaragdakis et al. 2012] sacrifice this principle and consider alternate executions that are not sync-preserving. Naturally, this typically results in performance degradation, as the problem is in general NP-hard [Mathur et al. 2020a], and considerable efforts are made towards improving the scalability of such techniques [Roemer and Bond 2019; Roemer et al. 2020].

Although happens-before only detects races whose exposure preserves the ordering of synchronization primitives, it can still miss simple races that adhere to this pattern. For example, consider the trace σ_1 shown in Figure 1a. Let us name the events of this trace based on the order in which they appear in the trace; thus, e_i denotes the i^{th} event of the trace. Here, the partial order happens-before orders the first $w(x)$ (event e_1) and the last $w(x)$ (event e_6), and therefore, does not detect any race in this execution. However events e_1 and e_6 are in race. This can be exposed by the alternate execution shown in Figure 1b, which is obtained by dropping the critical section of lock ℓ performed by thread t_1 . Notice that the order of synchronization events (namely, $\text{acq}(\ell)$ and $\text{rel}(\ell)$ events) *that appear in the trace* of Figure 1b, are in the same order as in the trace of Figure 1a, and hence this is a sync-preserving execution. Thus, the notion of sync-preservation captures races beyond standard happens-before races.

¹Conflicting data accesses come from different threads, access a common memory location, and at least one is a write.

Fig. 1. (1a) shows a trace σ_1 with a sync-preserving race (e_1, e_6) missed by $\leq_{HB}$. (1b) shows a witness that exposes the race. (1c) shows a sync-preserving race (e_1, e_6) that is non-consecutive, due to the intermediate event e_3 .

Another important limitation of happens-before and virtually all partial-order methods [Kini et al. 2017; Mathur et al. 2018; Roemer et al. 2018, 2020; Smaragdakis et al. 2012] is highlighted in Figure 1c. The trace σ_2 has a race between e_1 and e_6 , both conflicting on variable x . Notice, however, that the intermediate event e_3 also accesses x , but is not in race with either e_1 or e_6 . Partial-order methods for race prediction are limited to capturing races *only between successive* conflicting accesses². Hence, distant races that are interjected with intermediate conflicting but non-racy events, are missed by such methods. On the other hand, sync-preservation is not bound to such limitations: (e_1, e_6) is characterized as a race under this criterion, regardless of the intermediate, non-racy e_3 , and is exposed by a witness that omits the critical section on lock ℓ in the thread t_1 .

Our Contributions. Motivated by the above observations, we make the following contributions.

- (1) We introduce the novel notion of *sync(hronization)-preserving data races*. This is a *sound* notion of predictable races, and it *strictly subsumes* the standard notion of happens-before races. Moreover, it characterizes races between events that can be arbitrarily far apart in the input trace, as opposed to happens-before and other partial-order methods that only characterize races between *successive* conflicting accesses. Our notion is applicable to all concurrency settings, and interestingly, it is also *complete* for systems with synchronization-deterministic concurrency [Aguado et al. 2018; Bocchino et al. 2009; Cui et al. 2015; Zhao et al. 2019].
- (2) We develop an efficient, single-pass, nearly linear time algorithm SyncP that, given a trace σ , detects whether σ contains a sync-preserving race. In fact, our algorithm soundly reports *all* events e_2 which are in a sync-preserving race with an event e_1 that appears earlier in σ . Given N events in σ , our algorithm spends $\tilde{O}(N)$ time, where $\tilde{O}$ hides factors poly-logarithmic in N , when other parameters of the input (e.g., number of threads) are $\tilde{O}(1)$.
- (3) Although our algorithm performs a single pass of the trace, in the worst case, it might use space that is nearly linear in the length of the trace, i.e., $\tilde{O}(N)$ space. Hence follows a natural question: is there an efficient algorithm for sync-preserving race prediction that uses considerably less space? We answer this question in negative, by showing that *any* single-pass algorithm for detecting even a single sync-preserving race must use nearly linear space. Hence, our algorithm SyncP has nearly *optimal* performance in both time and space.
- (4) We next study the complexity of race prediction with respect to the number of synchronization reversals that might occur when constructing a witness that exposes the race. In the case

²When the earlier access is a read instead of a write, this statement is true *per thread*.

of synchronization via locks, this number corresponds to the number of critical sections whose order is reversed in the witness trace. We prove that the problem of predicting races which can be witnessed by *a single* reversal (of two critical sections) is NP-complete and even $W[1]$ -hard when parameterized by the number of threads. Thus, sync-preservation characterizes *exactly* the tractability boundary of race prediction, and our algorithm is nearly *optimal* for the tractable side. Moreover, our result shows that *any level* of synchronization suffices to make the problem of race prediction as hard as in the general case.

- (5) Finally, we have implemented our race prediction algorithm SyncP and evaluated its performance on standard benchmarks. Our results show that sync-preservation characterizes many races that are missed by state-of-the-art methods, and SyncP detects them efficiently.

2 PRELIMINARIES

In this section we establish notation useful throughout of the paper. The exposition follows other related works in the literature.

2.1 Background on Dynamic Data Race Prediction

Traces and Events. Our objective is to develop a dynamic analysis technique which works over execution traces, or simply *traces* of concurrent programs. We work with the sequential consistency memory model. In this setting, traces are sequences of events. We will use $\sigma, \sigma', \dots, \sigma_1, \sigma_2, \dots$ to denote traces. Every event of σ can be represented as a tuple $e = \langle i, t, \text{op} \rangle$, where i is a unique identifier of e in σ , t is the thread that performs e and op is the operation performed in the event e . We often omit the unique identifier of such a tuple and simply write $e = \langle t, \text{op} \rangle$. We use $\text{thr}(e)$ and $\text{op}(e)$ to denote the thread performing e and the operation performed by e . An operation can be one of read from or write to a shared memory location or *variable* x , denoted $\text{r}(x)$ and $\text{w}(x)$, and acquisition or release of a lock ℓ , denoted $\text{acq}(\ell)$ or $\text{rel}(\ell)$. Forks and joins can be naturally handled, but we avoid introducing them here for notational convenience. We denote by Events_σ the set of events in a trace σ . We use Thr_σ , Vars_σ and Locks_σ to denote respectively the threads, variables and locks that appear in σ . Likewise, we use $\text{Acquires}_\sigma(\ell)$ and $\text{Releases}_\sigma(\ell)$ to denote the set of acquire and release events of σ on lock $\ell \in \text{Locks}_\sigma$.

We require that traces obey lock semantics. In particular, every lock ℓ is released by a thread t only if there is an earlier matching acquire event by the same thread t , and that each such lock is held by at most one thread at a time. Formally, let $\sigma|_\ell$ denote the projection of σ to the set of events $\text{Acquires}_\sigma(\ell) \cup \text{Releases}_\sigma(\ell)$. We require that for every lock ℓ , the sequence $\sigma|_\ell$ is a prefix of some sequence that belongs to the language of the regular expression $(\sum_{t \in \text{Thr}_\sigma} \langle t, \text{acq}(\ell) \rangle \cdot \langle t, \text{rel}(\ell) \rangle)^*$.

For an acquire event e , we use $\text{match}_\sigma(e)$ to denote the matching release event of e if one exists (and $\perp$ otherwise). Similarly, for a release event e , $\text{match}_\sigma(e)$ is the matching acquire of e on the same lock. For an acquire event e , the critical section protected by e , denoted $\text{CS}_\sigma(e)$, is the set of events e' such that $\text{thr}(e') = \text{thr}(e)$ and e' occurs after e and before the matching release $\text{match}_\sigma(e)$ (if it exists) in σ . For a release event e , we have $\text{CS}_\sigma(e) = \text{CS}_\sigma(\text{match}_\sigma(e))$.

Orders on Traces. A partial order $\leq_P^\sigma$ defined over a trace σ is a reflexive, anti-symmetric and transitive binary relation on Events_σ ; the symbol P is an optional identifier for the partial order. We write $e_1 \leq_P^\sigma e_2$ to denote $(e_1, e_2) \in \leq_P^\sigma$, where $e_1, e_2 \in \text{Events}_\sigma$. For a partial order $\leq_P^\sigma$, we use $<_P^\sigma$ to denote the strict order $\leq_P^\sigma \setminus \{(e, e) \mid e \in \text{Events}_\sigma\}$. We write $e_1 \not\leq_P^\sigma e_2$ to denote that $(e_1, e_2) \notin \leq_P^\sigma$. Events $e_1, e_2 \in \text{Events}_\sigma$ are said to be *unordered* by $\leq_P^\sigma$, denoted $e_1 \parallel_P^\sigma e_2$ if $e_1 \not\leq_P^\sigma e_2$ and $e_2 \not\leq_P^\sigma e_1$; otherwise, we write $e_1 \Vdash_P^\sigma e_2$, denoting that e_1 and e_2 are ordered by $\leq_P^\sigma$ in one or the other way. When σ is clear from context, we will use $\leq_P$, $<_P$, $\not\leq_P$, $\parallel_P$ and $\Vdash_P$ instead of respectively $\leq_P^\sigma$, $<_P^\sigma$,

Fig. 2. Traces, data races and predictable data races

$\not\leq_p^\sigma$, $\parallel_p^\sigma$ and $\parallel_p^\sigma$. For a partial order $\leq_p^\sigma$, a set $S \subseteq \text{Events}_\sigma$ is said to be *downward-closed with respect to* $\leq_p^\sigma$ if for every $e, e' \in \text{Events}_\sigma$, if $e \leq_p^\sigma e'$ and $e' \in S$, then $e \in S$.

The *trace-order* $\leq_{tr}^\sigma$ defined by σ is the total order on Events_σ imposed by the sequence σ , i.e., $e_1 \leq_{tr}^\sigma e_2$ iff the event e_1 occurs before e_2 in σ . The *thread-order* (or *program-order*) $\leq_{TO}^\sigma$ of σ is the partial order on Events_σ that orders events in the same thread: for two events $e_1, e_2 \in \text{Events}_\sigma$, $e_1 \leq_{TO}^\sigma e_2$ iff $e_1 \leq_{tr}^\sigma e_2$ and $\text{thr}(e_1) = \text{thr}(e_2)$.

Conflicting Events and Data Races. Let σ be a trace. Two events $e_1, e_2 \in \text{Events}_\sigma$ are said to be *conflicting*, denoted $e_1 \asymp e_2$, if $\text{thr}(e_1) \neq \text{thr}(e_2)$, and there is a common variable $x \in \text{Vars}_\sigma$ such that $\text{op}(e_1), \text{op}(e_2) \in \{r(x), w(x)\}$ and at least one of $\text{op}(e_1)$ and $\text{op}(e_2)$ is $w(x)$. Let ρ be a trace with $\text{Events}_\rho \subseteq \text{Events}_\sigma$. An event $e \in \text{Events}_\sigma$ is said to be σ -*enabled* in ρ if $e \notin \text{Events}_\rho$ and for all events $e' \in \text{Events}_\sigma$ such that $e' <_{TO}^\sigma e$, we have $e' \in \text{Events}_\rho$. A pair of conflicting events (e_1, e_2) in σ is said to be a *data race* of σ if σ has a prefix σ' such that both e_1 and e_2 are σ -enabled in σ' . The trace σ is said to have a data race if there is a pair of conflicting events (e_1, e_2) in σ that constitutes a data race of σ .

Example 1. Consider the trace σ_3 in Figure 2a. The set of events of σ_3 is $\text{Events}_{\sigma_3} = \{e_1, e_2, \dots, e_6\}$, $\text{Thr}_\sigma = \{t_1, t_2\}$, $\text{Vars}_\sigma = \{x\}$ and $\text{Locks}_\sigma = \{\ell\}$. For the event $e_1 = \langle t_1, \text{acq}(\ell) \rangle$, we have $\text{thr}(e_1) = t_1$ and $\text{op}(e_1) = \text{acq}(\ell)$. The trace order of this trace is $\leq_{tr}^{\sigma_3} = \{(e_i, e_j) \mid i \leq j\}$ and the thread-order is $\leq_{TO}^{\sigma_3} = \{(e_1, e_2), (e_1, e_5), (e_2, e_5), (e_3, e_4), (e_3, e_6), (e_4, e_6)\}$. Events e_5 and e_6 conflict because they access the same variable x and are performed by different threads. For the prefix trace $\sigma'_3 = e_1 \cdot e_2 \cdot e_3 \cdot e_4$, both e_5 and e_6 are σ'_3 -enabled in σ'_3 . Thus, (e_5, e_6) constitutes a data race of σ_3 .

Correct Reorderings. Execution traces of concurrent programs are sensitive to thread scheduling, and looking for a trace with a specific pattern is like searching for a needle in a haystack. In terms of data race detection, this means that a dynamic analysis that looks for executions with enabled conflicting events (data races) is likely to miss many data races that might have otherwise been captured in alternate executions of the same program that arise due to slightly different thread scheduling. The notion of data race *prediction* attempts to alleviate this problem by capturing a more robust notion of data races. The idea here is to infer data races that might occur in alternate reorderings of an observed trace, thereby detecting data races beyond those in just the execution that was observed. The set of allowable reorderings of an observed trace σ is defined in a manner that ensures that data races can be detected agnostic of the program that generated σ in the first place. Such a notion is captured by a *correct reordering* which we define next.

For a trace σ and a read event e , we use $\text{lw}_\sigma(e)$ to denote the write event observed by e . That is, $e' = \text{lw}_\sigma(e)$ is the last (according to the trace order $\leq_{tr}^\sigma$) write event e' of σ such that e and e' access the same variable and $e' \leq_{tr}^\sigma e$; if no such e' exists, then we write $\text{lw}_\sigma(e) = \perp$.

Given the above notation, a trace ρ is said to be a correct reordering of trace σ if

- (a) $\text{Events}_\rho \subseteq \text{Events}_\sigma$
- (b) Events_ρ is downward closed with respect to $\leq_{\text{TO}}^\sigma$, and further $\leq_{\text{TO}}^\rho \subseteq \leq_{\text{TO}}^\sigma$,
- (c) for every read event $e \in \text{Events}_\rho$, $\text{lw}_\rho(e) = \text{lw}_\sigma(e)$.

The above definition ensures that if ρ is a correct reordering of σ , then every program that generates the execution trace σ also generates ρ . This is because ρ preserves both intra-thread ordering, as well as the values read by every read occurring in ρ , thereby preserving any control flow that might have been taken by σ . This style of formalizing alternative executions based on semantics of concurrent objects was popularized by [Herlihy and Wing 1990] and by prior race detection works [Said et al. 2011; Şerbănuță et al. 2012]. Our definition of correct reordering has been derived from [Smaragdakis et al. 2012], which has subsequently also been used in the literature [Genç et al. 2019; Kini et al. 2017; Mathur et al. 2018, 2020a; Pavlogiannis 2019; Roemer et al. 2018].

Data Race Prediction. Armed with the notion of correct reorderings, we can now define a more robust notion of data races. A pair of conflicting events (e_1, e_2) in σ is said to be a *predictable* data race of σ if there is a correct reordering ρ of σ such that e_1, e_2 are σ -enabled in ρ . We remark that a pair of conflicting events (e_1, e_2) in trace σ may not be a data race of σ , but nevertheless may still be a *predictable* data race of σ .

Example 2. Consider the trace σ_4 in Figure 2b. Observe that there is no prefix of σ_4 in which both e_1 and e_6 are enabled. However, (e_1, e_6) is a predictable race of σ_4 that is witnessed by the singleton correct reordering $\sigma_4^{\text{CR}} = e_5$ in which both e_1 and e_6 are enabled; σ_4^{CR} is both downward closed with respect to, and respects $\leq_{\text{TO}}^{\sigma_4}$. Further, it has no read events and thus vacuously every read observes the same last write as in σ_4 . The other pair of conflicting events in σ_4 , namely (e_3, e_6) , however, is not a predictable race. These events are protected by a common lock, and there is no correct reordering in which e_3 and e_6 are simultaneously enabled — any attempt at doing so will lead to overlapping critical sections on ℓ , thereby violating lock semantics.

Example 3. Now, consider σ_5 in Figure 2c. Here, the conflicting pair (e_3, e_6) cannot be a predictable race as in the case of σ_4 — the lock ℓ protects both e_3 and e_6 . Now consider the other conflicting pair (e_1, e_8) . Let ρ be a correct reordering of σ_5 in which e_8 is enabled. We must have $e_6 \in \text{Events}_\rho$ (ρ must be $\leq_{\text{TO}}^{\sigma_5}$ -downward closed) and further $e_3 \in \text{Events}_\rho$ (as $e_3 = \text{lw}_{\sigma_5}(e_6) = \text{lw}_\rho(e_6)$). Clearly, e_1 cannot be enabled in any such trace ρ , and thus, the trace σ_5 has no predictable data race.

The central theme of race prediction is to solve the problem below.

Problem 1 (Data Race Prediction). Given a trace σ , determine if σ has a predictable data race.

A Note on Soundness. We say that an algorithm for data race prediction is *sound* if whenever the algorithm reports a YES answer, then the given trace has a predictable data race. Likewise, an algorithm is *complete* if the algorithm reports YES whenever the input trace has a data race. Our convention for this nomenclature ensures that no false positives are reported by a *sound* algorithm [Sergey 2019] and is consistent with prior work on data race prediction [Genç et al. 2019; Kini et al. 2017; Pavlogiannis 2019; Roemer et al. 2018; Smaragdakis et al. 2012]. Soundness is often a desirable property for dynamic race predictors for widespread adoption [Gorogiannis et al. 2019].

2.2 Synchronization-Preserving Data Races

In general, the problem of data race prediction is intractable [Mathur et al. 2020a], and a sound and complete algorithm for data race prediction is unlikely to scale beyond programs of even moderate size. A recent trend in predictive analysis for race detection instead, aims to develop

Fig. 3. Sync-preserving correct reordering and sync-preserving races

techniques that are sound but incomplete, with successively better prediction power (ability to report more data races) than previous techniques [Genç et al. 2019; Kini et al. 2017; Pavlogiannis 2019; Roemer et al. 2018; Smaragdakis et al. 2012]. Most of these techniques are either based on partial orders [Kini et al. 2017; Pozniansky and Schuster 2003; Smaragdakis et al. 2012] or use graph-based algorithms [Pavlogiannis 2019; Roemer et al. 2018]. In this paper, we characterize a class of predictable data races, called *sync(hronization)-preserving* races, which we define shortly. We will later (Section 4) present an algorithm that reports a race iff the input trace has a sync-preserving race. Since sync-preserving races are predictable races, our algorithm will be sound for race prediction.

Sync-Preserving Correct Reordering. A correct reordering of a trace is called *sync(hronization)-preserving* if it does not reverse the order of synchronization constructs; in our formalism, traces use locks as synchronization primitives to enforce mutual exclusion. Formally, a correct reordering ρ of a given trace σ is *sync-preserving* with respect to σ if for every lock ℓ and for any two acquire events $e_1, e_2 \in \text{Acquires}_\rho(\ell)$, we have $e_1 \leq_{\text{tr}}^\rho e_2$ iff $e_1 \leq_{\text{tr}}^\sigma e_2$. In other words, the order of two critical sections on the same lock is the same in σ and ρ . Let us illustrate this notion on an example.

Example 4. Consider trace σ_6 in Figure 3a. This trace has 3 critical sections on lock ℓ . Now consider the correct reordering σ_6^{CR} (Figure 3b) of σ_6 . Here, the critical section in thread t_1 is not present. But, nevertheless, the order amongst the remaining critical sections on ℓ (in threads t_2 and t_3) is the same as in σ_6 , making σ_6^{CR} a sync-preserving correct reordering of σ_6 . This example also demonstrates that the order of read and write events may be different in a trace and its sync-preserving correct reordering (as in Figure 3).

A pair of conflicting events (e_1, e_2) of a trace σ is said to be a *sync(hronization)-preserving race* of σ if there is a sync-preserving correct reordering ρ of σ in which e_1 and e_2 are σ -enabled.

Example 5. Let us again consider traces from Figure 3. Events e_1 and e_8 in σ_6 (Figure 3a) correspond respectively to events e_6 and e_7 in σ_6^{CR} (Figure 3b). These two events are σ_6 -enabled in the prefix $\rho = e_1 \cdot e_2 \cdot e_3 \cdot e_4 \cdot e_5$ of σ_6^{CR} . As a result, (e_1, e_8) is a sync-preserving race of σ_6 . Likewise, (e_1, e_4) is also a sync-preserving race of σ_6 witnessed by the singleton sync-preserving correct reordering $\rho' = \langle t_2, \text{acq}(\ell) \rangle$, in which both e_1 and e_4 are enabled.

In this paper we present a linear time algorithm for the following decision problem, giving a *sound* algorithm for Problem 1.

Problem 2 (Sync-Preserving Race Prediction). Given trace σ , determine if there is a pair of conflicting events (e_1, e_2) in σ such that (e_1, e_2) is a sync-preserving data race of σ .

Comparison with Other Approaches. Here we briefly compare sync-preserving races with other approaches in the literature for sound dynamic race prediction. Races reported using the famous *happens-before* (HB) partial order [Pozniansky and Schuster 2003], and its extension to *schedulable-happens-before* (SHB) [Mathur et al. 2018] are strictly subsumed by this notion. That is, these techniques only compute sync-preserving races, but can also miss simple cases of sync-preservation, as already illustrated in the examples of Figure 1. The *causally precedes* (CP) partial order [Smaragdakis et al. 2012], and its extension to the *weak causally precedes* (WCP) partial order [Kini et al. 2017] are capable of predicting races that reverse critical sections. However, they are closed under composition with HB, and as such can miss even simple sync-preserving races, even on two-threaded traces. The *doesn't commute* (DC) partial order [Roemer et al. 2018] is an unsound weakening to WCP, that further undergoes a vindication phase to filter out unsound reports. Nevertheless, DC is somewhat similar to WCP and also misses sync-preserving races. The recently introduced partial order *strong-dependently-precedes* (SDP) [Genç et al. 2019], while claimed to be sound in that paper, is, in fact, unsound. In our technical report [Mathur et al. 2020b], we show a counter-example to the soundness theorem of SDP, which we confirmed with the authors [Genç et al. 2020]. The partial order WDP [Genç et al. 2019] is a (unsound) weakening of DC, and can nevertheless miss sync-preserving races in the *vindication* phase employed for ruling out false positives. We further refer to [Mathur et al. 2020b] for a few examples that illustrate the above comparison.

3 SUMMARY OF MAIN RESULTS

Here we give an outline of the main results of this paper. In later sections we present the details, i.e., algorithms, proofs and examples. Due to limited space, some technical proofs are relegated to our companion technical report [Mathur et al. 2020b]. Our first result is an algorithm for dynamic prediction of sync-preserving races. We show the following theorem.

Theorem 3.1. *Sync-preserving race prediction is solvable in $O(N \cdot \mathcal{T}^2 + \mathcal{A} \cdot \mathcal{V} \cdot \mathcal{T}^3)$ time and $O(N + \mathcal{T}^3 \cdot \mathcal{V} \cdot \mathcal{L})$ space, for a trace σ with length N , $\mathcal{T}$ threads, $\mathcal{A}$ acquires, and $\mathcal{V}$ variables.*

In many settings the number of events N and number of acquires $\mathcal{A}$ are the dominating parameters, whereas the other parameters are much smaller, i.e., $\mathcal{T}, \mathcal{V} = \tilde{O}(1)$, where $\tilde{O}$ hides poly-logarithmic factors. Hence, the complexity of our algorithm becomes $\tilde{O}(N)$ for both time and space. Our next result shows that a linear space complexity is essentially unavoidable when predicting sync-preserving races with one-pass streaming algorithms.

Theorem 3.2. *Any one-pass algorithm for sync-preserving race prediction on traces with ≥ 2 threads, N events and $\Omega(\log N)$ locks uses $\Omega(N/\log^2 N)$ space.*

Clearly, any algorithm must spend linear time, while Theorem 3.2 shows that the algorithm must also use (nearly) linear space. As our algorithm uses $\tilde{O}(N)$ time and space, it is optimal for both resources, modulo poly-logarithmic improvements. Our next theorem shows a combined time-space lower bound for the problem, which highlights that reducing the space usage must lead to an increased running time, given that the algorithm is executed on the Turing Machine model.

Theorem 3.3. *Consider the problem of sync-preserving race prediction on traces with ≥ 2 threads, N events and $\Omega(\log N)$ locks. Consider any Turing Machine algorithm for the problem with time and space complexity $T(N)$ and $S(N)$, respectively. Then we have $T(N) \cdot S(N) = \Omega(N^2/\log^2 N)$.*

Finally, we study the complexity of general race prediction as a function of the number of reversals of synchronization operations. Given our positive result in Theorem 3.1, can we relax

our restriction of sync-preservation while retaining a tractable definition of predictable races? Our next theorem answers this question in negative.

Theorem 3.4. *Dynamic race prediction on traces with a single lock and two critical sections is $W[1]$ -hard parameterized by the number of threads.*

Note that $W[1]$ -hardness implies NP-hardness. The theorem has two important implications.

- (1) Any witness of predictable races in the setting of Theorem 3.4 is either a sync-preserving reordering, or reverses the order of a single pair of acquire events. Thus, Theorem 3.1 and Theorem 3.4 establish a *tight dichotomy* on the tractability of the problem, based on the number of synchronization reversals: the problem is as hard as in the general case for just 1 reversal, while it is efficiently solvable for no reversals.
- (2) The general problem of dynamic race prediction was shown to be $W[1]$ -hard in [Mathur et al. 2020a]. However, that proof requires traces with $\Omega(N)$ critical sections, and hence it applies to traces that essentially comprise of synchronization events entirely. In contrast, the class of traces in Theorem 3.4 has the smallest level of synchronization possible, i.e., just a single lock and two critical sections on that lock. Hence, Theorem 3.4 shows that *any* amount of lock-based synchronization suffices to make the problem as hard as in the general case.

Together, Theorem 3.1, Theorem 3.2 and Theorem 3.4 characterize *exactly* the tractability horizon of race prediction, and show that our algorithm is *time and space optimal* for the tractable side.

4 DETECTING SYNCHRONIZATION-PRESERVING RACES

In this section, we discuss our algorithm SyncP for detecting sync-preserving data races. The complete algorithm is presented in Section 4.4. The algorithm may appear complex at first glance and to make the exposition simple, we first present a high-level overview of the algorithm in Section 4.1. In the overview, we highlight important observations and algorithmic insights for solving smaller subproblems of the main problem of sync-preserving race prediction. Section 4.2 and Section 4.3 present details of the algorithms for the smaller subproblems, and pave the way for the final algorithm in Section 4.4. In Section 4.5, we present a matching space lowerbound result for detecting sync-preserving races, thereby showing the optimality of our algorithm.

4.1 Insights and Overview of the Algorithm

Our algorithm, SyncP, relies on several important observations that are crucial for detecting sync-preserving races in linear time. In order to present these observations, it is helpful to define intermediate subproblems.

Problem 3 (Sync-Preserving Race Prediction Given Pair). Given a trace σ and a pair of conflicting events (e_1, e_2) of σ , determine if (e_1, e_2) is a sync-preserving data race of σ .

Problem 4 (Sync-Preserving Race Prediction Given Event and Thread). Given a trace σ , an event e in σ and a thread $t \neq \text{thr}(e)$, check if there is an event $e' \leq_{\text{tr}}^{\sigma} e$ such that $\text{thr}(e') = t$ and (e', e) constitutes a sync-preserving race of σ .

Observe that a trace with N events can have $O(N^2)$ conflicting pairs of events. Thus, an algorithm for Problem 3 that runs in time $O(T(N))$ can be used to obtain an algorithm for Problem 4 (resp. Problem 2) that runs in time $O(N \cdot T(N))$ (resp. $O(N^2 \cdot T(N))$) by checking if every other event of the given thread t that conflicts with e is also in race with e (resp. every conflicting pair of events is a race). We will, however, present algorithms for all three problems that run in $O(N)$ time.

Efficiently Solving Problem 3. Our first important observation towards a full fledged solution to Problem 3 is that when checking for the existence of a sync-preserving reordering (of a trace σ)

that witnesses a race on a given pair (e_1, e_2) , it, suffices to only search for those reorderings ρ of σ which impose the same order as in σ , on *all of its events*, and not just on the critical sections. We formalize this in Lemma 4.1.

Lemma 4.1. *If (e_1, e_2) is a sync-preserving race of σ , then there is a correct reordering ρ of σ such that both e_1, e_2 are σ -enabled in ρ and $\leq_{tr}^\rho \subseteq \leq_{tr}^\sigma$.*

The implication of Lemma 4.1 is the following. When we are searching for a correct reordering ρ of σ that witnesses a race (e_1, e_2) , and if we already have access to the set of events $S \subseteq \text{Events}_\sigma$ of a candidate reordering ρ , a simple check suffices — does S form a correct reordering of σ when linearized according to $\leq_{tr}^\sigma$? In other words, we do not need to enumerate and check all the (exponentially many) permutations of events in S . Thus, Problem 3 — ‘search for a sync-preserving correct reordering ρ ’ — reduces to a simpler problem — ‘search for an appropriate set of events’. Of course, not all sets $S \subseteq \text{Events}_\sigma$ of events can be linearized (according to $\leq_{tr}^\sigma$) to obtain a correct reordering of σ . At the very least, S should satisfy some sanity conditions which we outline next.

Definition 1 (Thread-Order and Last-Write Closure). Let σ be a trace. A set $S \subseteq \text{Events}_\sigma$ is said to be $(\leq_{TO}^\sigma, \text{lw}_\sigma)$ -closed if (a) S is downward-closed with respect to $\leq_{TO}^\sigma$, and (b) for any read event $r \in \text{Events}_\sigma$, if $r \in S$ and if $\text{lw}_\sigma(r)$ exists, then $\text{lw}_\sigma(r) \in S$. The $(\leq_{TO}^\sigma, \text{lw}_\sigma)$ -closure of a set $S \subseteq \text{Events}_\sigma$, denoted $\text{TLClosure}_\sigma(S)$ is the smallest set $S' \subseteq \text{Events}_\sigma$ such that $S \subseteq S'$ and S' is $(\leq_{TO}^\sigma, \text{lw}_\sigma)$ -closed.

We remark that any correct reordering ρ of σ that contains events in the set S must also contain all the events in $\text{TLClosure}_\sigma(S)$.

Definition 2 (Sync-Preserving Closure). Let σ be a trace. A set $S \subseteq \text{Events}_\sigma$ is said to be sync-preserving closed if

- (a) S is $(\leq_{TO}^\sigma, \text{lw}_\sigma)$ -closed, and
- (b) for any two acquire events $a_1, a_2 \in \text{Acquires}_\sigma(\ell)$ with $a_1 \leq_{tr}^\sigma a_2$, if both $a_1, a_2 \in S$, then $\text{match}_\sigma(a_1) \in S$

The sync-preserving closure of a set $S \subseteq \text{Events}_\sigma$, denoted $\text{SPClosure}_\sigma(S)$ is the smallest set $S' \subseteq \text{Events}_\sigma$ such that $S \subseteq S'$ and S' is sync-preserving closed.

Intuitively, the set $S' = \text{SPClosure}_\sigma(S)$ captures the additional set of events that must be present in any sync-preserving correct reordering ρ of σ given that ρ contains all events in S . First, any correct reordering of σ containing S will contain $\text{TLClosure}_\sigma(S)$ and thus $\text{TLClosure}_\sigma(S) \subseteq S'$ (Condition a). Second, if a correct reordering ρ is sync-preserving and contains two acquires $a_1 \leq_{tr}^\sigma a_2$ on the same lock ℓ , then we must also have $a_1 \leq_{tr}^\rho a_2$. Then, in order to ensure well-formedness of ρ , $\text{CS}_\sigma(a_1)$ must also finish entirely before a_2 in ρ , and thus ρ must contain $\text{match}_\sigma(a_1)$ (Condition b).

For two events $e_1, e_2 \in \text{Events}_\sigma$, we define

$$\text{SPIdeal}_\sigma(e_1, e_2) = \text{SPClosure}_\sigma(\{\text{prev}_\sigma(e_1)\} \cup \{\text{prev}_\sigma(e_2)\}).$$

Here, we use $\text{prev}_\sigma(e)$ to denote the last event e' in σ such that $e' \leq_{TO}^\sigma e$; if no such event exists, then $\text{prev}_\sigma(e) = \perp$ (and further we let $\{\perp\} = \emptyset$). In essence, $\text{SPIdeal}_\sigma(e_1, e_2)$ contains the necessary set of events that must be present in any sync-preserving correct reordering that witnesses the race (e_1, e_2) . We next show that, in fact, it is also a sufficient set of events, given that it is disjoint from $\{e_1, e_2\}$.

Lemma 4.2. *(e_1, e_2) is a sync-preserving race of σ iff $\{e_1, e_2\} \cap \text{SPIdeal}_\sigma(e_1, e_2) = \emptyset$.*

Lemma 4.2 gives us a straightforward algorithm for Problem 3 — compute $I = \text{SPIdeal}_\sigma(e_1, e_2)$ and check if neither $e_1, e_2 \notin I$. In Section 4.2 we show how to compute this in linear time.

Efficiently Solving Problem 4. As noted before, a linear time algorithm for Problem 3 guarantees a *quadratic* time algorithm for Problem 4. In order to design a more efficient *linear time* algorithm, we will exploit *monotonicity* of $\text{SPIdeal}_\sigma(\cdot, \cdot)$, which we formalize next.

Lemma 4.3. *Let σ be a trace and let $e_1, e_2, e'_1, e'_2 \in \text{Events}_\sigma$ such that $e_1 \leq_{\text{TO}}^\sigma e'_1$ and $e_2 \leq_{\text{TO}}^\sigma e'_2$. Then, $\text{SPIdeal}_\sigma(e_1, e_2) \subseteq \text{SPIdeal}_\sigma(e'_1, e'_2)$.*

Our linear time algorithm for Problem 4 exploits Lemma 4.3 as follows. Suppose we are checking if a given event e in the given trace σ is in sync-preserving race with some earlier conflicting event of thread t . To accomplish this, we can scan σ and enumerate the list L of events that belong to t and conflict with e . When checking for a race with the first such event e'_{first} , we compute $I_{\text{first}} = \text{SPIdeal}_\sigma(e'_{\text{first}}, e)$. If a race is found, we are done. Otherwise, we analyze the next event e'_{next} in L and compute $I_{\text{next}} = \text{SPIdeal}_\sigma(e'_{\text{next}}, e)$. Here Lemma 4.3 ensures that $I_{\text{first}} \subseteq I_{\text{next}}$. Our algorithm exploits this observation by computing the latter set I_{next} incrementally, spending time that is proportional only to the number of *extra* events (i.e., events in $I_{\text{next}} \setminus I_{\text{first}}$). This principle is applied repeatedly to subsequent events of L , giving us an overall linear time algorithm (Section 4.3).

Efficiently Solving Problem 2. A final ingredient in our incremental linear time algorithm for Problem 2 is the following observation which builds on Lemma 4.3.

Lemma 4.4. *Let σ be a trace and let $e_1, e_2, e'_2 \in \text{Events}_\sigma$ such that $e_1 \leq_{\text{tr}}^\sigma e'_2 \leq_{\text{TO}}^\sigma e'_2$, $e_1 \asymp e_2$ and $e_1 \asymp e'_2$. If (e_1, e_2) is not a sync-preserving race, then (e_1, e'_2) is also not a sync-preserving race of σ .*

Intuitively, this observation suggests the following. Suppose that, when looking for a sync-preserving race, the algorithm determines that e_2 is not in race with any earlier conflicting event. Then, for an event e'_2 (that appears later in the thread of e_2), we only need to investigate if e'_2 is in race with conflicting events e'_1 that appear *after* e_2 in the trace (i.e., $e_2 \leq_{\text{tr}}^\sigma e'_1 \leq_{\text{tr}}^\sigma e'_2$), instead of additionally looking for races (e'_1, e'_2) where $e'_1 \leq_{\text{tr}}^\sigma e'_2$.

High-Level Overview of SyncP. Equipped with Lemma 4.3 and Lemma 4.4, we now describe our incremental algorithm for Problem 2 that works in linear time. For ease of exposition, let us focus on the question — is there a write-write race on some fixed variable $x \in \text{Vars}$ when accessed in two fixed threads $t_1, t_2 \in \text{Vars}$. The algorithm scans the trace in a streaming forward pass and analyzes every event $e = \langle t_2, w(x) \rangle$, checking if there is an earlier conflicting event $e' = \langle t_1, w(x) \rangle \leq_{\text{tr}} e$ so that (e', e) is a sync-preserving race. In doing so, it computes $I = \text{SPIdeal}_\sigma(e', e)$ in linear time and checks if $e' \notin I$. If not, (e', e) is not a race and the algorithm checks if there is a different event $e'_{\text{next}} \leq_{\text{tr}} e$ so that (e'_{next}, e) is a race. This continues until there are no earlier events remaining that conflict with e . Each time, the ideal computation is performed incrementally, by using the previously computed ideals. After this, the algorithm moves to the next write event $e_{\text{next}} = \langle t_2, w(x) \rangle$ in t_2 and checks if it is in race with some earlier event e'' of thread t_1 , where this time, e'' appears in the trace after the previously discarded event e of t_2 . Again, the closed set $\text{SPIdeal}_\sigma(e'', e_{\text{next}})$ is computed incrementally. We show that all this incremental computation can be performed efficiently and present an outline for our algorithm for Problem 2 in Section 4.4.

Next, we present high-level descriptions of the intermediate steps that we outlined above, and discuss important algorithmic insights and data-structures that help achieve efficiency.

4.2 Checking if a Given Pair of Conflicting Events is a Sync-Preserving Race

Algorithm 1 outlines our solution to Problem 3 (check if a given pair of events (e_1, e_2) is a sync-preserving race of σ). This algorithm computes the closure $\text{SPIdeal}_\sigma(e_1, e_2)$ in an iterative fashion and checks if it contains neither e_1 nor e_2 (see Lemma 4.2). We remark that when $e_1 \leq_{\text{tr}}^\sigma e_2$, Definition 2 ensures that $e_2 \notin \text{SPIdeal}_\sigma(e_1, e_2)$. Consequently, the check ' $e_1 \notin \text{SPIdeal}_\sigma(e_1, e_2)$ '

Algorithm 1: *Checking if a Given Conflicting Pair Constitutes a Sync-Preserving Race***Input:** Trace σ , Conflicting events e_1 and e_2 with $e_1 \leq_{\text{tr}}^\sigma e_2$

```

1 function ComputeSPIdeal( $\sigma, e_1, e_2, I_0$ )
2    $I \leftarrow I_0 \cup \text{TLClosure}_\sigma(\{\text{prev}_\sigma(e_1)\}) \cup \text{TLClosure}_\sigma(\{\text{prev}_\sigma(e_2)\})$ 
3   repeat
4     if  $\exists \ell \in \text{Locks}_\sigma, \exists a_1, a_2 \in \text{Acquires}_\sigma(\ell)$  such that  $a_1 \leq_{\text{tr}}^\sigma a_2$  and  $a_1, a_2 \in I$  then
5        $I \leftarrow I \cup \text{TLClosure}_\sigma(\{\text{match}_\sigma(a_1)\})$ 
6   until  $I$  does not change
7   return  $I$ 
8  $I \leftarrow \text{ComputeSPIdeal}(\sigma, e_1, e_2, \emptyset)$ 
9 if  $e_1 \notin I$  then
10  declare ‘race’

```

(Line 9 in Algorithm 1) is equivalent to the condition ‘ $\{e_1, e_2\} \cap \text{SPIdeal}_\sigma(e_1, e_2) = \emptyset$ ’ (due to Lemma 4.2). The function `ComputeSPIdeal` performs a fixpoint computation, starting from the set $I = \bigcup_{i \in \{1,2\}} \text{TLClosure}_\sigma(\{\text{prev}_\sigma(e_i)\})$ (when $I_0 = \emptyset$). The correctness of the algorithm follows from the correctness of the function `ComputeSPIdeal`, which we formalize below.

Lemma 4.5. *Let σ be a trace, $e_1, e_2 \in \text{Events}_\sigma$ and $I_0 \subseteq \text{Events}_\sigma$ be a $(\leq_{\text{TO}}^\sigma, \text{lw}_\sigma)$ -closed set. Let I be the set returned by `ComputeSPIdeal` (σ, e_1, e_2, I_0) in Algorithm 1. Then, $I = \text{SPClosure}_\sigma(I_0 \cup \{\text{prev}_\sigma(e_1)\} \cup \{\text{prev}_\sigma(e_2)\})$.*

We now discuss the data-structures used to ensure linear time and space for Algorithm 1.

Vector Timestamps. Vector timestamps [Fidge 1991; Mattern 1988] are routinely used in distributed computing and also in prior work on race prediction [Flanagan and Freund 2009; Kini et al. 2017; Pozniansky and Schuster 2003; Roemer et al. 2018]. We use vector timestamps to represent sets of events that are $(\leq_{\text{TO}}^\sigma, \text{lw}_\sigma)$ -closed; a formal definition is deferred to Section 4.4. In Algorithm 1, the sets $\text{TLClosure}_\sigma(\{\text{prev}_\sigma(e_i)\})$ are $(\leq_{\text{TO}}^\sigma, \text{lw}_\sigma)$ -closed (Line 2). Further, the initial value $I_0 = \emptyset$ ensures that all subsequent values of I in `ComputeSPIdeal` are $(\leq_{\text{TO}}^\sigma, \text{lw}_\sigma)$ -closed. All these sets can be represented as vector timestamps. The advantage of using vector timestamps is two-fold. First, these timestamps provide a succinct representation of sets — instead of representing a set explicitly as a collection of events, a vector timestamp only uses $\mathcal{T}$ integers (where $\mathcal{T} = |\text{Thr}_\sigma|$). Second, the vector timestamps for $(\leq_{\text{TO}}^\sigma, \text{lw}_\sigma)$ -closed sets can be computed in a streaming fashion, incrementally, using vector timestamps of smaller subsets.

Projecting a Trace to Threads and Locks. Let us consider the check in Line 4. Here, we look for two acquire events $a_1 \leq_{\text{tr}}^\sigma a_2$ in the current ideal I that acquire the same lock ℓ . How do we efficiently discover two such acquires? A straightforward but naive way is to enumerate all pairs of events in I and check if they are acquire events of the above kind. But this can take $O(N^3)$ time, where $N = |\text{Events}_\sigma|$ because the number of such pairs can be $O(N^2)$ in the worst case and the number of times the ideal can change is $O(N)$. Instead, we rely on the following observation:

Proposition 4.6. *Let $I \subseteq \text{Events}_\sigma$ be downward closed with respect to $\leq_{\text{TO}}^\sigma$. For every $t \in \text{Thr}_\sigma$ and every $\ell \in \text{Locks}_\sigma$, there is at most one acquire event a with $\text{thr}(a) = t$, $\text{op}(a) = \text{acq}(\ell)$ such that $a \in I$ and $\text{match}_\sigma(a) \notin I$. When such an event a exists, then $\text{match}_\sigma(a') \in I$ for every other acquire $a' <_{\text{TO}}^\sigma a$ of the form $a' = \langle t, \text{acq}(\ell) \rangle$.*

The above observation can be exploited as follows. Let I be the current ideal and let $e_{t,\ell}^I$ be the last acquire on lock ℓ performed by thread t such that $e_{t,\ell}^I \in I$. Let $\text{Acq}_\ell^I = \{e_{t,\ell}^I\}_{t \in \text{Thr}_\sigma}$ be the set of last acquire events in I of every thread. Let e_ℓ^I be the last event (according to trace order $\leq_{\text{tr}}^\sigma$) in Acq_ℓ^I . Then, for every other acquire $e' \in \text{Acq}_\ell^I \setminus \{e_\ell^I\}$, the matching release $\text{match}_\sigma(e')$ must be added in I . Hence, if we can efficiently determine the events $e_{t,\ell}^I$ each time, then we can also efficiently determine e_ℓ^I and thus efficiently perform the closure each time. So, how do we determine $e_{t,\ell}^I$ efficiently each time? We achieve this by maintaining a FIFO sequence $\text{CSHist}_{t,\ell}$, for every thread $t \in \text{Thr}_\sigma$ and lock $\ell \in \text{Locks}_\sigma$. For every critical section on lock ℓ in thread t , there is a corresponding entry in the list $\text{CSHist}_{t,\ell}$, and the order of these entries is the same as the order in which these critical sections were performed in the trace σ . Every entry (corresponding to critical section with acquire a) is a pair $(\text{TLCClosure}_\sigma(a), \text{TLCClosure}_\sigma(\text{match}_\sigma(a)))$, represented as a pair of vector timestamps $(C_a, C_{\text{match}(a)})$.

Let us now see how we perform the check in Line 4 using these data structures $\{\text{CSHist}_{t,\ell}\}_{t,\ell}$. For the current ideal I , we essentially need to determine the last acquire $e_{t,\ell}^I$ of each thread t and lock ℓ that belongs to I . This can be done by traversing the list $\text{CSHist}_{t,\ell}$ starting from the earliest entry, until we encounter the entry corresponding to the last acquire $e_{t,\ell}^I$ that belongs to I (this corresponds to a simple timestamp comparison). All entries in $\text{CSHist}_{t,\ell}$ prior to the identified event $e_{t,\ell}^I$ can then be discarded from $\text{CSHist}_{t,\ell}$, because the ideal now contains their information and only grows monotonically through the course of the fixpoint computation and the discarded entries will not be of use from now on. Thus, every entry in the lists $\{\text{CSHist}_{t,\ell}\}_{t,\ell}$ is traversed only once and the overall fixpoint computation runs in linear time when the number of $\mathcal{T}$ is constant.

4.3 Checking for a Sync-Preserving Race on a Given Event with a Given Thread

Algorithm 2: *Checking if There Is a Sync-Preserving Race on a Given Event with a Given Thread*

Input: Trace σ , Event $e = \langle t, a(x) \rangle$, Thread u

```

1 for each  $b \in \{\text{r}, \text{w}\}$  such that  $b \asymp a^3$ 
2   let  $L_b$  be the list of events  $e'$  of the form  $\langle u, b(x) \rangle$  such that  $e' \leq_{\text{tr}}^\sigma e$ 
3   for each  $b \in \{\text{r}, \text{w}\}$  such that  $b \asymp a$ 
4      $I_b \leftarrow \emptyset$ 
5     for each  $e' \in L_b$ 
6        $I_b \leftarrow \text{ComputeSPIdeal}(\sigma, e', e, I_b)$ 
7       if  $e' \notin I_b$  then
8         declare 'race' and exit
```

Let us now consider Algorithm 2. This algorithm takes as input a trace σ , an event $e = \langle t, a(x) \rangle$ and a threads $u \neq t$, and checks if there is an event e' of thread u such that $e' \leq_{\text{tr}}^\sigma e$ and (e', e) is a sync-preserving race. Algorithm 2 works as follows. For the sake of simplicity, assume that e is a read event, i.e., $a = \text{r}$. The algorithm assumes access to the list L_w of write events $e' = \langle u, \text{w}(x) \rangle$ that appear prior to e (Line 2); these lists can be constructed in linear time, in a single pass traversal of the trace. The algorithm simply traverses L_w (according to trace order $\leq_{\text{tr}}^\sigma$) and checks for races with

³ $b \asymp a$ whenever not both b and a are read (r) operations

each event in L_w , by computing the fixpoint closure sets as in Algorithm 1. Instead of computing the ideal from scratch, the algorithm exploits the monotonicity property outlined earlier in Lemma 4.3 by reusing the ideal I_w computed so far. As with Algorithm 1, this algorithm also uses vector timestamps and maintains lists $\{\text{CSHist}_{t,\ell}\}_{t \in \text{Thr}, \ell \in \text{Locks}}$ for computing successive ideals efficiently. Overall again, each entry in $\{\text{CSHist}_{t,\ell}\}_{t \in \text{Thr}, \ell \in \text{Locks}}$ is visited a constant number of times and thus Algorithm 2 runs in linear time and uses linear space.

4.4 Algorithm SyncP for Sync-Preserving Race Prediction

The pseudo-code for SyncP is presented in Algorithm 3. This is a one pass streaming algorithm that processes events as they appear in the trace, modifying its state and detecting races on the fly. The algorithm maintains several data structures including vector clocks and FIFO queues in its state. We first describe these data structures, then discuss how the algorithm initializes and modifies them as it processes the trace, and finally discuss the time and space usage for this algorithm; many of these details have already been spelt out in Sections 4.1-4.3.

Let us first briefly explain the notion of vector timestamps and vector clocks [Fidge 1991; Mattern 1988]. A vector timestamp is a mapping $V : \text{Thr}_\sigma \rightarrow \mathbb{N}$ from the threads of a trace to natural numbers, and can be represented as a vector of length $|\text{Thr}_\sigma|$. The *join* of two vector timestamps V_1 and V_2 , denoted $V_1 \sqcup V_2$ is the vector timestamp $\lambda t, \max(V_1(t), V_2(t))$. Vector timestamps can be compared in a pointwise fashion: $V_1 \sqsubseteq V_2$ iff $\forall t, V_1(t) \leq V_2(t)$. The minimum timestamp is $\perp = \lambda t, 0$. For a scalar $c \in \mathbb{N}$, we use $V[t \mapsto c]$ to denote the timestamp $\lambda u, \text{if } u = t \text{ then } c \text{ else } V(u)$. Vector clocks are variables that take values from the space of vector timestamps. We use normal font for timestamps ($C, C', I \dots$) and boldfaced font for vector clocks ($\mathbb{C}, \mathbb{LW}, \mathbb{I}, \dots$).

Data Structures and Initialization. The algorithm maintains the following data structures.

- (1) **Vector Clocks.** The algorithm uses vector clocks primarily for two purposes. First, we assign timestamps to all events in the trace and use the following vector clocks for this purpose — for every thread, a dedicated clock $\mathbb{C}_t$, and for every variable x , a dedicated clock $\mathbb{LW}_x$. The timestamp of an event e is essentially a succinct representation of the set $\text{TLClosure}(e)$. Next, the algorithm computes $\text{SPIdeal}(\cdot, \cdot)$ sets and represents them as timestamps. These are stored in vector clocks $\mathbb{I}^{(t_1, t_2, a_1, a_2, x)}$, one for every tuple $(t_1, t_2, a_1, a_2, x) \in \text{Thr} \times \text{Thr}_\times \{r, w\} \times \{r, w\} \times \text{Vars}$. All vector clocks are initialized with the timestamp $\perp = \lambda t, 0$.
- (2) **Scalars.** For every lock ℓ , the algorithm maintains a scalar variable g_ℓ to record the index of the last acquire on ℓ seen in the trace seen so far. Each such scalar is initialized with 0.
- (3) **FIFO Queues.** The algorithm maintains several FIFO queues, whose entries correspond to different events in the trace. The algorithm ensures that an event appears only once, and additionally ensures that the entries respect the order of appearance of the corresponding events in the trace. The first kind of FIFO queues are used in the fixpoint computation. For this, the algorithm maintains queues $\text{CSHist}_{t,\ell}^{(t_1, t_2, a_1, a_2, x)}$, one for every thread t , lock ℓ and tuple (t_1, t_2, a_1, a_2, x) . Each entry of $\text{CSHist}_{t,\ell}^{(t_1, t_2, a_1, a_2, x)}$ is a triplet (g_e, C_e, C'_e) and corresponds to an acquire event e of the form $\langle t, \text{acq}(\ell) \rangle$ — here, g_e is the index of e in the trace, C_e is the timestamp of e and C'_e is the timestamp of the matching release match(e). The second kind of queues are of the form $\text{AccessHist}_{t,a,x}^{(u)}$, one for each $t, u \in \text{Thr}$, $a \in \{r, w\}$ and $x \in \text{Vars}$. Entry in such a FIFO queue corresponds to access events of the form $e = \langle t, a(x) \rangle$. Each entry is of the form $(C_{\text{prev}(e)}, C_e)$ where $C_{\text{prev}(e)}$ is the timestamp of $\text{prev}(e)$ (and $\perp$ if $\text{prev}(e)$ does not exist) and C_e is the timestamp of e . All FIFO queues are empty ($\emptyset$) in the beginning. We write $L \cdot \text{first}()$ and $L \cdot \text{last}()$ to denote the first and last elements of the FIFO queue L . Further, $L \cdot \text{isEmpty}()$, $L \cdot \text{addLast}()$ and $L \cdot \text{removeFirst}()$ respectively represent functions

Algorithm 3: Detailed Streaming Algorithm for Checking Sync-Preserving Races

```

1  function Initialization()
2    foreach  $t \in \text{Thr}$  do
3       $\mathbb{C}_t := \perp$ 
4    foreach  $x \in \text{Vars}$  do
5       $\mathbb{LW}_x := \perp$ 
6    foreach  $\ell \in \text{Locks}$  do
7       $g_\ell := 0$ 
8    foreach  $t_1 \neq t_2 \in \text{Thr}, a_1 \asymp a_2 \in \{r, w\}, x \in \text{Vars}$  do
9       $\mathbb{I}^{(t_1, t_2, a_1, a_2, x)} := \perp$ 
10     foreach  $t \in \text{Thr}, \ell \in \text{Locks}$  do
11        $\text{CSHist}_{t, \ell}^{(t_1, t_2, a_1, a_2, x)} := \emptyset$ 
12     foreach  $u \in \text{Thr}$  do
13       foreach  $t \in \text{Thr}, a \in \{r, w\}, x \in \text{Vars}$  do
14          $\text{AccessHist}_{t, a, x}^{(u)} := \emptyset$ 
15  function maxLowerBound( $U, Lst$ )
16     $(g_{\max}, C_{\max}, C'_{\max}) := (0, \perp, \perp)$ 
17    while not  $Lst \cdot \text{isEmpty}()$  do
18       $(g, C, C') := Lst \cdot \text{first}()$ 
19      if  $C \sqsubseteq U$  then
20         $(g_{\max}, C_{\max}, C'_{\max}) := (g, C, C')$ 
21      else
22        break
23       $Lst \cdot \text{removeFirst}()$ 
24    return  $(g_{\max}, C_{\max}, C'_{\max})$ 
25  function ComputeSPIdeal( $I, \langle t_1, t_2, a_1, a_2, x \rangle$ )
26    repeat
27      foreach  $\ell \in \text{Locks}$  do
28        foreach  $t \in \text{Thr}$  do
29           $(g_{\ell, t}, C_{\ell, t}, C'_{\ell, t}) := \text{maxLowerBound}(I,$ 
30             $\text{CSHist}_{t, \ell}^{(t_1, t_2, a_1, a_2, x)})$ 
31           $t_{\max} := \text{argmax}_{t \in \text{Thr}} \{g_{\ell, t}\}$ 
32           $I := I \sqcup \bigsqcup_{t \neq t_{\max} \in \text{Thr}} C'_{\ell, t}$ 
33    until  $I$  does not change
34    return  $I$ 
35  function checkRace( $Lst, I, \langle t_1, t_2, a_1, a_2, x \rangle$ )
36    while not  $Lst \cdot \text{isEmpty}()$  do
37       $(C_{\text{prev}}, C) := Lst \cdot \text{first}()$ 
38       $I := \text{ComputeSPIdeal}(I \sqcup C_{\text{prev}}, \langle t_1, t_2, a_1, a_2, x \rangle)$ 
39      if  $C \not\sqsubseteq I$  then
40        declare  $\langle a_1, a_2 \rangle$  race on  $x$ 
41        break
42       $Lst \cdot \text{removeFirst}()$ 
43  handler read( $t, x$ )
44     $C_{\text{prev}} := \mathbb{C}_t$ 
45     $\mathbb{C}_t := \mathbb{C}_t[t \mapsto \mathbb{C}_t(t) + 1] \sqcup \mathbb{LW}_x$ 
46    foreach  $u \in \text{Thr}$  do
47       $\text{AccessHist}_{t, r, x}^{(u)} \cdot \text{addLast}((C_{\text{prev}}, \mathbb{C}_t))$ 
48    foreach  $u \neq t \in \text{Thr}$  do
49       $I := \mathbb{I}^{(u, t, w, r, x)} \sqcup C_{\text{prev}}$ 
50       $\mathbb{I}^{(u, t, w, r, x)} := \text{checkRace}(\text{AccessHist}_{u, w, x}^{(u)}, I,$ 
51         $\langle u, t, w, r, x \rangle)$ 
52  handler write( $t, x$ )
53     $C_{\text{prev}} := \mathbb{C}_t$ 
54     $\mathbb{C}_t := \mathbb{C}_t[t \mapsto \mathbb{C}_t(t) + 1]; \quad \mathbb{LW}_x := \mathbb{C}_t$ 
55    foreach  $u \in \text{Thr}$  do
56       $\text{AccessHist}_{t, w, x}^{(u)} \cdot \text{addLast}((C_{\text{prev}}, \mathbb{C}_t))$ 
57    foreach  $u \neq t \in \text{Thr}, a \in \{r, w\}$  do
58       $I := \mathbb{I}^{(u, t, a, w, x)} \sqcup C_{\text{prev}}$ 
59       $\mathbb{I}^{(u, t, a, w, x)} := \text{checkRace}(\text{AccessHist}_{u, a, x}^{(u)}, I,$ 
60         $\langle u, t, a, w, x \rangle)$ 
61  handler acquire( $t, \ell$ )
62     $\mathbb{C}_t := \mathbb{C}_t[t \mapsto \mathbb{C}_t(t) + 1]; \quad g_\ell := g_\ell + 1$ 
63    foreach  $t_1, t_2 \in \text{Thr}, a_1, a_2 \in \{r, w\}, x \in \text{Vars}$  do
64       $\text{CSHist}_{t, \ell}^{(t_1, t_2, a_1, a_2, x)} \cdot \text{addLast}((g_\ell, \mathbb{C}_t, \perp))$ 
65  handler release( $t, \ell$ )
66     $\mathbb{C}_t := \mathbb{C}_t[t \mapsto \mathbb{C}_t(t) + 1]$ 
67    foreach  $t_1, t_2 \in \text{Thr}, a_1, a_2 \in \{r, w\}, x \in \text{Vars}$  do
68       $\text{CSHist}_{t, \ell}^{(t_1, t_2, a_1, a_2, x)} \cdot \text{last}() \cdot$ 
69         $\text{updateRelease}(\mathbb{C}_t)$ 

```

that check for emptiness of L , add an element at the end of L and remove the earliest (first) element of L .

Let us now describe the working of SyncP. The algorithm works in a streaming fashion and processes each event e as soon as it appears, by calling the appropriate **handler** depending upon the operation performed in e (read, write, acquire or release). The argument for each handler is the thread performing the event and the object (variable or lock) accessed in the event. In each handler, the algorithm updates vector clocks to compute timestamps, and maintains the invariants of the FIFO queues. In addition, inside read and write handlers, the algorithm also checks for races using a fixpoint computation (function `ComputeSPIdeal`). We explain some of these briefly.

Computing Vector Timestamps. The algorithm computes the timestamp of an event e , denoted C_e when processing the event e . The timestamps computed in the algorithm are such that

$$\forall t \in \text{Thr}, C_e(t) = |\{f \in \text{Events} \mid \text{thr}(f) = t, f \in \text{TLClosure}(e)\}|$$

Observe that, with this invariant, we have $C_e \sqsubseteq C_{e'}$ iff $e \in \text{TLClosure}(e')$. The algorithm uses vector clocks $\{\mathbb{C}_t\}_{t \in \text{Thr}}$ and $\{\mathbb{LW}_x\}_{x \in \text{Vars}}$ and ensures that after processing an event e , (1) $\mathbb{C}_t$ stores the timestamp C_{e_t} , where e_t is the last event by thread t that occurs before e , and (2) $\mathbb{LW}_x$ stores the timestamp C_{e_x} , where e_x is the last event with $\text{op}(e_x) = w(x)$ that occurs before e . The algorithm correctly maintains these values by appropriate vector clock operations on Lines 45, 53, 60 and 64.

Let us now describe the invariant for the vector clock $\mathbb{I}^{(t_1, t_2, a_1, a_2, x)}$ (given tuple (t_1, t_2, a_1, a_2, x)). Let $e_{t_i, a_i, x}$ be the last event with $\text{thr}(e_{t_i, a_i, x}) = t_i$ and $\text{op}(e_{t_i, a_i, x}) = a_i(x)$ ($i \in \{1, 2\}$). Let $I \subseteq \text{Events}_\sigma$ be defined as follows. If $e_{t_1, a_1, x}$ does not exist, then $I = \text{SPClosure}_\sigma(\{\text{prev}_\sigma(e_2)\})$. Otherwise, let e be the first event in σ (according to trace order $\leq_{\text{tr}}^\sigma$) with $\text{thr}(e) = t_1$ and $\text{op}(e) = a_1(x)$ such that (e, e_2) is a sync-preserving race of σ . If no such event exists, then let $e = e_{t_1, a_1, x}$. Then, $I = \text{SPIdeal}_\sigma(e_1, e_2)$. Then, the timestamp stored in $\mathbb{I}^{(t_1, t_2, a_1, a_2, x)}$ is $\bigsqcup_{u \in \text{Thr}_\sigma} C_{e_u^I}$, where e_u^I is the last event of thread u which is in I .

Checking Races. When processing an access event $e = \langle t, a(x) \rangle$, the algorithm checks for a race as follows. For every other thread u and for every other conflicting type $b \asymp a$, the algorithm calls `checkRace` with the list $Lst = \text{AccessHist}_{u, b, x}^{(t)}$ and the timestamp representation of the current ideal as argument. This function, similar to Algorithm 2, scans Lst and reports races by repeatedly performing fixpoint computations and checking membership in some set (using the timestamp comparison in Line 38). The closure computation is performed using the optimizations discussed in Section 4.2 (use of FIFO queues $\text{CSHist}_{t, \ell}^{(\dots)}$).

Space Optimizations. Observe that, for a given thread t and lock ℓ , the algorithm, as presented, maintains $4\mathcal{T}^2 \cdot \mathcal{V}$ FIFO queues $\{\text{CSHist}_{t, \ell}^{(t_1, t_2, a_1, a_2, x)}\}_{t_1, t_2 \in \text{Thr}, a_1, a_2 \in \{r, w\}, x \in \text{Vars}}$. The total number of entries across these queues will then be $O(\mathcal{A} \cdot 4\mathcal{T}^2 \cdot \mathcal{V})$. To this end, we observe that all the above data structures essentially have the same content, and are suffixes of a common queue corresponding to the critical sections on ℓ in thread t . Indeed, we exploit this redundancy and instead maintain a common underlying data-structure that stores all entries corresponding to acquires and releases on ℓ in t , and maintain a pointer for each $\langle t_1, t_2, a_1, a_2, x \rangle$. Such a pointer keeps track of the starting index of the FIFO queue $\text{CSHist}_{t, \ell}^{(t_1, t_2, a_1, a_2, x)}$. With this optimization, we only store $O(\mathcal{A})$ entries along with additional $4\mathcal{T}^3 \cdot \mathcal{V} \cdot \mathcal{L}$ pointers, one for every tuple $\langle t_1, t_2, a_1, a_2, x \rangle$ and every shared queue indexed by (t, ℓ) . The same observations also apply to the FIFO queues $\{\text{AccessHist}_{t, a, x}^{(u)}\}_{u \in \text{Thr}}$.

Lemma 4.7 (Correctness). *For every access event e in the input trace σ , Algorithm 3 declares a race when processing e iff there is an event e' such that $e' \leq_{\text{tr}}^\sigma e$ and (e', e) is a sync-preserving race of σ .*

The proof of Lemma 4.7 follows directly from the correctness of the semantics of clocks and other observations outlined in Section 4.1.

	t_1	t_2		t_1	t_2		t_1	t_2		t_1	t_2
1	acq(b_1)		13	acq(a_2)		25		acq(a_1)	37		acq(b_2)
2	acq(b_2)		14	acq(b_1)		26		acq(a_2)	38		acq(a_1)
3	acq(c)		15	r ₃ (x)		27		acq(c)	39		acq(c)
4	w ₁ (x)		16	rel(b_1)		28		w ₁ (x)	40		w ₃ (x)
5	rel(c)		17	rel(a_2)		29		rel(c)	41		rel(c)
6	rel(b_2)		18	acq(a_1)		30		rel(a_2)	42		rel(a_1)
7	rel(b_1)		19	acq(a_2)		31		rel(a_1)	43		rel(b_2)
8	acq(a_1)		20	acq(c)		32		acq(b_1)	44		acq(b_1)
9	acq(b_2)		21	w ₄ (x)		33		acq(a_2)	45		acq(b_2)
10	r ₂ (x)		22	rel(c)		34		r ₂ (x)	46		acq(c)
11	rel(b_2)		23	rel(a_2)		35		rel(a_2)	47		w ₄ (x)
12	rel(a_1)		24	rel(a_1)		36		rel(b_1)	48		rel(c)
									49		rel(b_2)
									50		rel(b_1)

Fig. 4. Construction of the trace σ on input $s = u\#v$ where $u = 1001$ and $v = 1011$. Observe that (e_{15}, e_{40}) is a sync-preserving race, which encodes that $e[3] \neq v[3]$.

The time complexity of the algorithm can be determined as follows. The algorithm visits each entry in the FIFO queues $\text{AccessHist}_{t,a,x}^{(u)}$ once, performing constant number of vector clock operations, each running in $O(\mathcal{T})$ time. The total length of all these queues is $O(\mathcal{T} \cdot \mathcal{N})$ (more precisely, the number of access events in the trace). Similarly, the algorithm visits each entry in the FIFO queues $\text{CSHist}_{t,\ell}^{(t_1, t_2, a_1, a_2, x)}$ once, performing constantly many vector clock operations. The total number of entries in these queues is $O(\mathcal{T}^2 \cdot \mathcal{V} \cdot \mathcal{A})$. This gives us the following complexity for SyncP.

Lemma 4.8 (Complexity). *Let σ be a trace with $\mathcal{T}$ threads, $\mathcal{L}$ locks, $\mathcal{V}$ variables and $\mathcal{N}$ events, of which $\mathcal{A}$ are acquire events. Then, Algorithm 3 runs in time $O(\mathcal{N} \cdot \mathcal{T}^2 + \mathcal{A} \cdot \mathcal{V} \cdot \mathcal{T}^3)$ and uses space $O(\mathcal{N} + \mathcal{T}^3 \cdot \mathcal{V} \cdot \mathcal{L})$ on input σ .*

The proof of Theorem 3.1 follows from Lemma 4.7 and Lemma 4.8.

4.5 Linear Space Lower Bound

In this section we prove the lower-bounds of Theorem 3.2 and Theorem 3.3, i.e., that any streaming algorithm for sync-preserving race prediction must essentially use linear space, while the time-space product of any algorithm for the problem must be quadratic in the length of the input trace.

The language $\mathcal{L}_n$. Given a natural number n , we define the equality language $\mathcal{L}_n = \{u\#v : u, v \in \{0, 1\}^n \text{ and } u = v\}$, i.e., it is the language of two n -bit strings that are separated by $\#$ and are equal.

Lemma 4.9. *Any streaming algorithm that recognizes $\mathcal{L}_n$ uses $\Omega(n)$ space.*

Reduction from $\mathcal{L}_n$ recognition to sync-preserving race prediction. Consider the language $\mathcal{L}_n$ for some n . We describe a transducer $\mathcal{A}_n$ such that, on input a string $s = u\#v$, the output $\mathcal{A}_n(s)$ is a trace σ with 2 threads, $O(n \cdot \log n)$ events, $2 \cdot \log n + 1$ locks and a single variable such that the following hold.

- (1) If $s \notin \mathcal{L}_n$, then σ has no predictable race.
- (2) If $s \in \mathcal{L}_n$, then σ has a single predictable race, which is a sync-preserving race.

Moreover, $\mathcal{A}_n$ uses $O(\log n)$ working space. The transducer $\mathcal{A}_n$ uses a single variable x , two sets of locks $A = \{a_1, \dots, a_{\log n}\}$ and $B = \{b_1, \dots, b_{\log n}\}$, plus one additional lock c . The trace σ consists of two local traces π_1, π_2 of threads t_1 and t_2 which encode the bits of u and v , respectively.

- (1) The local trace π_1 is constructed as follows. For every $i \in [n]$, π_1 contains an event e_i^1 , which is a write event $w(x)$ if $u[i] = 1$, and a read event $r(x)$ otherwise. The events e_i^1 are surrounded by locks from A and B arbitrarily, as long as the following holds. For any $i < j$, we have

$$\begin{aligned} \text{locksHeld}_\sigma(e_i^1) \cap A &\not\subseteq \text{locksHeld}_\sigma(e_j^1) \cap A \\ \text{and } \text{locksHeld}_\sigma(e_i^1) \cap B &\not\subseteq \text{locksHeld}_\sigma(e_j^1) \cap B. \end{aligned}$$

Here, the locks held at an event e has the obvious meaning: $\text{locksHeld}_\sigma(e) = \{\ell \in \text{Locks}_\sigma \mid \exists a \in \text{Acquires}_\sigma(\ell) \text{ such that } e \in \text{CS}_\sigma(a)\}$.

The above property can be easily met, for example, by making $\mathcal{A}_n$ perform a breadth-first traversal of the subset-lattice of A (resp., B) starting from the top (resp., bottom). Given the current i , the transducer surrounds e_i^1 with the locks of the current element in the corresponding lattice. Finally, every write event e_i^1 is surrounded by the lock c .

- (2) The local trace π_2 is similar to π_1 , i.e., we have an event e_i^2 for each $i \in [n]$, which is a write event $w(x)$ if $u[i] = 1$, otherwise it is a read event $r(x)$. The locks that protect e_i^2 satisfy

$$\text{locksHeld}_\sigma(e_i^2) \cap (A \cup B) = A \cup B \setminus \text{locksHeld}_\sigma(e_j^1).$$

Finally, similarly to π_1 , every write event e_i^2 is surrounded by the lock c .

See Figure 4 for an illustration. Observe that $\mathcal{A}_n$ uses $O(\log n)$ bits of memory, for storing a bit-set of locks for each set A and B that must surround the current event e_i^1 and e_i^2 .

The key idea of the construction is the following. Any two events e_i^1, e_j^2 are surrounded by a common lock from the set $A \cup B$ iff $i \neq j$. Hence, (e_i^1, e_j^2) may be a predictable race of σ only if $i = j$. In turn, if $u[i] = v[j]$, then either both events are read events, or both are write events. In the former case the events are not conflicting, while in the latter case the two events are surrounded by lock c . In both cases no race occurs between e_i^1 and e_j^2 . On the other hand, if $u[i] \neq v[j]$, then one event is a read and the other is a write event. Hence, the two events conflicting, and one of them is not surrounded by lock c , thereby constituting a predictable race. The following lemma makes the above insight formal and establishes the correctness of our reduction.

Lemma 4.10. *The following assertions hold.*

- (1) *If $s \in \mathcal{L}_n$, then σ has no predictable race.*
- (2) *If $s \notin \mathcal{L}_n$, then σ has a single predictable race, which is a sync-preserving race.*

We now prove Theorem 3.2, and refer to [Mathur et al. 2020b] for the proof of Theorem 3.3.

PROOF OF THEOREM 3.2. Consider any algorithm A_1 for sync-preserving race prediction, executed in the family of traces σ constructed in our above reduction. Let $m = n/\log n$, and assume towards contradiction that A_1 uses $o(m)$ space. Then we can pair A_1 with the transducer $\mathcal{A}_m$, and obtain a new algorithm A_2 for recognizing $\mathcal{L}_m$. Since $\mathcal{A}_m$ uses $O(\log m)$ space, the space complexity of A_2 is $o(m)$. However, this contradicts Lemma 4.9. The desired result follows. $\square$

5 BEYOND SYNCHRONIZATION-PRESERVING RACES

In this section we explore the problem of dynamic race prediction beyond sync-preservation. We show Theorem 3.4, i.e., that even when just two critical sections are present in the input trace, predicting races with witnesses that might reverse the order of the critical sections becomes intractable. Our reduction is from the realizability problem of Rf-posets, which we present next.

Rf-Posets. An rf-poset is a triplet $\mathcal{P} = (X, P, \text{RF})$, where X is a set of read and write events, P defines a partial order $\leq_P$ over X , and $\text{RF}: \text{Rds}(X) \rightarrow \text{Wts}(X)$ is a reads-from function that maps

every read event of X to a write event of X . Given two distinct events $e_1, e_2 \in X$, we write $e_1 \parallel_P e_2$ to denote $e_1 \not\leq_P e_2$ and $e_2 \not\leq_P e_1$. Given a set $Y \subseteq X$, we denote by $P|Y$ the *projection* of P on Y , i.e., we have $\leq_{P|Y} \subseteq Y \times Y$, and for all $e_1, e_2 \in Y$, we have $e_1 \leq_{P|Y} e_2$ iff $e_1 \leq_P e_2$. Given a partial order Q over X , we say that Q *refines* P denoted $Q \sqsubseteq P$ if for every two events $e_1, e_2 \in X$, if $e_1 \leq_P e_2$ then $e_1 \leq_Q e_2$. We consider that each event of X belongs to a unique thread, and there is thread order $\leq_{TO}$ that defines a total order on the events of X that belong to the same thread, and P agrees with $\leq_{TO}$. The number of threads of $\mathcal{P}$ is the number of threads of the events of X .

The Realizability problem of rf-posets. Given an rf-poset $\mathcal{P} = (X, P, \text{RF})$, the *realizability problem* is to decide whether P can be linearized to a total order σ such that $\text{lw}_\sigma = \text{RF}$. It has long been known that the problem is NP-complete [Gibbons and Korach 1997], while it was recently shown that it is even W[1]-hard [Mathur et al. 2020a].

Our proof of the lower bound of Theorem 3.4 is by a two-step reduction. First we define a variant of the realizability problem for rf-posets, namely *reverse rf-realizability*, and show that it is W[1]-hard when parameterized by the number of threads. Afterwards, we reduce reverse rf-realizability to the decision problem of dynamic race prediction, which concludes the hardness of the latter.

Rf-Triplets. Given an RF-poset $\mathcal{P} = (X, P, \text{RF})$, an *rf-triplet* of $\mathcal{P}$ is a tuple $\lambda = (\bar{w}, r, \bar{w}')$ such that (i) r is a read event, (ii) $\text{RF}(r) = \bar{w}$, and (iii) $\bar{w} \asymp \bar{w}'$. We refer to $\bar{w}$, r and $\bar{w}'$ as the *write*, *read*, and *interfering write* event of λ , respectively. We denote by $\text{Triplets}(\mathcal{P})$ the set of rf-triplets of $\mathcal{P}$.

We next define a variant of rf-poset realizability, and show that, like the original problem, it is W[1]-hard parameterized by the number of threads.

Reverse Rf-Poset Realizability. The input is a tuple $(\mathcal{P}, \lambda, \sigma)$, where $\mathcal{P} = (X, P, \text{RF})$ is an rf-poset, $\lambda = (\bar{w}, r, \bar{w}')$ is a distinguished triplet of $\mathcal{P}$, and σ is a witness to the realizability of $\mathcal{P}$ such that $\bar{w}' <_\sigma \bar{w}$. The task is to determine whether $\mathcal{P}$ has a linearization σ' with $\bar{w} <_{\sigma'} \bar{w}'$. In words, $\mathcal{P}$ is already realizable by a witness that orders $\bar{w}'$ before $\bar{w}$, and the task is to decide whether $\mathcal{P}$ also has a witness in which this order is reversed.

Hardness of Reverse Rf-Poset Realizability. We show that the problem is W[1]-hard when parameterized by the number of threads of the rf-poset. Our reduction is from rf-realizability. We first present the construction and then argue about its correctness.

Construction. Consider an rf-poset $\mathcal{P} = (X, P, \text{RF})$ with k threads, and we construct an instance of reverse rf-poset realizability $(\mathcal{P}' = (X', P', \text{RF}'), \lambda, \sigma)$ with $k' = O(k^2)$ threads. We refer to Figure 5 for an illustration. For simplicity of presentation, we assume wlog that the following hold.

- (1) X contains only the events of the triplets of $\mathcal{P}$.
- (2) For every read event r , we have $\text{thr}(r) = \text{thr}(\bar{w})$, i.e., every read observes a local write event.

Let $\{X_i\}_{1 \leq i \leq k}$ be a partitioning of X such that each $\pi_i = P|X_i$ is a total order containing all events of thread i (i.e., it is the thread order for thread i). We first construct the rf-poset $\mathcal{P}' = (X', P', \text{RF}')$. The threads of $\mathcal{P}'$ are defined implicitly by the sets of events for each thread. In particular, X' is partitioned in the sets X_i (which are the events of $\mathcal{P}$), as well as two sets X_i^j and Y_i^j for each $i, j \in [k]$ with $i \neq j$, where each such X_i^j and Y_i^j contains events of a unique thread of $\mathcal{P}'$. Finally, we have two threads containing the events of the distinguished triplet λ . Hence, $\mathcal{P}'$ has $k' = k + k \cdot (k - 1) + 2 = k^2 + 2$ threads.

We first define the set of triplets $\text{Triplets}(\mathcal{P})$, which defines the event set X' and the observation function RF' . We have $X \subseteq X'$ and $\text{Triplets}(\mathcal{P}) \subseteq \text{Triplets}(\mathcal{P}')$. In addition, we create a distinguished triplet $\lambda = (\bar{w}, r, \bar{w}')$, and all its events are in X' . Finally, for every $i, j \in [k]$ with $i \neq j$, we have $X_i^j, Y_i^j \subseteq X'$, where the sets X_i^j and Y_i^j are constructed as follows. We call a pair of events

Fig. 5. Reduction of rf-poset realizability (5a) to reverse rf-poset realizability (5b).

$(e_1, e_2) \in X_i \times X_j$ with $e_1 <_P e_2$ *dominant* if for any pair $(e'_1, e'_2) \in X_i \times X_j$ such that $e_1 \leq_P e'_1$, and $e'_2 \leq_P e_2$, and $e'_1 <_P e'_2$, we have $e'_i = e_i$ for each $i \in [2]$. In words, a dominant pair identifies an ordering in P that cannot be inferred transitively by other orderings. For every dominant pair $(e_1, e_2) \in X_i \times X_j$ we create a triplet $(w_{e_1, e_2}, r_{e_1, e_2}, w'_{e_1, e_2})$, and let $w_{e_1, e_2}, r_{e_1, e_2} \in X_i^j$ and $w'_{e_1, e_2} \in Y_i^j$.

We now define the partial order P' . For every triplet (w, r, w') of $\mathcal{P}'$, we have $w <_{P'} r$. For every $i, j \in [k]$ with $i \neq j$, for every two events $e_1, e'_1 \in X_i$ such that $e_1 <_P e'_1$, for every two events $e_2, e'_2 \in X_j$ such that $e_2 <_P e'_2$, if r_{e_1, e_2} and $w'_{e'_1, e'_2}$ are events of X' (i.e., (e_1, e_2) and (e'_1, e'_2) are dominant pairs), we have (i) $r_{e_1, e_2} <_{P'} w'_{e'_1, e'_2}$ and (ii) $w'_{e_1, e_2} <_{P'} w'_{e'_1, e'_2}$. Finally, for every triplet of the form $(w_{e_1, e_2}, r_{e_1, e_2}, w'_{e_1, e_2})$, we have

$$e_1 <_{P'} r_{e_1, e_2} \quad \text{and} \quad w'_{e_1, e_2} <_{P'} e_2 \quad \text{and} \quad w_{e_1, e_2} <_{P'} \bar{r} \quad \text{and} \quad \bar{w}' <_{P'} w'_{e_1, e_2}.$$

The following lemma establishes that P' is indeed a partial order.

Lemma 5.1. P' is a partial order.

We now turn our attention to the solution σ of $\mathcal{P}'$, which is constructed in two steps. First, we construct a partial order $Q \subseteq P'$ over X' which orders in every triplet the interfering write before the write of the triplet. That is, for every triplet (w, r, w') of $\mathcal{P}'$, we have $w' <_Q w$. Then, we obtain σ by linearizing Q arbitrarily. The following lemma states that σ witnesses the realizability of $\mathcal{P}'$.

Lemma 5.2. The trace σ realizes $\mathcal{P}'$.

Observe that the size of $\mathcal{P}'$ is polynomial in the size of $\mathcal{P}$. The following lemma states the correctness of the reduction. Here we sketch the argument while the detailed proof is in our companion technical report [Mathur et al. 2020b].

Lemma 5.3. Reverse rf-poset realizability is $W[1]$ -hard parameterized by the number of threads.

Correctness. We now present the key insight behind the correctness of the reduction. Consider any dominant pair of events (e_1, e_2) in the initial rf-poset, i.e., we have $e_1 <_P e_2$. Observe that the two events are unordered in P' . Now consider any trace σ^* that solves the reverse rf-poset realizability problem for $\mathcal{P}'$. By definition, σ^* must reverse the order of the two writes of the conflicting triplet, i.e., we must have $\bar{w} <_{\sigma^*} \bar{w}'$.

- (1) Since $\bar{w} <_{\sigma^*} \bar{w}'$, we have $\bar{r} <_{\sigma^*} \bar{w}'$, so the last write of $\bar{r}$ is not violated in σ^* .
- (2) Since $w_{e_1, e_2} <_{P'} \bar{r}$, by the previous item we also have transitively $w_{e_1, e_2} <_{P'} \bar{w}'$, and since $\bar{w}' <_{P'} \bar{w}'_{e_1, e_2}$, we have, transitively $w_{e_1, e_2} <_{\sigma^*} w'_{e_1, e_2}$.

Fig. 6. Example of our reduction of an instance of reverse rf-poset realizability (Figure 6a) to an instance of dynamic data-race prediction (Figure 6b) on the event pair (e_4, e_{18}) .

(3) Since $w_{e_1, e_2} <_{\sigma^*} w'_{e_1, e_2}$, we have $r_{e_1, e_2} <_{\sigma^*} w'_{e_1, e_2}$, so the last write of r_{e_1, e_2} is not violated in σ^* .

(4) Finally, since $e_1 <_{P'} r_{e_1, e_2}$ and $w'_{e_1, e_2} < e_2$, we also have, transitively, that $e_1 <_{\sigma^*} e_2$.

Hence, the witness σ^* also respects the partial order P , and thus also serves as a witness of the realizability of $\mathcal{P}$ (when projected to the set of events X). Thus, if reverse rf-poset realizability holds for $\mathcal{P}'$, then rf-poset realizability holds for $\mathcal{P}$. The inverse direction is similar.

Hardness of 1-Reversal Dynamic Race Prediction. We are now ready to prove our second step of the reduction, i.e., to establish an FPT reduction from reverse rf-poset realizability to the decision problem of dynamic race prediction. We first describe the construction.

Consider an instance $(\mathcal{P} = (X, P, RF), \lambda = (\bar{w}, \bar{r}, \bar{w}'), \sigma)$ of reverse rf-poset realizability, and we construct a trace σ' such a specific event pair of σ' is a predictable race iff $\mathcal{P}$ is realizable by a witness that reverses λ . We assume wlog that X contains only events that appear in triplets of $\mathcal{P}$. We construct σ' by inserting various events in σ , as follows. Figure 6 provides an illustration.

- (1) For every dominant pair (e_1, e_2) of $\mathcal{P}$, we introduce a new variable x_{e_1, e_2} , and a write event $w(x_{e_1, e_2})$ and a read event $r(x_{e_1, e_2})$. We make $\text{thr}(w(x_{e_1, e_2})) = \text{thr}(e_1)$ and $\text{thr}(r(x_{e_1, e_2})) = \text{thr}(e_2)$. Finally, we thread-order $w(x_{e_1, e_2})$ after e_1 and $r(x_{e_1, e_2})$ before e_2 . Notice that any correct reordering σ^* of σ' must order $w(x_{e_1, e_2}) \leq_{\text{tr}}^{\sigma^*} r(x_{e_1, e_2})$, and thus, transitively, also order $e_1 \leq_{\text{tr}}^{\sigma^*} e_2$.
- (2) For every thread $t_i \neq \text{thr}(\bar{w})$, $t_i \neq \text{thr}(\bar{w}')$ we introduce a new variable x^i , a write event $w(x^i)$, and a read event $r(x^i)$. We make $\text{thr}(w(x^i)) = t_i$ and $\text{thr}(r(x^i)) = \text{thr}(\bar{r})$. We thread-order each $w(x^i)$ as the last event of t_i , and thread-order all $r(x^i)$ as final events of $\text{thr}(\bar{r})$ so far.
- (3) We introduce a new variable y , and a write event $w(y)$ and a read event $r(y)$. We make $\text{thr}(w(y)) = \text{thr}(\bar{w}')$ and $\text{thr}(r(y)) = \text{thr}(\bar{w})$. Finally, we thread-order $w(y)$ and $r(y)$ at the end of their respective threads. In particular, $r(y)$ is thread-ordered after the events $r(x^i)$

introduced in the previous item. Notice that because of this ordering and the previous item, any correct reordering σ^* of σ' must contain all events of X' .

- (4) We introduce a lock ℓ and two pairs of lock-acquire and lock-release events $(\text{acq}_i(\ell), \text{rel}_i(\ell))$, for each $i \in [2]$. We make $\text{thr}(\text{acq}_i(\ell)) = \text{thr}(\text{rel}_i(\ell)) = t_j$, where $t_j = \text{thr}(\bar{w}')$ if $i = 1$ and $t_j = \text{thr}(\bar{w})$ otherwise. Finally, we surround with the critical section of $\text{acq}_1(\ell), \text{rel}_1(\ell)$ all events of the corresponding thread, and surround with the critical section of $\text{acq}_2(\ell), \text{rel}_2(\ell)$ the event $\bar{w}$. Notice that any correct reordering σ^* that witnesses a race on $(\bar{w}(y), r(y))$ is missing $\text{rel}_1(\ell)$, and thus must order $\text{rel}_2(\ell) \leq_{\text{tr}}^{\sigma^*} \text{acq}_1(\ell)$. In turn, this leads to a transitive ordering $\bar{w} \leq_{\text{tr}}^{\sigma^*} \bar{w}'$, and since the last write of $\bar{r}$ must be $\text{lw}_{\sigma^*}(\bar{r}) = \bar{w}$, we must have $\bar{r} \leq_{\text{tr}}^{\sigma^*} \bar{w}'$.

We now outline the correctness of the reduction (proof in [Mathur et al. 2020b]). Consider any correct reordering σ^* that witnesses a predictable race $(\bar{w}(y), r(y))$ on σ' . Item 2 and Item 3 above guarantee that $X \subseteq \text{Events}_{\sigma^*}$, while Item 1 guarantees that σ^* linearizes P , and Item 4 guarantees that σ^* reverses λ , i.e., $\bar{r} \leq_{\text{tr}}^{\sigma^*} \bar{w}'$. Finally, note that σ' has size that is polynomial in n , while the number of threads of σ' equals the number of threads of $\mathcal{P}$. This concludes the proof of Theorem 3.4.

6 EXPERIMENTS

In this section we report on an implementation and experimental evaluation of the techniques presented in this work. Our objective is two-fold. The first goal is to quantify the practical relevance of sync-preservation, i.e., whether in practice the definition captures races that are missed by the standard notion of happens-before and WCP [Kini et al. 2017] races. The second goal is to evaluate the performance of our algorithm SyncP for detecting sync-preserving races.

6.1 Experimental Setup

We have implemented SyncP (Algorithm 3) for predicting all sync-preserving races in our tool RAPID [Mathur 2020], written in Java, and evaluated it on a standard set of benchmarks.

Benchmarks. Our benchmark set consists of standard benchmarks found in the recent literature [Huang et al. 2014; Kini et al. 2017; Mathur et al. 2018; Pavlogiannis 2019; Roemer et al. 2018; Yu et al. 2018]. It consists of 30 concurrent programs taken from standard benchmark suites: (i) the IBM Contest benchmark suite [Farchi et al. 2003], (ii) the Java Grande forum benchmark suite [Smith et al. 2001], (iii) the DaCapo benchmark suite [Blackburn et al. 2006], (iv) the Software Infrastructure Repository [Do et al. 2005], and (v) some standalone benchmarks. For each benchmark, we generated a single trace using RV-Predict [Rosu 2018] and evaluated all methods on the same trace.

Compared Methods. We compare our algorithm with state-of-the-art sound race detectors, namely, SHB [Mathur et al. 2018], WCP [Kini et al. 2017] and M2 [Pavlogiannis 2019]. Recall that SHB and WCP are linear time algorithms that perform a single pass of the input trace σ . SHB computes happens-before races and is sound even beyond the first race. On the other hand, WCP is only sound for the first race report. In order to allow WCP to soundly report more than one race, whenever a race is reported on an event pair (e_1, e_2) (i.e., we have $e_1 \parallel_{\text{WCP}}^{\sigma} e_2$), we force an order $e_1 \leq_{\text{WCP}}^{\sigma} e_2$ before proceeding with the next event of σ . This is a standard practice that has been followed in other works, e.g., [Pavlogiannis 2019; Roemer et al. 2018]. Finally, M2 is a more heavyweight algorithm that makes sound reports for all races by design, though its running time is a larger polynomial (of order n^4) [Pavlogiannis 2019].

Optimizations. In general, the benchmark traces can be huge and often scale to sizes of order as large as 10^8 . A closer inspection shows that many events, even though they perform accesses to shared memory, are non-racy and even totally ordered by fork-join mechanisms and data flows in the trace. We have implemented a lightweight, linear time, single-pass optimization of the input

trace σ that filters out such events. The optimization simply identifies memory locations x whose conflicting accesses are totally ordered in σ by thread and data-flow orderings, and ignores all such accesses in σ . For a fair comparison, we employ the optimization in all compared methods. This is similar to FastTrack-like optimizations [Flanagan and Freund 2009], which identify and ignore thread-local events. We note that, as each of the compared methods attempts to report as many races as possible, epoch-like optimizations were not applied.

Reported Results. Our experiments were conducted on a 2.6GHz 64-bit Linux machine with Java 1.8 as the JVM and 30GB heap space. Each of the compared methods is evaluated on the same input trace σ . For every such input, the respective method reports the following race warnings.

- (1) *Racy events.* We report the number of events e_2 such that there is an event e_1 with $e_1 <_{tr} e_2$ for which a race (e_1, e_2) is detected. We remark that this is the standard way of reporting race warnings [Flanagan and Freund 2009; Genç et al. 2019; Kini et al. 2017; Mathur et al. 2018; Roemer et al. 2018], as it allows for one-pass, linear time algorithms that avoid the overhead of testing for races between all possible $\Theta(n^2)$ pairs of events.
- (2) *Racy source-code lines.* We report the number of distinct source-code lines which correspond to events e_2 that are found as racy in Item 1. This is a meaningful measure, as the same source-code line might be reported by many different events e_2 .
- (3) *Racy memory locations.* We report the number of different memory locations that are accessed by all the events e_2 that are found as racy in Item 1.
- (4) *Running time.* We measure the time of the algorithm required to process each benchmark, while imposing a 1-hour timeout (TO).

6.2 Experimental Results

We now turn our attention to the experimental results. Table 1 shows the races and running times reported by each method on each benchmark.

Coverage of Sync-Preserving Races. We find that *every* race reported by SHB or WCP is a sync-preserving race, also reported by SyncP. On the other hand, bold-face entries highlight benchmarks which have sync-preserving races that are not happens-before races. We see that such races are found in 11 out of 30 benchmarks. Interestingly, in the 5 most challenging out of these 11 benchmarks, the same pattern occurs if we focus on source-code lines (i.e., the entries in the parentheses). Hence, for these benchmarks, sync-preservation is *necessary* to capture many racy source-code lines, which happens-before would completely miss. We also remark that the more heavyweight analysis M2 misses several of these races due to frequent timeouts. In total, we have 18 unique source-code lines that are racy but only detected by SyncP. On the other hand, there are only 2 source-code lines that are caught by M2 but not by SyncP.

Running Times. Our experimental times indicate that SyncP is quite efficient in practice. Among all algorithms, SyncP is the second fastest, being about 1.4 times slower than the fastest, lightweight SHB, while at the same time, being able to detect considerably more races than SHB (i.e., 1342 more racy events, and 21 more racy source-code lines). On the other hand, SyncP detects even more races than M2, due to timeouts, and even has almost equal detection capability with M2 on the cases that M2 does not time out. Due to the slow performance of M2 (i.e., over 8.5 hours and with several timeouts), we exclude it from the more refined analysis that follows.

Racy Memory Locations. We next proceed to evaluate the capability of SyncP in detecting racy memory locations. As all races detected by SHB or WCP are sync-preserving, the same follows for the racy memory locations, i.e., they are all detected as racy by SyncP. On the other hand, Table 2 shows a few cases in which SyncP has discovered racy variables that are missed by SHB and WCP.

Table 1. Dynamic race reports in our benchmarks. $\mathcal{N}$ and $\mathcal{T}$ denote the number of events and number of threads in the respective trace. For races, an entry ' r (s)' denotes the number r of events e_2 found to be in race with an earlier event e_1 , as well as the number s of unique source-code lines corresponding to such events e_2 . Bold-face entries highlight cases where there are sync-preserving races that are not happens-before races.

Benchmark	$\mathcal{N}$	$\mathcal{T}$	SHB		WCP		M2		SyncP	
			Races	Time	Races	Time	Races	Time	Races	Time
array	51	4	0 (0)	0.02s	0 (0)	0.03s	0 (0)	0.09s	0 (0)	0.04s
critical	59	5	3 (3)	0.19s	1 (1)	0.03s	3 (3)	0.11s	3 (3)	0.07s
account	134	5	3 (1)	0	3 (1)	0.06s	3 (1)	0.23s	3 (1)	0.09s
airtickets	140	5	8 (3)	0.02s	5 (2)	0.03s	8 (3)	0.13s	8 (3)	0.05s
pingpong	151	7	8 (3)	1.09s	8 (3)	0.04s	8 (3)	0.17s	8 (3)	0.06s
twostage	193	13	4 (1)	0.02s	4 (1)	0.09s	4 (1)	0.20s	4 (1)	0.10s
wronglock	246	23	12 (2)	0.02s	3 (2)	0.09s	25 (2)	0.43s	25 (2)	0.15s
bbuffer	332	3	3 (1)	0.01s	1 (1)	0.05s	3 (1)	0.11s	3 (1)	0.06s
prodcons	658	9	1 (1)	0.03s	1 (1)	0.09s	1 (1)	0.20s	1 (1)	0.10s
clean	1.0K	10	59 (4)	0.04s	33 (4)	0.14s	110 (4)	0.85s	60 (4)	0.17s
mergesort	3.0K	6	1 (1)	11m10s	1 (1)	0.12s	5 (2)	0.96s	3 (1)	0.13s
bubblesort	4.0K	13	269 (5)	0.03s	100 (5)	0.27s	374 (5)	8.05s	269 (5)	0.50s
lang	6.0K	8	400 (1)	0.10s	400 (1)	0.23s	400 (1)	1.31s	400 (1)	0.31s
readswrites	11K	6	92 (4)	0.12s	92 (4)	0.41s	228 (4)	12.74s	199 (4)	0.77s
raytracer	15K	4	8 (4)	0.02s	8 (4)	0.30s	8 (4)	0.40s	8 (4)	0.30s
bufwriter	22K	7	8 (4)	0.10s	8 (4)	0.70s	8 (4)	2.65s	8 (4)	0.84s
ftpserver	49K	12	69 (21)	6.91s	69 (21)	1.34s	85 (21)	4.11s	85 (21)	4.69s
moldyn	200K	4	103 (3)	0.05s	103 (3)	1.83s	103 (3)	1m25s	103 (3)	1.86s
linkedlist	1.0M	13	5.0K (4)	7.25s	5.0K (3)	27.07s	TO	TO	7.0K (4)	5m19s
derby	1.0M	5	29 (10)	0.01s	28 (10)	16.48s	30 (11)	22.49s	29 (10)	24.07s
jigsaw	3.0M	12	4 (4)	0.41s	4 (4)	19.53s	6 (6)	11.69s	6 (6)	17.30s
sunflow	11M	17	84 (6)	39.66s	58 (6)	47.14s	130 (7)	50.24s	119 (7)	55.30s
cryptorsa	58M	9	11 (5)	3m4s	11 (5)	6m35s	TO	TO	35 (7)	9m42s
xalan	122M	7	31 (10)	0.15s	21 (7)	15m30s	TO	TO	37 (12)	10m44s
lufact	134M	5	21K (3)	7m26s	21K (3)	14m57s	TO	TO	21K (3)	10m38s
batik	157M	7	10 (2)	9m49s	10 (2)	22m56s	TO	TO	10 (2)	11m59s
lusearch	217M	8	232 (44)	12.63s	119 (27)	13m40s	232 (44)	27m9s	232 (44)	14m5s
tsp	307M	10	143 (6)	15m2s	140 (6)	29m10s	TO	TO	143 (6)	20m19s
luindex	397M	3	1 (1)	24m40s	2 (2)	31m6s	TO	TO	15 (15)	31m46s
sor	606M	5	0 (0)	38m38s	0 (0)	TO	TO	TO	0 (0)	44m36s
Totals	2.0B	-	29520 (157)	1h51m	29133 (134)	≥ 3h15m	1846 (131)	≥ 8h30m	30862 (178)	2h40m

Hence, sync-preservation is more adequate to capture not only racy program locations, but also racy memory locations. We note that, in principle, many different racy memory locations could correspond to the same static race (e.g., if memory is allocated dynamically). Note, however, that the additional reports of SyncP in Table 2 occur on benchmarks where it also makes more race reports in Table 1. Together, the two experimental tables give confidence that the new reported races are on entirely different variables.

Race Distances. We examine the capability of SyncP to detect races that are far apart in the input trace. Table 3 shows maximum race distance of races (e_1, e_2) in various benchmarks, including the ones that contains sync-preserving races that are missed by happens-before. In each case, the distance is counted as the number of events in the input trace between e_1 and e_2 , for every event e_2 reported as racy. We see a sharp contrast between SHB/WCP and SyncP, with the latter being

Table 2. Numbers of different memory locations that are detected as racy.

Benchmark	SHB	WCP	SyncP
ftpsrvr	49	49	50
jigsaw	4	4	5
xalan	7	6	9
cryptorsa	4	4	5
luindex	1	2	9
sunflow	14	10	17
linkedlist	927	927	932
Total	1006	1002	1027

Table 3. Maximum race distances.

Benchmark	SHB	WCP	SyncP
tsp	11K	11K	224M
batik	1.7M	1.7M	4.8M
cryptorsa	7.9M	7.9M	8.3M
jigsaw	428	428	121K
sunflow	10M	1.0M	10M
xalan	4K	4K	13K
ftpsrvr	11K	11K	11K
linkedlist	165K	165K	165K
luindex	783	783	6.9K
mergesort	57	57	1.4K
clean	355	47	1.2K
readswrites	13	13	696
wronglock	50	6	113

able to detect races that are far more distant in the input. This is in direct alignment with our theoretical observations already illustrated earlier in Section 1 (see Figure 1c). Indeed, as partial orders, SHB/WCP can only detect races between conflicting accesses that are successive in the input trace. On the other hand, sync-preserving races may be interleaved with arbitrarily many conflicting, non-racy accesses, and our complete algorithm SyncP is guaranteed to detect them. Overall, all our experimental observations suggest that sync-preservation is an elegant notion: it finely characterizes almost all races that are efficiently detectable, while it captures several races that are beyond the standard happens-before relation.

Table 4. Statistics of the core of the benchmark traces after the lightweight optimization is applied. $\bar{N}$, $\bar{\mathcal{T}}$, $\bar{\mathcal{A}}$, and $\bar{\mathcal{V}}$ denotes respectively the number of events, threads, lock-acquire events, and variables in the core trace.

Benchmark	$\bar{N}$	$\bar{\mathcal{T}}$	$\bar{\mathcal{A}}$	$\bar{\mathcal{V}}$	Benchmark	$\bar{N}$	$\bar{\mathcal{T}}$	$\bar{\mathcal{A}}$	$\bar{\mathcal{V}}$	Benchmark	$\bar{N}$	$\bar{\mathcal{T}}$	$\bar{\mathcal{A}}$	$\bar{\mathcal{V}}$
array	14	4	2	2	mergesort	170	6	49	1	jigsaw	3.0K	12	1.0K	51
critical	14	5	0	1	bubblesort	1.0K	13	119	25	sunflow	3.0K	17	585	20
account	18	5	0	1	lang	1.0K	8	0	100	cryptorsa	1.0M	9	156K	18
airtickets	27	5	0	1	readswrites	9.0K	6	1.0K	6	xalan	671K	7	183K	72
pingpong	38	7	0	2	raytracer	529	4	0	3	lufact	891K	5	0	4
twostage	86	13	20	2	bufwriter	10K	7	1.0K	6	batik	132	7	0	5
wronglock	125	23	20	1	ftpsrvr	17K	12	4.0K	135	lusearch	751K	8	53	77
bbuffer	13	3	0	1	molodyn	21K	4	0	2	tsp	15M	10	91	189
prodcons	248	9	34	3	linkedlist	910K	13	1.0K	932	luindex	15K	3	6.0K	9
clean	871	10	239	2	derby	75K	5	21K	190	sor	1.0M	5	633K	4

Complexity of SyncP and Running Time. Recall the complexity of SyncP established in Theorem 3.1. We have argued that the complexity is $\tilde{O}(N)$, i.e., $\mathcal{T}, \mathcal{V} = \tilde{O}(1)$, meaning that the number of threads and variables are much smaller than N . Here we justify this assumption experimentally, by presenting these numbers for the benchmark traces in Table 4. For each trace σ , we report the parameters of the core trace $\bar{\sigma}$ resulting from our lightweight optimization discussed earlier. We see that SyncP (and the other algorithms) is, in reality, executed on the core trace $\bar{\sigma}$ where the number of threads $\bar{\mathcal{T}}$ and variables $\bar{\mathcal{V}}$ is indeed considerably smaller than $\bar{N}$. Hence, our theoretical treatment of $\mathcal{T}, \mathcal{V} = \tilde{O}(1)$ is justified.

7 RELATED WORK

The happens-before (HB) partial order, computable in linear time [Fidge 1991; Mattern 1988], is the basis of many race detectors [Bond et al. 2010; Christiaens and Bosschere 2001; Flanagan and Freund 2009; Pozniansky and Schuster 2003; Schonberg 1989]. However, it can miss many predictable races. Recent work improves upon this with a small increase of computational resources [Genç et al. 2019; Kini et al. 2017; Pavlogiannis 2019; Roemer et al. 2018, 2020; Smaragdakis et al. 2012].

Another common approach to race prediction is via lockset-based methods. At a high level, a lockset is a set of locks that guards all accesses to a given memory location. Such techniques report races when they discover a write access to a location which is not consistently protected (i.e., whose lockset is empty). They were introduced in [Dinning and Schonberg 1991] and equipped in Eraser [Savage et al. 1997]. The lockset criterion is complete but unsound, and various works attempt to reduce false positives by enhancements such as random testing [Sen 2008] and static analysis [Choi et al. 2002; von Praun and Gross 2001]. Locksets have also been combined with happens-before techniques [Elmas et al. 2007; Yu et al. 2005].

Another direction to dynamic race prediction is symbolic techniques that typically rely on SAT/SMT encodings of the condition of a correct reordering, and dispatch such encodings to the respective solver [Huang et al. 2014; Liu et al. 2016; Said et al. 2011; Wang et al. 2009]. The encodings are typically sound and complete in theory, but the solution takes exponential time. In practice, windowing techniques are used to fragment the trace into chunks and analyze each chunk independently. This introduces incompleteness, as races between events of different chunks are naturally missed. Dynamic techniques have also been used for predicting other types of errors, such as deadlocks, atomicity violations and synchronization errors [Chen et al. 2008; Farzan and Madhusudan 2009; Farzan et al. 2009; Flanagan et al. 2008; Kalhauge and Palsberg 2018; Mathur and Viswanathan 2020; Sen et al. 2005; Sorrentino et al. 2010].

8 CONCLUSION

In this work, we have introduced the new notion of synchronization-preserving races. Conceptually, this is a completion of the principle behind happens-before races, namely that such races can be witnessed without reversing the order in which synchronization operations are observed. We have shown that sync-preservation strictly subsumes happens-before, and can detect races that are far apart in the input trace. We have developed an algorithm SyncP that is sound and complete for sync-preserving races, and has nearly linear time and space complexity. In addition, we have shown that relaxing our definition even slightly, i.e., by allowing a single synchronization reversal suffices to make the problem $W[1]$ -hard. Finally, we have performed an extensive experimental evaluation of SyncP. Our experiments show that sync-preservation is an elegant notion that characterizes almost all races that are efficiently detectable, while it captures several races that are beyond the standard happens-before. Given the demonstrated relevance of this new notion, we identify as important future work the development of more efficient race detectors for sync-preserving races, in a similar manner that happens-before race detectors have been refined over the years.

ACKNOWLEDGMENTS

We thank anonymous reviewers for their constructive feedback on an earlier draft of this manuscript. Umang Mathur is partially supported by a Google PhD Fellowship. Mahesh Viswanathan is partially supported by grants NSF SHF 1901069 and NSF CCF 2007428.

REFERENCES

- Joaquín Aguado, Michael Mendler, Marc Pouzet, Partha Roop, and Reinhard von Hanxleden. 2018. Deterministic Concurrency: A Clock-Synchronised Shared Memory Approach. In *Programming Languages and Systems*, Amal Ahmed (Ed.). Springer International Publishing, Cham, 86–113.
- Stephen M. Blackburn, Robin Garner, Chris Hoffmann, Asjad M. Khang, Kathryn S. McKinley, Rotem Bentzur, Amer Diwan, Daniel Feinberg, Daniel Frampton, Samuel Z. Guyer, Martin Hirzel, Antony Hosking, Maria Jump, Han Lee, J. Eliot B. Moss, Aashish Phansalkar, Darko Stefanović, Thomas VanDrunen, Daniel von Dincklage, and Ben Wiedermann. 2006. The DaCapo Benchmarks: Java Benchmarking Development and Analysis. In *Proceedings of the 21st Annual ACM SIGPLAN Conference on Object-oriented Programming Systems, Languages, and Applications* (Portland, Oregon, USA) (OOPSLA '06). ACM, New York, NY, USA, 169–190. <https://doi.org/10.1145/1167473.1167488>
- Robert L. Bocchino, Vikram S. Adve, Sarita V. Adve, and Marc Snir. 2009. Parallel Programming Must Be Deterministic by Default. In *Proceedings of the First USENIX Conference on Hot Topics in Parallelism* (Berkeley, California) (HotPar'09). USENIX Association, USA, 4.
- Hans-J. Boehm. 2011. How to Mismatch Programs with “Benign” Data Races. In *Proceedings of the 3rd USENIX Conference on Hot Topic in Parallelism* (Berkeley, CA) (HotPar'11). USENIX Association, USA, 3.
- Hans-J. Boehm. 2012. Position Paper: Nondeterminism is Unavoidable, but Data Races Are Pure Evil. In *Proceedings of the 2012 ACM Workshop on Relaxing Synchronization for Multicore and Manycore Scalability* (Tucson, Arizona, USA) (RACES '12). Association for Computing Machinery, New York, NY, USA, 9–14. <https://doi.org/10.1145/2414729.2414732>
- Michael D. Bond, Katherine E. Coons, and Kathryn S. McKinley. 2010. PACER: Proportional Detection of Data Races. In *Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation* (Toronto, Ontario, Canada) (PLDI '10). ACM, New York, NY, USA, 255–268. <https://doi.org/10.1145/1806596.1806626>
- Feng Chen, Traian Florin Șerbănuță, and Grigore Roșu. 2008. jPredictor: a predictive runtime analysis tool for Java. In *ICSE '08: Proceedings of the 30th International Conference on Software Engineering* (Leipzig, Germany). ACM, New York, NY, USA, 221–230.
- Jong-Deok Choi, Keunwoo Lee, Alexey Loginov, Robert O’Callahan, Vivek Sarkar, and Manu Sridharan. 2002. Efficient and Precise Datarace Detection for Multithreaded Object-oriented Programs. In *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation* (Berlin, Germany) (PLDI '02). ACM, New York, NY, USA, 258–269. <https://doi.org/10.1145/512529.512560>
- Mark Christiaens and Koenraad De Bosschere. 2001. TRaDe: Data Race Detection for Java. In *Proceedings of the International Conference on Computational Science-Part II (ICCS '01)*. Springer-Verlag, London, UK, UK, 761–770.
- Heming Cui, Rui Gu, Cheng Liu, Tianyu Chen, and Junfeng Yang. 2015. Paxos Made Transparent. In *Proceedings of the 25th Symposium on Operating Systems Principles* (Monterey, California) (SOSP '15). Association for Computing Machinery, New York, NY, USA, 105–120. <https://doi.org/10.1145/2815400.2815427>
- Anne Dinning and Edith Schonberg. 1991. Detecting Access Anomalies in Programs with Critical Sections. In *Proceedings of the 1991 ACM/ONR Workshop on Parallel and Distributed Debugging* (Santa Cruz, California, USA) (PADD '91). ACM, New York, NY, USA, 85–96. <https://doi.org/10.1145/122759.122767>
- Hyunsook Do, Sebastian G. Elbaum, and Gregg Rothermel. 2005. Supporting Controlled Experimentation with Testing Techniques: An Infrastructure and its Potential Impact. *Empirical Software Engineering: An International Journal* 10, 4 (2005), 405–435.
- Tayfun Elmas, Shaz Qadeer, and Serdar Tasiran. 2007. Goldilocks: A Race and Transaction-aware Java Runtime. In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation* (San Diego, California, USA) (PLDI '07). ACM, New York, NY, USA, 245–255. <https://doi.org/10.1145/1250734.1250762>
- Eitan Farchi, Yarden Nir, and Shmuel Ur. 2003. Concurrent Bug Patterns and How to Test Them. In *Proceedings of the 17th International Symposium on Parallel and Distributed Processing (IPDPS '03)*. IEEE Computer Society, Washington, DC, USA, 286.2–.
- Azadeh Farzan and P. Madhusudan. 2009. The Complexity of Predicting Atomicity Violations. In *Proceedings of the 15th International Conference on Tools and Algorithms for the Construction and Analysis of Systems: Held As Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, (York, UK) (TACAS '09)*. Springer-Verlag, Berlin, Heidelberg, 155–169. https://doi.org/10.1007/978-3-642-00768-2_14
- Azadeh Farzan, P. Madhusudan, and Francesco Sorrentino. 2009. Meta-analysis for Atomicity Violations Under Nested Locking. In *Proceedings of the 21st International Conference on Computer Aided Verification* (Grenoble, France) (CAV '09). Springer-Verlag, Berlin, Heidelberg, 248–262. https://doi.org/10.1007/978-3-642-02658-4_21
- Colin Fidge. 1991. Logical Time in Distributed Computing Systems. *Computer* 24, 8 (Aug. 1991), 28–33. <https://doi.org/10.1109/2.84874>
- Cormac Flanagan and Stephen N. Freund. 2009. FastTrack: Efficient and Precise Dynamic Race Detection. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Dublin, Ireland) (PLDI '09). ACM, New York, NY, USA, 121–133. <https://doi.org/10.1145/1542476.1542490>

- Cormac Flanagan, Stephen N. Freund, and Jaeheon Yi. 2008. Velodrome: A Sound and Complete Dynamic Atomicity Checker for Multithreaded Programs. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Tucson, AZ, USA) (PLDI '08). ACM, New York, NY, USA, 293–303. <https://doi.org/10.1145/1375581.1375618>
- Kaan Genç, Jake Roemer, Yufan Xu, and Michael D. Bond. 2019. Dependence-Aware, Unbounded Sound Predictive Race Detection. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 179 (Oct. 2019), 30 pages. <https://doi.org/10.1145/3360605>
- Kaan Genç, Yufan Xu, and Michael D. Bond. 2020. Personal Communication. (2020).
- Phillip B. Gibbons and Ephraim Korach. 1997. Testing Shared Memories. *SIAM J. Comput.* 26, 4 (Aug. 1997), 1208–1244. <https://doi.org/10.1137/S0097539794279614>
- Nikos Gorogiannis, Peter W. O'Hearn, and Ilya Sergey. 2019. A True Positives Theorem for a Static Race Detector. *Proc. ACM Program. Lang.* 3, POPL, Article 57 (Jan. 2019), 29 pages. <https://doi.org/10.1145/3290370>
- Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: A Correctness Condition for Concurrent Objects. *ACM Trans. Program. Lang. Syst.* 12, 3 (July 1990), 463–492. <https://doi.org/10.1145/78969.78972>
- Jeff Huang, Patrick O'Neil Meredith, and Grigore Rosu. 2014. Maximal Sound Predictive Race Detection with Control Flow Abstraction. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Edinburgh, United Kingdom) (PLDI '14). ACM, New York, NY, USA, 337–348. <https://doi.org/10.1145/2594291.2594315>
- Jeff Huang and Arun K. Rajagopalan. 2016. Precise and Maximal Race Detection from Incomplete Traces. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications* (Amsterdam, Netherlands) (OOPSLA 2016). ACM, New York, NY, USA, 462–476. <https://doi.org/10.1145/2983990.2984024>
- Christian Gram Kalhauge and Jens Palsberg. 2018. Sound Deadlock Prediction. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 146 (Oct. 2018), 29 pages. <https://doi.org/10.1145/3276516>
- Baris Kasikci, Cristian Zamfir, and George Candea. 2013. RaceMob: Crowdsourced Data Race Detection. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles* (Farmington, Pennsylvania) (SOSP '13). ACM, New York, NY, USA, 406–422.
- Dileep Kini, Umang Mathur, and Mahesh Viswanathan. 2017. Dynamic Race Prediction in Linear Time. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Barcelona, Spain) (PLDI 2017). ACM, New York, NY, USA, 157–170. <https://doi.org/10.1145/3062341.3062374>
- Leslie Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* 21, 7 (July 1978), 558–565.
- Peng Liu, Omer Tripp, and Xiangyu Zhang. 2016. IPA: Improving Predictive Analysis with Pointer Analysis. In *Proceedings of the 25th International Symposium on Software Testing and Analysis* (Saarbrücken, Germany) (ISSTA 2016). Association for Computing Machinery, New York, NY, USA, 59–69. <https://doi.org/10.1145/2931037.2931046>
- Shan Lu, Soyeon Park, Eunsoo Seo, and Yuan Yuan Zhou. 2008. Learning from Mistakes: A Comprehensive Study on Real World Concurrency Bug Characteristics. In *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems* (Seattle, WA, USA) (ASPLOS XIII). ACM, New York, NY, USA, 329–339. <https://doi.org/10.1145/1346281.1346323>
- Umang Mathur. 2020. RAPID. <https://github.com/umangm/rapid> Accessed: 2020-10-25.
- Umang Mathur, Dileep Kini, and Mahesh Viswanathan. 2018. What Happens-after the First Race? Enhancing the Predictive Power of Happens-before Based Dynamic Race Detection. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 145 (Oct. 2018), 29 pages. <https://doi.org/10.1145/3276515>
- Umang Mathur, Andreas Pavlogiannis, and Mahesh Viswanathan. 2020a. The Complexity of Dynamic Data Race Prediction. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science* (Saarbrücken, Germany) (LICS '20). Association for Computing Machinery, New York, NY, USA, 713–727. <https://doi.org/10.1145/3373718.3394783>
- Umang Mathur, Andreas Pavlogiannis, and Mahesh Viswanathan. 2020b. Optimal Prediction of Synchronization-Preserving Races. *CoRR* abs/2010.16385 (2020). arXiv:2010.16385 <https://arxiv.org/abs/2010.16385>
- Umang Mathur and Mahesh Viswanathan. 2020. Atomicity Checking in Linear Time Using Vector Clocks. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (ASPLOS '20). Association for Computing Machinery, New York, NY, USA, 183–199. <https://doi.org/10.1145/3373376.3378475>
- Friedemann Mattern. 1988. Virtual Time and Global States of Distributed Systems. In *Parallel and Distributed Algorithms*. North-Holland, 215–226.
- Satish Narayanasamy, Zhenghao Wang, Jordan Tigani, Andrew Edwards, and Brad Calder. 2007. Automatically Classifying Benign and Harmful Data Races Using Replay Analysis. In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation* (San Diego, California, USA) (PLDI '07). Association for Computing Machinery, New York, NY, USA, 22–31. <https://doi.org/10.1145/1250734.1250738>
- Andreas Pavlogiannis. 2019. Fast, Sound, and Effectively Complete Dynamic Race Prediction. *Proc. ACM Program. Lang.* 4, POPL, Article 17 (Dec. 2019), 29 pages. <https://doi.org/10.1145/3371085>

- Eli Pozniansky and Assaf Schuster. 2003. Efficient On-the-fly Data Race Detection in Multithreaded C++ Programs. In *Proceedings of the Ninth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (San Diego, California, USA) (PPoPP '03). ACM, New York, NY, USA, 179–190. <https://doi.org/10.1145/781498.781529>
- Jake Roemer and Michael D. Bond. 2019. Online Set-Based Dynamic Analysis for Sound Predictive Race Detection. *CoRR* abs/1907.08337 (2019). arXiv:1907.08337 <http://arxiv.org/abs/1907.08337>
- Jake Roemer, Kaan Genç, and Michael D. Bond. 2018. High-coverage, Unbounded Sound Predictive Race Detection. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Philadelphia, PA, USA) (PLDI 2018). ACM, New York, NY, USA, 374–389. <https://doi.org/10.1145/3192366.3192385>
- Jake Roemer, Kaan Genç, and Michael D. Bond. 2020. SmartTrack: Efficient Predictive Race Detection. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) (PLDI 2020). Association for Computing Machinery, New York, NY, USA, 747–762. <https://doi.org/10.1145/3385412.3385993>
- Grigore Rosu. 2018. *RV-Predict, Runtime Verification*. Accessed: 2018-04-01.
- Mahmoud Said, Chao Wang, Zijiang Yang, and Karem Sakallah. 2011. Generating Data Race Witnesses by an SMT-based Analysis. In *Proceedings of the Third International Conference on NASA Formal Methods* (Pasadena, CA) (NFM'11). Springer-Verlag, Berlin, Heidelberg, 313–327.
- Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro, and Thomas Anderson. 1997. Eraser: A Dynamic Data Race Detector for Multithreaded Programs. *ACM Trans. Comput. Syst.* 15, 4 (Nov. 1997), 391–411. <https://doi.org/10.1145/265924.265927>
- D. Schonberg. 1989. On-the-fly Detection of Access Anomalies. In *Proceedings of the ACM SIGPLAN 1989 Conference on Programming Language Design and Implementation* (Portland, Oregon, USA) (PLDI '89). ACM, New York, NY, USA, 285–297. <https://doi.org/10.1145/73141.74844>
- Koushik Sen. 2008. Race Directed Random Testing of Concurrent Programs. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Tucson, AZ, USA) (PLDI '08). ACM, New York, NY, USA, 11–21. <https://doi.org/10.1145/1375581.1375584>
- Koushik Sen, Grigore Roşu, and Gul Agha. 2005. Detecting Errors in Multithreaded Programs by Generalized Predictive Analysis of Executions. In *Proceedings of the 7th IFIP WG 6.1 International Conference on Formal Methods for Open Object-Based Distributed Systems* (Athens, Greece) (FMODS'05). Springer-Verlag, Berlin, Heidelberg, 211–226.
- Traian Florin Şerbănuţă, Feng Chen, and Grigore Roşu. 2012. Maximal causal models for sequentially consistent systems. In *International Conference on Runtime Verification*. Springer, 136–150.
- Ilya Sergey. 2019. *What Does It Mean for a Program Analysis to Be Sound?* <https://blog.sigplan.org/2019/08/07/what-does-it-mean-for-a-program-analysis-to-be-sound/>
- Yannis Smaragdakis, Jacob Evans, Caitlin Sadowski, Jaeheon Yi, and Cormac Flanagan. 2012. Sound Predictive Race Detection in Polynomial Time. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Philadelphia, PA, USA) (POPL '12). ACM, New York, NY, USA, 387–400. <https://doi.org/10.1145/2103656.2103702>
- L. A. Smith, J. M. Bull, and J. Obdrizalek. 2001. A Parallel Java Grande Benchmark Suite. In *SC '01: Proceedings of the 2001 ACM/IEEE Conference on Supercomputing*. 6–6. <https://doi.org/10.1145/582034.582042>
- Francesco Sorrentino, Azadeh Farzan, and P. Madhusudan. 2010. PENELOPE: Weaving Threads to Expose Atomicity Violations. In *Proceedings of the Eighteenth ACM SIGSOFT International Symposium on Foundations of Software Engineering* (Santa Fe, New Mexico, USA) (FSE '10). ACM, New York, NY, USA, 37–46. <https://doi.org/10.1145/1882291.1882300>
- Christoph von Praun and Thomas R. Gross. 2001. Object Race Detection. In *Proceedings of the 16th ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications* (Tampa Bay, FL, USA) (OOPSLA '01). ACM, New York, NY, USA, 70–82. <https://doi.org/10.1145/504282.504288>
- Chao Wang, Sudipta Kundu, Malay Ganai, and Aarti Gupta. 2009. Symbolic Predictive Analysis for Concurrent Programs. In *Proceedings of the 2nd World Congress on Formal Methods* (Eindhoven, The Netherlands) (FM '09). Springer-Verlag, Berlin, Heidelberg, 256–272.
- Misun Yu, Joon-Sang Lee, and Doo-Hwan Bae. 2018. AdaptiveLock: Efficient Hybrid Data Race Detection Based on Real-World Locking Patterns. *International Journal of Parallel Programming* (04 Jun 2018). <https://doi.org/10.1007/s10766-018-0579-5>
- Yuan Yu, Tom Rodeheffer, and Wei Chen. 2005. RaceTrack: Efficient Detection of Data Race Conditions via Adaptive Tracking. In *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles* (Brighton, United Kingdom) (SOSP '05). Association for Computing Machinery, New York, NY, USA, 221–234. <https://doi.org/10.1145/1095810.1095832>
- Qi Zhao, Zhengyi Qiu, and Guoliang Jin. 2019. Semantics-Aware Scheduling Policies for Synchronization Determinism. In *Proceedings of the 24th Symposium on Principles and Practice of Parallel Programming* (Washington, District of Columbia) (PPoPP '19). Association for Computing Machinery, New York, NY, USA, 242–256. <https://doi.org/10.1145/3293883.3295731>
- M. Zhivich and R. K. Cunningham. 2009. The Real Cost of Software Errors. *IEEE Security and Privacy* 7, 2 (March 2009), 87–90. <https://doi.org/10.1109/MSP.2009.56>