

1 Dynamic Data-Race Detection through the 2 Fine-Grained Lens

3 **Rucha Kulkarni** @

4 University of Illinois at Urbana-Champaign, USA

5 **Umang Mathur** @

6 University of Illinois at Urbana-Champaign, USA

7 **Andreas Pavlogiannis** @

8 Aarhus University, Denmark

9 — Abstract —

10 Data races are among the most common bugs in concurrency. The standard approach to data-race
11 detection is via dynamic analyses, which work over executions of concurrent programs, instead of
12 the program source code. The rich literature on the topic has created various notions of dynamic
13 data races, which are known to be detected efficiently when certain parameters (e.g., number of
14 threads) are small. However, the *fine-grained* complexity of all these notions of races has remained
15 elusive, making it impossible to characterize their trade-offs between precision and efficiency.

16 In this work we establish several fine-grained separations between many popular notions of dynamic
17 data races. The input is an execution trace σ with \mathcal{N} events, \mathcal{T} threads and \mathcal{L} locks. Our main results
18 are as follows. First, we show that *happens-before* (HB) races can be detected in $O(\mathcal{N} \cdot \min(\mathcal{T}, \mathcal{L}))$
19 time, improving over the standard $O(\mathcal{N} \cdot \mathcal{T})$ bound when $\mathcal{L} = o(\mathcal{T})$. Moreover, we show that even
20 reporting an HB race that involves a read access is hard for 2-orthogonal vectors (2-OV). This is the
21 first rigorous proof of the conjectured quadratic lower-bound in detecting HB races. Second, we
22 show that the recently introduced *synchronization-preserving races* are hard to detect for OV-3 and
23 thus have a cubic lower bound, when $\mathcal{T} = \Omega(\mathcal{N})$. This establishes a complexity separation from
24 HB races which are known to be less expressive. Third, we show that *lock-cover races* are hard for
25 2-OV, and thus have a quadratic lower-bound, even when $\mathcal{T} = 2$ and $\mathcal{L} = \omega(\log \mathcal{N})$. The similar
26 notion of *lock-set races* is known to be detectable in $O(\mathcal{N} \cdot \mathcal{L})$ time, and thus we achieve a complexity
27 separation between the two. Moreover, we show that lock-set races become hitting-set (HS)-hard
28 when $\mathcal{L} = \Theta(\mathcal{N})$, and thus also have a quadratic lower bound, when the input is sufficiently complex.
29 To our knowledge, this is the first work that characterizes the complexity of well-established dynamic
30 race-detection techniques, allowing for a rigorous comparison between them.

31 **2012 ACM Subject Classification** Software and its engineering → Software testing and debugging;
32 Theory of computation → Parameterized complexity and exact algorithms

33 **Keywords and phrases** dynamic analyses, data races, fine-grained complexity

34 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

35 **1** Introduction

36 Concurrent programs that communicate over shared memory are prone to *data races*. Two
37 events are *conflicting* if they access the same memory location and one (at least) modifies that
38 location. Data races occur when conflicting access happen concurrently between different
39 threads, and form one of the most common bugs in concurrency. In particular, data races are
40 often symptomatic of bugs in software like data corruption [5, 20, 26], and they have been
41 deemed *pure evil* [6] due to the problems they have caused in the past [43]. Moreover, many
42 compiler optimizations are unsound in the presence of data races [36, 40], while data-race



© Rucha Kulkarni and Umang Mathur and Andreas Pavlogiannis;
licensed under Creative Commons License CC-BY 4.0

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:28



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

43 freeness is often a requirement for assigning well-defined semantics to programs [7].

44 The importance of data races in concurrency has led to a multitude of techniques for detecting
 45 them efficiently [4, 39]. By far the most standard approach is via *dynamic analyses*. Instead
 46 of analyzing the full program, dynamic analyzers try to *predict* the existence of data races
 47 by observing and analyzing concurrent executions [37, 21, 28]. As full dynamic data race
 48 prediction is NP-hard in general [24], researchers have developed several approximate notions
 49 of dynamic races, accompanied by efficient techniques for detecting each notion.

50 *Happens-before races*. The most common technique for detecting data races dynamically
 51 is based on Lamport’s *happens-before (HB)* partial order [22]. Two conflicting events form
 52 an HB race if they are unordered by HB, as the lack of ordering between them indicates
 53 the fact that they may execute concurrently, thereby forming a data race. The standard
 54 approach to HB race detection is via the use of vector clocks [19], and has seen wide success
 55 in commercial race detectors [35]. As vector clock computation is known to require $\Theta(\mathcal{N} \cdot \mathcal{T})$
 56 time on traces of \mathcal{N} events and \mathcal{T} threads [10], HB race detection is often assumed to suffer
 57 the same bound, and has thus been a subject of further practical optimizations [29, 16].

58 *Synchronization preserving races*. HB races were recently generalized to sync(hronization)-
 59 preserving races [25]. Intuitively, two conflicting events are in a sync-preserving race if the
 60 observed trace can be soundly reordered to a witness trace in which the two events are
 61 concurrent, but without reordering synchronization events (e.g., locking events). Similar to
 62 HB races, sync-preserving races can be detected in linear time when the number of threads
 63 is constant. However, the dependence on the number of threads is cubic for sync-preserving
 64 races, as opposed to the linear dependence for HB races. On the other hand, sync-preserving
 65 races are known to offer better precision in program analysis.

66 *Races based on the locking discipline*. The locking discipline dictates that threads that access
 67 a common memory location must do so inside *critical sections*, using a common lock, when
 68 performing the access [39]. Although this discipline is typically not enforced, it is considered
 69 good practice, and hence instances that violate this principle are often considered indicators
 70 of erroneous behavior. For this reason, there have been two popular notions of data races
 71 based on the locking discipline, namely *lock-cover races* [14] and *lock-set races* [33]. Both
 72 notions are detectable in linear time when the number of locks is constant, however, lock-set
 73 race detection is typically faster in practice, which also comes at the cost of being less precise.

74 Observe that, although techniques for all aforementioned notions of races are generally
 75 thought to operate in linear time, they only do so assuming certain parameters, such as the
 76 number of threads, are constant. However, as these techniques are deployed in runtime, often
 77 with extremely long execution traces, they have to be as efficient as absolutely possible, often
 78 in scenarios when these parameters are very large. When a data-race detection technique is
 79 too slow for a given application, the developers face a dilemma: do they look for a faster
 80 algorithm, or for a simpler abstraction (i.e., a different notion of dynamic races)? For these
 81 reasons, it is important to understand the *fine-grained* complexity of the problem at hand
 82 with respect to such parameters. Fine-grained lower bounds can rule out the possibility of
 83 faster algorithms, and thus help the developers focus on new abstractions that are more
 84 tractable for the given application. Motivated by such questions, in this work we settle the
 85 fine-grained complexity of dynamically detecting several popular notions of data races.

86 1.1 Our Contributions

87 Here we give a full account of the main results of this work, while we refer to later sections
 88 for precise definitions and proofs. We also refer to Appendix A for relevant notions in
 89 fine-grained complexity and popular hypotheses. The input is always a concurrent trace σ of
 90 length \mathcal{N} , consisting of \mathcal{T} threads, \mathcal{L} locks, and \mathcal{V} variables.

91 **Happens-before races.** We first study the fine-grained complexity of HB races, as they
 92 form the most popular class of dynamic data races. The task of most techniques is to report
 93 all events in σ that participate in an HB race, which is known to take $O(\mathcal{N} \cdot \mathcal{T})$ time [19]. Note
 94 that the bound is quadratic when $\mathcal{T} = \Theta(\mathcal{N})$, and multiple heuristics have been developed
 95 to address it in practice (see e.g., [16]). Our first result shows that polynomial improvements
 96 below this quadratic bound are unlikely.

97 **► Theorem 1.** *For any $\epsilon > 0$, there is no algorithm that detects even a single HB race that
 98 involves a read in time $O(\mathcal{N}^{2-\epsilon})$, unless the OV hypothesis fails.*

99 Orthogonal vectors (OV) is a well-studied problem with a long-standing quadratic worst-case
 100 upper bound. The associated hypothesis states that there is no sub-quadratic algorithm
 101 for the problem [42]. It is also known that the strong exponential time hypothesis (SETH)
 102 implies the Orthogonal Vectors hypothesis [41]. Thus, under the OV hypothesis, Theorem 1
 103 establishes a quadratic lower bound for HB race detection.

104 Note that the hardness of Theorem 1 arises out of the requirement to detect HB races that
 105 involve a read. A natural follow-up question is whether detecting if the input contains *any*
 106 HB race (i.e., not necessarily involving a read) has a similar lower bound based on SETH.
 107 Our next theorem shows that under the non-deterministic SETH (NSETH) [9], there is no
 108 fine-grained reduction from SETH that proves any lower bound for this problem above $\mathcal{N}^{3/2}$.

109 **► Theorem 2.** *For any $\epsilon > 0$, there is no $(2^{\mathcal{N}}, \mathcal{N}^{3/2+\epsilon})$ -fine-grained reduction from SAT to
 110 the problem of detecting any HB race with bound, unless NSETH fails.*

111 Given the impossibility of Theorem 2, it would be desirable to at least show a super-linear
 112 lower bound for detecting any HB data race. To tackle this question, we show that detecting
 113 any HB race is hard for the general problem of model checking first-order formulas quantified
 114 by $\forall\exists\exists$ on structures of size n with m relational tuples (denoted $\text{FO}(\forall\exists\exists)$).

115 **► Theorem 3.** *For any $\epsilon > 0$, if there is an algorithm for detecting any HB race in time
 116 $O(\mathcal{N}^{1+\epsilon})$, then there is an algorithm for $\text{FO}(\forall\exists\exists)$ formulas in time $O(m^{1+\epsilon})$.*

117 It is known that $\text{FO}(\forall\exists\exists)$ can be solved in $O(m^{3/2})$ time [17], which yields a bound $O(n^3)$ for
 118 dense structures (i.e., when $m = \Theta(n^2)$). Theorem 3 implies that if $m^{3/2}$ is the best possible
 119 bound for $\text{FO}(\forall\exists\exists)$, then detecting any HB race cannot take $O(\mathcal{N}^{1+\epsilon})$ time for any $\epsilon < 1/2$.
 120 Although improvements for $\text{FO}(\forall\exists\exists)$ over the current $O(m^{3/2})$ bound might be possible,
 121 we find that a truly linear bound $O(m)$ would require major breakthroughs¹. Under this
 122 hypothesis, Theorem 3 implies a super-linear bound for HB races.

123 Finally, we give an improved upper bound for this problem when $\mathcal{L} = o(\mathcal{T})$.

124 **► Theorem 4.** *Deciding whether σ has an HB race can be done in time $O(\mathcal{N} \cdot \min(\mathcal{T}, \mathcal{L}))$.*

125 In fact, similar to existing techniques [16], the algorithm behind Theorem 4 detects *all*
 126 variables that participate in an HB race (instead of just reporting σ as racy).

127 **Synchronization-preserving races.** Next, we turn our attention to the recently introduced

¹ Even the well-studied problem of testing triangle freeness, which is a special case of the similarly flavored $\text{FO}(\exists\exists\exists)$, has the super-linear bound $O(n^\omega)$.

128 sync-preserving races [24]. It is known that detecting sync-preserving races takes $O(\mathcal{N} \cdot \mathcal{V} \cdot \mathcal{T}^3)$
 129 time. As sync-preserving races are known to be more expressive than HB races, the natural
 130 question is whether sync-preserving races can be detected more efficiently, e.g., by an
 131 algorithm that achieves a bound similar to Theorem 4 for HB races. Our next theorem
 132 answers this question in negative.

133 ► **Theorem 5.** *For any $\epsilon > 0$, there is no algorithm that detects even a single sync-preserving*
 134 *race in time $O(\mathcal{N}^{3-\epsilon})$, unless the 3-OV hypothesis fails. Moreover, the statement holds even*
 135 *for traces over a single variable.*

136 As HB races take at most quadratic time, Theorem 5 shows that the increased expressiveness
 137 of sync-preserving races incurs a complexity overhead that is unavoidable in general.

138 **Races based on the locking discipline.** We now turn our attention to data races based on
 139 the locking discipline, namely *lock-cover races* and *lock-set races*. It is known that lock-cover
 140 races are more expressive than lock-set races. On the other hand, existing algorithms run
 141 in $O(\mathcal{N}^2 \cdot \mathcal{L})$ time for lock-cover races and in $O(\mathcal{N} \cdot \mathcal{L})$ time for lock-set races, and thus
 142 hint that the former are computationally harder to detect. Our first theorem makes this
 143 separation formal, by showing that even with just two threads, having slightly more than
 144 logarithmically many locks implies a quadratic hardness for lock-cover races.

145 ► **Theorem 6.** *For any $\epsilon > 0$, any $\mathcal{T} \geq 2$ and any $\mathcal{L} = \omega(\log \mathcal{N})$, there is no algorithm that*
 146 *detects even a single lock-cover race in time $O(\mathcal{N}^{2-\epsilon})$, unless the OV hypothesis fails.*

147 Observe that the $O(\mathcal{N} \cdot \mathcal{L})$ bound for lock-set races also becomes quadratic, when the number
 148 of locks is unbounded (i.e., $\mathcal{L} = \Theta(\mathcal{N})$). Is there a SETH-based quadratic lower bound similar
 149 to Theorem 6 for this case? Our next theorem rules out this possibility, again under NSETH.

150 ► **Theorem 7.** *For any $\epsilon > 0$, there is no $(2^{\mathcal{N}}, \mathcal{N}^{1+\epsilon})$ -fine-grained reduction from SAT to*
 151 *the problem of detecting any lock-set race, unless NSETH fails.*

152 Hence, even though we desire a quadratic lower bound, Theorem 7 rules out any super-linear
 153 lower-bound based on SETH. Alas, our next theorem shows that a quadratic lower bound for
 154 lock-set races does exist, based on the hardness of the hitting set (HS) problem.

155 ► **Theorem 8.** *For any $\epsilon > 0$ and any $\mathcal{T} = \omega(\log n)$, there is no algorithm that detects even*
 156 *a single lock-cover race in time $O(\mathcal{N}^{2-\epsilon})$, unless the HS hypothesis fails.*

157 Hitting set is a problem similar to OV, but has different quantifier structure. Just like the
 158 OV hypothesis, the HS hypothesis states that there is no sub-quadratic algorithm for the
 159 problem [3]. Although HS implies OV, the opposite is not known, and thus Theorem 8
 160 does not contradict Theorem 7. In conclusion, we have that both lock-cover and lock-set
 161 races have (conditional) quadratic lower bounds, though the latter is based on a stronger
 162 hypothesis (HS), and requires more threads and locks for hardness to arise.

163 Finally, on our way to Theorem 7, we obtain the following theorem.

164 ► **Theorem 9.** *Deciding whether a trace σ has a lock-set race on a variable x can be*
 165 *performed in $O(\mathcal{N})$ time. Thus, deciding whether σ has a lock-set race can be performed in*
 166 *$O(\mathcal{N} \cdot \min(\mathcal{L}, \mathcal{V}))$ time.*

167 Hence, Theorem 9 strengthens the $O(\mathcal{N} \cdot \mathcal{L})$ upper bound for lock-set races when $\mathcal{V} = o(\mathcal{L})$.

168 1.2 Related Work

169 **Dynamic data-race detection.** There exists a rich literature in dynamic techniques for
 170 data race detection. Methods based on vector clocks (DJIT algorithm [19]) using Lamport's
 171 Happens Before (HB) [22] and the lock-set principle in Eraser [33] were the first ones to
 172 popularize dynamic analysis for detecting data races. Later work attempted to increase the

173 performance of these notions using optimizations as in [29] and FASTTRACK [16], altogether
 174 different algorithms (e.g., the GoldiLocks algorithm [15]), and hybrid techniques [27]. HB
 175 and lock-set based race detection are respectively sound (but incomplete) and complete (but
 176 unsound) variants of the more general problem of data-race *prediction* [34]. While earlier
 177 work on data race prediction focused on explicit [34] or symbolic [31, 32] enumeration, recent
 178 efforts have focused on scalability [37, 23, 21, 28, 30, 38]. The more recent notion [25] of
 179 sync-preserving races generalizes the notion of HB. As the complexity of race prediction is
 180 prohibitive (NP-hard in general [24]), this work characterizes the fine-grained complexity of
 181 popular, more relaxed notions of dynamic races that take polynomial time.

182 **Fine-grained complexity.** Traditional complexity theory usually shows a problem is
 183 intractable by proving it NP-hard, and tractable by showing it is in P. For algorithms with
 184 large input sizes, this distinction may be too coarse. It becomes important to understand,
 185 even for problems in P, whether algorithms with smaller degree polynomials than the known
 186 are possible, or if there are fine-grained lower bounds making this unlikely. Fine-grained
 187 complexity involves proving such lower bounds, by showing relationships between problems
 188 in P, with an emphasis on the degree of the complexity polynomial, and is nowadays a field
 189 of very active study. We refer to [8] for an introductory, and to [42] for a more extensive
 190 exposition on the topic. Fine-grained arguments have also been instrumental in characterizing
 191 the complexity of various problems in concurrency, such as bounded context-switching [11],
 192 safety verification [12], data-race prediction [24] and consistency checking [13].

193 2 Preliminaries

194 2.1 Concurrent Program Executions and Data Races

195 **Traces and Events.** We consider execution traces (or simply *traces*) generated by concurrent
 196 programs, under the sequential consistency memory model. Under this memory model, a
 197 trace σ is a sequence of events. Each event e is labeled with a tuple $\text{lab}(e) = \langle t, op \rangle$, where t
 198 is the (unique) identifier of the thread that performs the event e , and op is the operation
 199 performed in e . We will often abuse notation and write $e = \langle t, op \rangle$ instead of $\text{lab}(e) = \langle t, op \rangle$.
 200 For the purpose of this presentation, an operation can be one of (a) read ($\text{r}(x)$) from, or
 201 write ($\text{w}(x)$) to, a shared memory variable x , (b) $\text{acq}(\ell)$ or $\text{rel}(\ell)$ of a lock ℓ .

202 For an event $e = \langle t, op \rangle$, we use $\text{tid}(e)$ and $\text{op}(e)$ to denote respectively the thread identifier t
 203 and the operation op . For a trace σ , we use Events_σ to denote the set of events that appear
 204 in σ . Similarly, we will use Threads_σ , Locks_σ and Vars_σ to denote respectively the set of
 205 threads, locks and shared variables that appear in trace σ . We denote by $\mathcal{N} = |\text{Events}_\sigma|$,
 206 $\mathcal{T} = |\text{Threads}_\sigma|$, $\mathcal{L} = |\text{Locks}_\sigma|$, and $\mathcal{V} = |\text{Vars}_\sigma|$. The set of read events and write events
 207 on variable $x \in \text{Vars}_\sigma$ will be denoted by $\text{Reads}_\sigma(x)$ and $\text{Writes}_\sigma(x)$, and further we let
 208 $\text{Accesses}_\sigma(x) = \text{Reads}_\sigma(x) \cup \text{Writes}_\sigma(x)$. Similarly, we let $\text{Acquires}_\sigma(\ell)$ and $\text{Releases}_\sigma(\ell)$
 209 denote the set of lock-acquire and lock-release events, respectively, of σ on lock ℓ . The *trace*
 210 *order* of σ , denoted \leq_{tr}^σ , is the total order on Events_σ induced by the sequence σ . Finally,
 211 the *thread-order* of σ , denoted \leq_{TO}^σ is the smallest partial order on Events_σ such that for any
 212 two events $e_1, e_2 \in \text{Events}_\sigma$, if $e_1 \leq_{\text{tr}}^\sigma e_2$ and $\text{tid}(e_1) = \text{tid}(e_2)$, then $e_1 \leq_{\text{TO}}^\sigma e_2$.

213 Traces are assumed to be well-formed in that critical sections on the same lock do not
 214 overlap. For a lock $\ell \in \text{Locks}_\sigma$, let $\sigma|_\ell$ be the projection of the trace σ on the set of events
 215 $\{e \mid \text{op}(e) \in \{\text{acq}(\ell), \text{rel}(\ell)\}\}$. Also, let t_1, \dots, t_k be the thread identifiers in Threads_σ . Well-
 216 formedness then entails that for each lock ℓ , the projection $\sigma|_\ell$ is a prefix of some string
 217 in the language of the grammar with production rules $S \rightarrow \varepsilon \mid S \cdot S_{t_1} \mid S \cdot S_{t_2} \mid \dots \mid S \cdot S_{t_k}$ and

218 $S_{t_i} \rightarrow \langle t_i, \text{acq}(\ell) \rangle \cdot \langle t_i, \text{rel}(\ell) \rangle$ and start symbol S . Thus, every release event e has a unique
 219 matching acquire event, which we denote by $\text{match}_\sigma(e)$. Likewise for an acquire event e ,
 220 $\text{match}_\sigma(e)$ denotes the unique matching release event if one exists. For an acquire event e , the
 221 critical section of e is the set of events $\text{CS}_\sigma(e) = \{f \mid e \leq_{\text{TO}}^\sigma f \leq_{\text{TO}}^\sigma \text{match}_\sigma(e)\}$ if $\text{match}_\sigma(e)$
 222 exists, and $\text{CS}_\sigma(e) = \{f \mid e \leq_{\text{TO}}^\sigma f\}$ otherwise.

223 **Data Races.** Two events $e_1, e_2 \in \text{Events}_\sigma$ are said to be *conflicting* if they are performed by
 224 different threads, they are access events touching the same memory location, and at least one
 225 of them is a write access. Formally, we have (i) $\text{tid}(e_1) \neq \text{tid}(e_2)$, (ii) $e_1, e_2 \in \text{Accesses}_\sigma(x)$
 226 for some $x \in \text{Vars}_\sigma$, and (iii) $\{e_1, e_2\} \cap \text{Writes}_\sigma(x) \neq \emptyset$. An event $e \in \text{Events}_\sigma$ is said to be
 227 *enabled* in a prefix ρ of σ , if for every event $e' \neq e$ with $e' \leq_{\text{TO}}^\sigma e$, we have $e' \in \text{Events}_\rho$. A
 228 data race in σ is a pair of conflicting events (e_1, e_2) such that there is a prefix ρ in which
 229 both e_1 and e_2 are simultaneously enabled. Depending on the type of access of e_1 and e_2 ,
 230 we often distinguish between write-write races and write-read races.

231 2.2 Notions of Dynamic Data Races

232 As the problem of determining whether a concurrent program has an execution with a data
 233 race is undecidable, dynamic techniques observe program traces and report whether certain
 234 events indicate the presence of a race. Depending on the technique, such reports can be
 235 sound (i.e., they guarantee the presence of a race in the program), Here we describe in detail
 236 some popular approaches to dynamic race detection that are the subject of this work.

237 **Happens-Before Races.** Given a trace σ , the *happens before* order \leq_{HB}^σ is the smallest
 238 partial order on Events_σ such that (a) $\leq_{\text{TO}}^\sigma \subseteq \leq_{\text{HB}}^\sigma$, and (b) for any lock $\ell \in \text{Locks}_\sigma$ and for
 239 events $e \in \text{Acquires}_\sigma(\ell)$ and $f \in \text{Releases}_\sigma(\ell)$, if $e \leq_{\text{tr}}^\sigma f$ then $e \leq_{\text{HB}}^\sigma f$. A pair of conflicting
 240 events (e_1, e_2) is an **HB-race** in σ if they are unordered by **HB**, i.e., $e_1 \not\leq_{\text{HB}}^\sigma e_2$ and $e_2 \not\leq_{\text{HB}}^\sigma e_1$.
 241 The associated decision question is, *given a trace σ , determine whether σ has an **HB race**.*
 242 Typically **HB** race detectors are tasked to report all events that form **HB** race with an earlier
 243 event in the trace [35, 2, 1]). That is, they solve the following function problem: *given a trace*
 244 *σ , determine all events $e_2 \in \text{Events}_\sigma$ for which there exists an event $e_1 \in \text{Events}_\sigma$ such that*
 245 *$e_1 \leq_{\text{tr}}^\sigma e_2$, and (e_1, e_2) is an **HB** race of σ .* The standard algorithm for solving both versions
 246 of the problem is a vector-clock algorithm that runs in $O(\mathcal{N} \cdot \mathcal{T})$ time [19].

247 **Synchronization Preserving Races.** Next, we present the notion of *sync(hronization)-*
 248 *preserving races* [24]. For a trace σ and a read event e , we use $\text{lw}_\sigma(e)$ to denote the write event
 249 observed by e . That is, $e' = \text{lw}_\sigma(e)$ is the last (according to the trace order \leq_{tr}^σ) write event e'
 250 of σ such that e and e' access the same variable and $e' \leq_{\text{tr}}^\sigma e$; if no such e' exists, then we write
 251 $\text{lw}_\sigma(e) = \perp$. A trace ρ is said to be a correct reordering of trace σ , if (a) $\text{Events}_\rho \subseteq \text{Events}_\sigma$
 252 (b) Events_ρ is downward closed with respect to \leq_{TO}^σ , and further $\leq_{\text{TO}}^\rho \subseteq \leq_{\text{TO}}^\sigma$, and (c) for every
 253 read event $e \in \text{Events}_\rho$, $\text{lw}_\rho(e) = \text{lw}_\sigma(e)$. We say that ρ is *sync-preserving* with respect to σ
 254 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
 255 $e_1 \leq_{\text{tr}}^\sigma e_2$. That is, the order of two critical sections on the same lock is the same in σ and ρ .

256 A pair of conflicting events (e_1, e_2) is a *sync-preserving race* in σ if σ has a sync-preserving
 257 correct reordering ρ such that (e_1, e_2) is a data race of ρ . The associated decision question
 258 is, *given a trace σ , determine whether σ has a sync-preserving race.* As with **HB** races, we
 259 are typically interested in reporting all events $e_2 \in \text{Events}_\sigma$ for which there exists an event
 260 $e_1 \in \text{Events}_\sigma$ such that $e_1 \leq_{\text{tr}}^\sigma e_2$, and (e_1, e_2) is an **HB** race of σ . It is known one can report
 261 all such events e_2 in time $O(\mathcal{N} \cdot \mathcal{V} \cdot \mathcal{T}^3)$.

262 **Lock-Cover and Lock-Set Races.** Lock-cover and lock-set races indicate violations

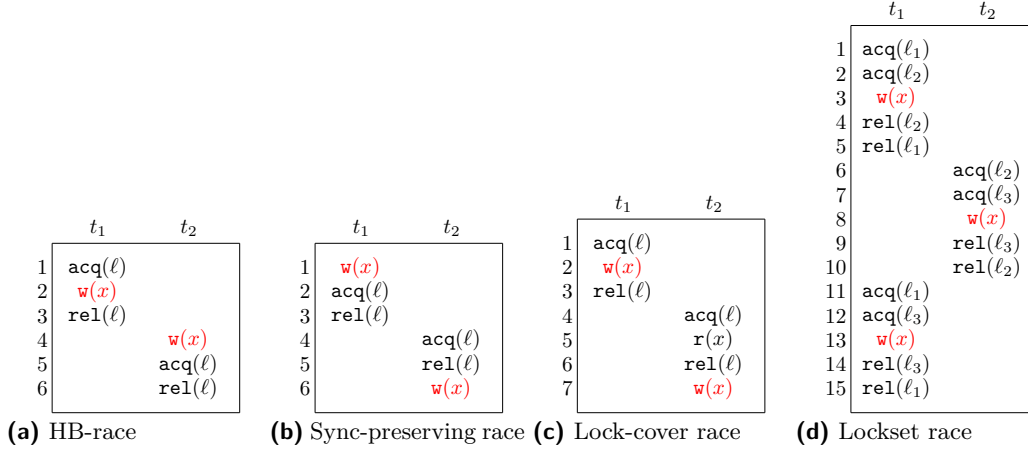


Figure 1 Types of data races.

of the *locking discipline*. For an event e in a trace σ , let $\text{locksHeld}_\sigma(e) = \{\ell \mid \exists f \in \text{Acquires}_\sigma(\ell), \text{ such that } e \in \text{CS}_\sigma(f)\}$, i.e., $\text{locksHeld}_\sigma(e)$ is the set of locks held by thread $\text{tid}(e)$ when e is executed. A pair of conflicting events might indicate a data race if $\text{locksHeld}_\sigma(e_1) \cap \text{locksHeld}_\sigma(e_2) = \emptyset$. Although this condition does not guarantee the presence of a race, it constitutes a violation of the locking discipline and can be further investigated.

A pair of conflicting events (e_1, e_2) is a *lock-cover race* if $\text{locksHeld}_\sigma(e_1) \cap \text{locksHeld}_\sigma(e_2) = \emptyset$. The decision question is, *given a trace σ , determine if σ has a lock-cover race*. The problem is solvable in $O(\mathcal{N}^2 \cdot \mathcal{L})$ time, by checking the above condition over all conflicting event pairs.

As the algorithm for lock-cover races takes quadratic time, developers often look for less expensive indications of violations of locking discipline, called lock-set races (as proposed by ERASER race detector [33]). A trace σ has a *lock-set race* on variable $x \in \text{Vars}_\sigma$ if

- (a) there exists a pair of conflicting events $(e_1, e_2) \in \text{Writes}_\sigma(x) \times \text{Accesses}_\sigma(x)$, and
- (b) $\bigcap_{e \in \text{Accesses}_\sigma(x)} \text{locksHeld}_\sigma(e) = \emptyset$.

The associated decision question is, *given a trace σ , determine if σ has a lock-set race*. Note that a lock-cover race implies a lock-set race, but not the other way around. On the other hand, determining whether σ has a lock-set race is easily performed in $O(\mathcal{N} \cdot \mathcal{L})$ time.

Example. We illustrate the different notions of races in Figure 1. We use e_i to denote the i^{th} event of the trace in consideration. First consider the trace σ_a in Figure 1a. The events e_2 and e_4 are conflicting and unordered by $\leq_{\text{HB}}^{\sigma_a}$, thus (e_2, e_4) is an HB-race. Second, in trace σ_b of Figure 1b, the pair (e_1, e_6) is not an HB-race as $e_1 \leq_{\text{HB}}^{\sigma_b} e_6$. But this is a sync-preserving race witnessed by the correct reordering e_4, e_5 , as both e_1 and e_6 are enabled. Third, in trace σ_c of Figure 1c, the pair (e_2, e_7) is neither a sync-preserving race nor an HB race, but is a lock-cover race as $\text{locksHeld}_{\sigma_c}(e_2) \cap \text{locksHeld}_{\sigma_c}(e_7) = \emptyset$. Finally, the trace σ_d in Figure 1d has no HB, sync-preserving or lock-cover race, as all $w(x)$ are protected by a common lock. But there is a lock-set race on x as there is no single lock that protects all $w(x)$.

3 Happens-Before Races

In this section we prove the results for detecting HB races, i.e., Theorem 1 to Theorem 4.

290 **3.1 Algorithm for HB Races**

291 In this section, we outline our $O(\mathcal{N} \cdot \mathcal{L})$ -time algorithm for checking if a trace σ has an HB-race,
 292 thereby proving Theorem 4. As with the standard vector clock algorithm [19], our algorithm
 293 is based on computing timestamps for each event. However, unlike the standard algorithm
 294 that assigns thread-indexed timestamps, we use *lock-indexed* timestamps, or *lockstamps*,
 295 which we formalize next. We fix the input trace σ in the rest of the discussion.

296 **Lockstamps** A lockstamp is a mapping from locks to natural numbers (including infinity)
 297 $L : \text{Locks}_\sigma \rightarrow \mathbb{N} \cup \{\infty\}$. Given lockstamps L, L_1, L_2 and lock ℓ , we use the notation
 298 (i) $L[\ell \mapsto c]$ to denote the the lockstamp $\lambda m \cdot$ if $m = \ell$ then c else $L(m)$, (ii) $L_1 \sqcup L_2$ to
 299 denote the pointwise maximum, i.e., $(L_1 \sqcup L_2)(\ell) = \max(L_1(\ell), L_2(\ell))$ for every ℓ , (iii) $L_1 \sqcap L_2$
 300 to denote the pointwise minimum, and (iv) $L_1 \sqsubseteq L_2$ to denote the predicate $\forall \ell \cdot L_1(\ell) \leq L_2(\ell)$.

301 Our algorithm computes *acquire* and *release* lockstamps AcqLS_e^σ and RelLS_e^σ for every event
 302 $e \in \text{Events}_\sigma$. Let us formalize these next. For a lock ℓ and acquire event $f \in \text{Acquires}_\sigma(\ell)$
 303 (resp. release event $g \in \text{Releases}_\sigma(\ell)$), let $\text{pos}_\sigma(f) = |\{f' \in \text{Acquires}_\sigma(\ell) \mid f' \leq_{\text{tr}}^\sigma f\}|$ (resp.
 304 $\text{pos}_\sigma(g) = |\{f' \in \text{Releases}_\sigma(\ell) \mid f' \leq_{\text{tr}}^\sigma f\}|$) denote the relative position of f (resp. g) among
 305 all acquire events (resp. release events) of ℓ . Then, for an event $e \in \text{Events}_\sigma$ the lockstamps
 306 AcqLS_e^σ and RelLS_e^σ are defined as follows (we assume that $\max \emptyset = 0$ and $\min \emptyset = \infty$.)

$$\text{AcqLS}_e^\sigma(\ell) = \lambda \ell \cdot \max\{\text{pos}_\sigma(f) \mid f \in \text{Acquires}_\sigma(\ell), f \leq_{\text{HB}}^\sigma e\} \quad (1)$$

$$\text{RelLS}_e^\sigma(\ell) = \lambda \ell \cdot \min\{\text{pos}_\sigma(f) \mid f \in \text{Releases}_\sigma(\ell), e \leq_{\text{HB}}^\sigma f\}$$

308 Our $O(\mathcal{N} \cdot \mathcal{L})$ algorithm now relies on the following observations. First, the HB partial order
 309 can be inferred using by comparing lockstamps of events (Lemma 10). Second, there is an
 310 $O(\mathcal{N} \cdot \mathcal{L})$ time algorithm that computes the acquire and release lockstamps for each event
 311 in the input trace. Third, the existence of an HB race can be determined by examining
 312 only $O(\mathcal{N})$ pairs of conflicting events (using their lockstamps), instead of all possible $O(\mathcal{N}^2)$
 313 pairs (Lemma 11). Finally, we can also examine all the $O(\mathcal{N})$ pairs in time $O(\mathcal{N} \cdot \mathcal{L})$ (using
 314 $O(\mathcal{N})$ lockstamp comparisons) and thus determine the existence of an HB race in the same
 315 asymptotic running time. Let us first state how we use lockstamps to infer the HB relation.

316 **► Lemma 10.** *Let $e_1 \leq_{\text{tr}}^\sigma e_2$ be events in σ such that $\text{tid}(e_1) \neq \text{tid}(e_2)$. We have, $e_1 \leq_{\text{HB}}^\sigma$
 317 $e_2 \iff \neg(\text{AcqLS}_{e_2}^\sigma \sqsubseteq \text{RelLS}_{e_1}^\sigma)$*

318 **Computing Lockstamps.** We now illustrate how to compute the acquire lockstamps for
 319 all events, by processing the trace σ in a forward pass. For each thread t and lock ℓ , we
 320 maintain lockstamp variables \mathbb{C}_t and \mathbb{L}_ℓ . We also maintain an integer variable \mathbf{p}_ℓ for each
 321 lock ℓ that stores the index of the latest $\text{acq}(\ell)$ event in σ . Initially, we set each \mathbb{C}_t and \mathbb{L}_m
 322 to the *bottom* map $\lambda \ell \cdot 0$, and \mathbf{p}_m to 0, for each thread t and lock m . We traverse σ left to
 323 right, and perform updates to the data structures as described in Algorithm 1, by invoking
 324 the appropriate *handler* based on the thread and operation of the current event $e = \langle t, \text{op} \rangle$.
 325 At the end of each handler, we assign the lockstamp AcqLS_e^σ to e . The computation of release
 326 lockstamps is similar, albeit in a reverse pass, and presented in Appendix B.1. Observe that
 327 each step takes $O(\mathcal{L})$ time giving us a total running time of $O(\mathcal{N} \cdot \mathcal{L})$ to assign lockstamps.

328 We say that a pair of conflicting access events (e_1, e_2) (with $e_1 \leq_{\text{tr}}^\sigma e_2$) to a variable x is a
 329 *consecutive conflicting pair* if there is no event $f \in \text{Writes}_\sigma(x)$ such that $e_1 <_{\text{tr}}^\sigma f <_{\text{tr}}^\sigma e_2$. We
 330 make the following observation.

331 **► Lemma 11.** *A trace σ has an HB-race iff there is pair of consecutive conflicting events in
 332 σ that is an HB-race. Moreover, σ has $O(\mathcal{N})$ many consecutive conflicting pairs of events.*

333 **Checking for an HB race.** We now describe the algorithm for checking for an HB race in

■ **Algorithm 1** *Assigning acquire lockstamps to events in the trace*

| | | |
|--|--|---|
| 1 acquire (t, ℓ): 2 $p_\ell \leftarrow p_\ell + 1$ 3 $C_t \leftarrow C_t[\ell \mapsto p_\ell] \sqcup L_\ell$ 4 $\text{AcqLS}_e^\sigma \leftarrow C_t$ | 5 release (t, ℓ): 6 $L_\ell \leftarrow C_t$ 7 $\text{AcqLS}_e^\sigma \leftarrow C_t$ | 8 read (t, x): 9 $\text{AcqLS}_e^\sigma \leftarrow C_t$ 10 write (t, x): 11 $\text{AcqLS}_e^\sigma \leftarrow C_t$ |
|--|--|---|

| | |
|-------|-------|
| A_1 | A_2 |
| 101 | 111 |
| 100 | 011 |
| 010 | 110 |

| | | | | | | | |
|-------------|---------------------|-------------|---------------------|-------------|---------------------|-------------|-----------------------------|
| $t_{(1,0)}$ | | $t_{(1,1)}$ | | $t_{(1,2)}$ | | $t_{(1,3)}$ | |
| 1 | $w(z)$ | 4 | $\text{cs}(\ell^x)$ | 8 | $\text{cs}(\ell^x)$ | 10 | $\text{cs}(\ell^x)$ |
| 2 | $\text{cs}(\ell^x)$ | 6 | $\text{cs}(\ell_1)$ | | | 12 | $\text{cs}(\ell_3)$ |
| | | | | | | 14 | $\text{cs}(\ell_{(y_1,3)})$ |
| | | | | | | 16 | $\text{cs}(\ell_{(y_2,3)})$ |
| | | | | | | 18 | $\text{cs}(\ell_{(y_3,3)})$ |

OV instance

■ **Figure 2** Reducing OV to detecting a write-read HB-races. Illustration of the threads $t(x, i)$, where x is the first vector of A_1 . $\text{cs}(\ell)$ denotes the sequence $\text{acq}(\ell), \text{rel}(\ell)$. Event numbers indicate the relative order in which these threads execute in σ .

334 σ . We perform a forward pass on σ while storing the release lockstamps of some of the earlier
335 events. When processing an access event e , we check if it is in race with an earlier event by
336 comparing the acquire lockstamp of e with a previously stored release lockstamp. More
337 precisely, we maintain a variable \mathbb{W}_x to store the release lockstamp of the last write event on
338 x , a variable t_x^w to store the thread that performed this write and set S_x to store pairs (t, L)
339 of threads and release lockstamps of all the read events performed since the last write on x
340 was observed. Initially, $t_x^w = \text{NIL}$, $\mathbb{W}_x = \lambda \ell \cdot \infty$ and $S_x = \emptyset$. The update performed at each
341 event $e = \langle t, op \rangle$ are presented in the corresponding handler in Algorithm 2.

■ **Algorithm 2** *Determining the existence of an HB-race using lockstamps*

| | |
|--|--|
| 1 read (t, x): 2 if $t_x^w \notin \{\text{NIL}, t\} \wedge \text{AcqLS}_e^\sigma \sqsubseteq \mathbb{W}_x$ then 3 declare ‘race’ and exit 4 $S_x \leftarrow S_x \cup \{(t, \text{RelLS}_e^\sigma)\}$ | 5 write (t, x): 6 if $t_x^w \notin \{\text{NIL}, t\} \wedge \text{AcqLS}_e^\sigma \sqsubseteq \mathbb{W}_x$ then 7 declare ‘race’ and exit 8 if $\exists (u, L) \in S_x, t \neq u \wedge \text{AcqLS}_e^\sigma \sqsubseteq L$ then 9 declare ‘race’ and exit 10 $t_x^w = t; S_x \leftarrow \emptyset; \mathbb{W}_x \leftarrow \text{RelLS}_e^\sigma$ |
|--|--|

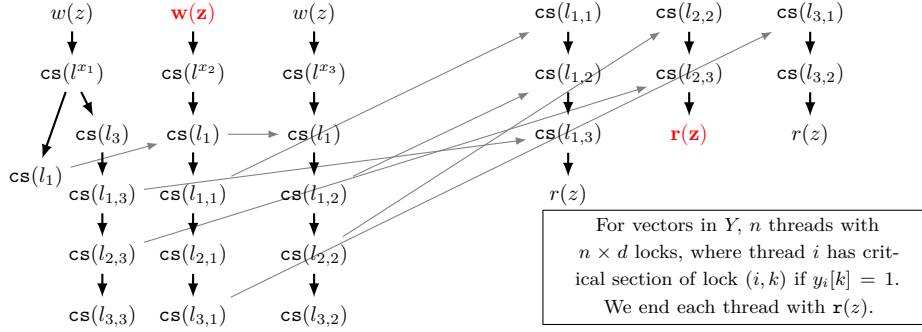
342 We refer to Appendix B.1 for the correctness, which concludes the proof of Theorem 4.

343 3.2 Hardness Results for HB

344 We now turn our attention to the hardness results for HB race detection. To this end, we
345 prove Theorem 1, Theorem 2, and Theorem 3. We start with defining the graph $G(\leq_{\text{HB}}^\sigma)$,
346 which can be thought of as a form of transitive reduction of the HB relation.

347 **The graph $G(\leq_{\text{HB}}^\sigma)$.** Given a trace σ , the graph $G(\leq_{\text{HB}}^\sigma)$ is a graph with node set Events_σ ,
348 and we have an edge (e_1, e_2) in $G(\leq_{\text{HB}}^\sigma)$ iff (i) e_2 is the immediate successor of e_1 wrt the
349 thread order \leq_{TO}^σ , or (ii) e_1 is a $\text{rel}(\ell)$ event, e_2 is a $\text{acq}(\ell)$ event, $e_1 \leq_{\text{tr}}^\sigma e_2$, and there is no
350 intermediate event in σ that accesses lock ℓ . It follows easily that for any two distinct events
351 e_1, e_2 , we have $e_1 \leq_{\text{HB}}^\sigma e_2$ iff e_2 is reachable from e_1 in $G(\leq_{\text{HB}}^\sigma)$. Moreover, every node has
352 out-degree ≤ 2 and thus $G(\leq_{\text{HB}}^\sigma)$ is sparse, while it can be easily constructed in $O(\mathcal{N})$ time.

353 **OV hardness of write-read HB races.** Given a OV instance $\text{OV}(n, d)$ on two vector sets
354 A_1, A_2 , we create a trace σ as follows. For the part A_1 of OV, we introduce $n \cdot (d + 1)$ threads
355 denoted by $t(x, i)$, for $x \in [n], i \in \{0\} \cup [d]$, and d locks, each denoted by ℓ_i , for $i \in [d]$. For



■ **Figure 3** Reducing OV to finding HB races using the instance of Figure 2. For simplicity, we show the graph $G(\leq_{\text{HB}}^{\sigma})$ instead of the trace σ . The HB race is marked in red, corresponding to the orthogonal pair (x_2, y_2) .

356 the second part A_2 we introduce $n \cdot d$ locks denoted by $l(y, i)$, for $y \in [n], i \in [d]$, and n
 357 threads, denoted by t_y , for $y \in [n]$. Finally, we have a single variable z .

358 We first describe the threads $t(x, i)$. We order the vectors in A_1 arbitrarily. For each vector
 359 x , for each $i \in [d]$, we introduce a critical section on the lock l_i . If x is the last vector of
 360 A_1 with $x[i] = 1$, we also insert the critical sections $l_{(y,i)}$ for all $y \in [n]$, to $t(x, i)$ after the
 361 critical section of l_x . Finally, we construct a thread $t_{x,0}$ which starts with a write event $w(z)$,
 362 followed by a critical section on lock l^x . We also insert a critical section on lock l^x to all
 363 threads $t(x, i)$, for $i \in [d]$. Hence the $w(z)$ event is ordered by HB before all other events of
 364 $t(x, i)$. See Figure 2 for an illustration.

365 Now we describe the threads t_y . For each $i \in [d]$, if $y[i] = 1$, we add a critical section of the
 366 lock $l(y, i)$ in t_y . We end the thread with a read event $r(z)$.

367 Finally, we construct σ by first executing each thread $t(x, i)$ in the pre-determined order of
 368 $x \in A_1$, followed by executing the traces t_y in any order. See Figure 3 for an illustration.
 369 We refer to Appendix B for the correctness, which concludes the proof of Theorem 1.

370 We now turn our attention to the problem of detecting a single HB race (i.e., not necessarily
 371 involving a read event). We define a useful multi-connectivity problem on graphs.

372 **► Problem 1.** [MCONN] Given a directed graph G with n nodes and m edges, and k pairs
 373 of nodes $(s_i, t_i), i \in [k]$, decide if there is a path in G from every s_i to the corresponding t_i .

374 Due to Lemma 11, detecting whether there is an HB race in σ reduces to testing MCONN
 375 between all $O(\mathcal{N})$ pairs of consecutive conflicting events in σ .

376 **Short witnesses for HB races.** We now prove Theorem 2. Following [9, Corollary 2], it
 377 suffices to show that deciding MCONN can be done in $\text{NTIME}[\mathcal{N}^{3/2}] \cap \text{coNTIME}[\mathcal{N}^{3/2}]$. At
 378 a first glance, the bound $\text{NTIME}[\mathcal{N}^{3/2}]$ may seem too optimistic, as there are $\Theta(\mathcal{N})$ paths
 379 $P_i: s_i \rightsquigarrow t_i$, and each of them can have size $\Theta(\mathcal{N})$. Hence even just guessing these paths
 380 appears to take quadratic time. Our proof shows that more succinct witnesses exist.

381 **Proof of Theorem 2.** First consider the simpler case where σ has an HB-race. Phrased as a
 382 MCONN problem on $G(\leq_{\text{HB}}^{\sigma})$, it suffices to show that there is a pair (s_i, t_i) such that s_i does
 383 not reach t_i . We construct a non-deterministic algorithm for this task that simply guesses
 384 the pair (s_i, t_i) , and verifies that there is no $s_i \rightsquigarrow t_i$ path. Since $G(\leq_{\text{HB}}^{\sigma})$ is sparse, this can
 385 be easily verified in $O(\mathcal{N})$ time.

386 Now consider the case when there is no HB-race. Phrased as a MCONN problem on $G(\leq_{\text{HB}}^{\sigma})$,
 387 it suffices to verify that for every pair (s_i, t_i) , we have that s_i reaches t_i . We construct a
 388 non-deterministic algorithm for this task, as follows. The algorithm operates in two phases,

389 using a set A , initialized as $A = \{(s_i, t_i)\}_{i \in k}$.

- 390 1. In the first phase, the algorithm repeatedly guesses a node u that lies on at least $\mathcal{N}^{1/2}$
391 paths $s_i \rightsquigarrow t_i$, for $(s_i, t_i) \in A$. It verifies this guess via a backward and a forward traversal
392 from u . The algorithm then removes all such (s_i, t_i) from A , and repeats.
- 393 2. In the second phase, the algorithm guesses for every remaining $(s_i, t_i) \in A$ a path
394 $P_i: s_i \rightsquigarrow t_i$, and verifies that P_i is a valid path.

395 Phase 1 can be execute at most $\mathcal{N}^{1/2}$ iterations, while each iteration takes $O(\mathcal{N})$ time since
396 $G(\leq_{\text{HB}}^\sigma)$ is sparse. Hence the total time for phase 1 is $O(\mathcal{N}^{3/2})$. Phase 2 takes $O(\mathcal{N}^{3/2})$ time,
397 as every node of $G(\leq_{\text{HB}}^\sigma)$ appears in at most $\mathcal{N}^{1/2}$ paths P_i . The desired result follows. ◀

398 **A super-linear lower bound for general HB races.** Finally, we turn our attention to
399 Theorem 3. The problem $\text{FO}(\forall\exists\exists)$ takes as input a first-order formula ϕ with quantifier
400 structure $\forall\exists\exists$ and whose atoms are tuples, and the task is to verify whether ϕ has a model
401 on a structure of n elements and m relational tuples. For simplicity, we can think of the
402 structure as a graph G of n nodes and m edges, and ϕ a formula that characterizes the
403 presence/absence of edges (e.g., $\phi = \forall x \exists y \exists z e(x, y) \wedge \neg e(y, z)$).

404 The crux of the proof of Theorem 3 is showing the following lemma.

405 ▶ **Lemma 12.** *FO($\forall\exists\exists$) reduces to MCONN on a graph G with $O(n)$ nodes in $O(n^2)$ time.*

406 Finally, we arrive at Theorem 3 by constructing in $O(n^2)$ time a trace σ with $\mathcal{N} = \Theta(n^2)$
407 such that $G(\leq_{\text{HB}}^\sigma)$ is similar in structure to the graph G of Lemma 12. In the end, detecting
408 an HB race in σ in $O(\mathcal{N}^{1+\epsilon})$ time yields an algorithm for $\text{FO}(\forall\exists\exists)$ in $\mathcal{N} = \Theta(n^{2+\epsilon})$ time.
409 We refer to Appendix B for the details, which conclude the proof of Theorem 3.

410 4 Synchronization-Preserving Races

411 In this section we discuss the dynamic detection of sync-preserving races, and prove Theorem 5.

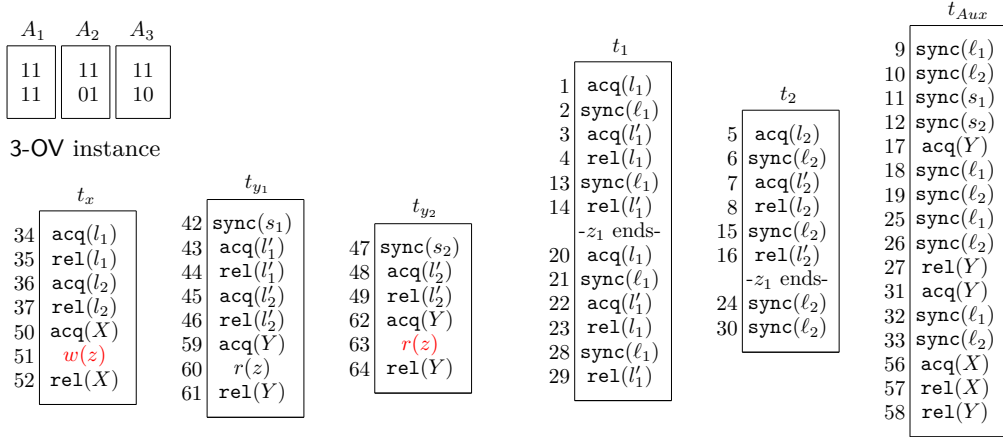
412 For notational convenience, we will frequently use the composite *sync* events. A $\text{sync}(\ell)$
413 event represents the sequence $\text{acq}(\ell), \text{r}(x_\ell), \text{w}(x_\ell), \text{rel}(\ell)$. The key idea behind sync events is
414 as follows. Assume that in a trace σ we have two $\text{sync}(\ell)$ events e_1 and e_2 with $e_1 <_{\text{tr}}^\sigma e_2$.
415 Then any correct reordering ρ of σ with $e_2 \in \text{Events}_\rho$ satisfies the following.

- 416 (a) We have $e_1 \in \text{Events}_\rho$, as the read event of e_2 must read from the write event of e_1 .
- 417 (b) For every $e'_1, e'_2 \in \text{Events}_\rho$ such that $e'_1 \leq_{\top 0}^\sigma e_1$ and $e_2 \leq_{\top 0}^\sigma e'_2$, we have $e_1 <_{\text{tr}}^\rho e_2$.

418 We hence use sync events to ensure certain orderings in any sync-preserving correct reordering
419 of σ that exposes a sync-preserving data race.

420 **Intuition.** Before we proceed with the detailed reduction, we provide a high-level description.
421 The input to 3-OV is three sets of vectors $A_1 = \{x_i\}_{i \in [n]}$, $A_2 = \{y_i\}_{i \in [n]}$, and $A_3 =$
422 $\{z_i\}_{i \in [n]}$. Every vector $x \in A_1$ is represented by a thread t_x , ending with the critical section
423 $\text{acq}(X), \text{w}(z), \text{rel}(X)$. Similarly, every vector $y \in A_1$ is represented by a thread t_y , ending
424 with the critical section $\text{acq}(Y), \text{r}(z), \text{rel}(Y)$. Notice that we can only have a race between
425 the write event of a thread t_x and the read event of a thread t_y . The search for such a race
426 corresponds to the search of the corresponding vectors $x \in A_1$ and $y \in A_2$ such that there is
427 a vector $z \in A_3$ which makes the triplet x, y, z orthogonal.

428 To establish this correspondence, we insert in t_x empty critical sections on locks l_k , for $k \in [d]$
429 that represent the coordinates k for which $x[k] = 1$. We use a similar encoding with locks l'_k
430 for the threads t_y , capturing that $y[k] = 1$. To encode the vectors in A_3 , we use k threads t_k ,
431 for $k \in [d]$, such that the i^{th} segment of t_k encodes $z_i[k]$: we have two interleaved critical



■ **Figure 4** Example reduction from 3-OV to `sync` race detection. The trace orders events as shown by their numbering. We only show one thread t_x , as the two x vectors are identical.

sections on locks l_k and l'_k iff $z_i[k] = 1$.

Finally, we use some `sync` events to force all threads t_k be partially executed whenever we want to execute the write event of any thread t_x . Hence, any correct reordering of σ that exposes a data race in σ , must execute all t_k at least partially. We make all threads t_k execute before all t_x and t_y in σ . The notion of sync-preservation ensures that if we have a correct reordering that exposes a race between two threads t_x and t_y , then the following holds. For every coordinate $k \in [d]$ in which $x[k] = y[k] = 1$, since the corresponding threads t_x and t_y have critical sections on locks l_k and l'_k , the thread t_k must execute up to a point where it does not have critical sections on these locks. This means that we have found a vector z with $z[k] = 0$, and thus the triplet x, y, z is orthogonal on that coordinate.

Reduction. Given an 3-OV instance $\text{OV}(n, d, 3)$ on vector sets $A_1 = \{x_i\}_{i \in [n]}$, $A_2 = \{y_i\}_{i \in [n]}$, and $A_3 = \{z_i\}_{i \in [n]}$, we create a trace σ as follows (see Figure 4). We have $\mathcal{T} = 2 \cdot n + d + 1$ threads, while all access events (not counting the `sync` events) are of the form $w(z)/r(z)$ in a single variable z . We first describe the threads, and then how they interleave in σ .

Threads. We introduce a thread t_x for every vector $x \in A_1$ and a lock l_k for every $k \in [d]$. Each thread t_x consists of two segments t_x^1 and t_x^2 . We create t_x^1 as follows. For every $k \in [d]$ where $x[k] = 1$, we add an empty critical section $\text{acq}(l_k), \text{rel}(l_k)$ in t_x^1 . We create t_x^2 as the sequence $\text{acq}(X), w(z), \text{rel}(X)$, where X is a new lock, common for all t_x^2 .

For the vectors in A_2 , we introduce threads similar to those of part A_1 , as follows. We have a thread t_y for every vector $y \in A_2$ and a lock l'_k for every $k \in [d]$. Each thread t_y consists of two segments t_y^1 and t_y^2 . For every $k \in [d]$ where $y[k] = 1$, we add an empty critical section $\text{acq}(l'_k), \text{rel}(l'_k)$ in t_y^1 . In contrast to the t_x^1 , every t_y^1 also has an event `sync`(s_y) at the very beginning. We create t_y^2 as the sequence $\text{acq}(Y), r(z), \text{rel}(Y)$, where Y is a new lock, common for all t_y^2 .

The construction of the threads corresponding to the vectors in A_3 is more involved. We have one thread t_k for every $k \in [d]$. Each thread has some fixed `sync` events, as well as critical sections corresponding to one coordinate of all n vectors in A_3 . In particular, we construct each t_k as follows. We iterate over all z_i , and if $z_i[k] = 0$, we simply append two events `sync`(ℓ_k), `sync`(ℓ_k) to t_k . On the other hand, if $z_i[k] = 1$, we interleave these sync events with two critical sections, by appending the sequence $\text{acq}(l_k), \text{sync}(\ell_k), \text{acq}(l'_k), \text{rel}(l_k), \text{sync}(\ell_k), \text{rel}(l'_k)$.

463 Lastly, we have a single auxiliary trace t that consists of three parts t^1 , t^2 and t^3 , where

$$\begin{aligned}
 464 \quad t^1 &= \text{sync}(\ell_1), \dots, \text{sync}(\ell_k), \text{sync}(s_{y_1}), \dots, \text{sync}(s_{y_n}) \\
 465 \quad t^2 &= (\text{acq}(Y), \text{sync}(\ell_1), \dots, \text{sync}(\ell_k), \text{sync}(\ell_1), \dots, \text{sync}(\ell_k), \text{rel}(Y))^{n-1} \\
 466 \quad t^3 &= \text{acq}(Y), \text{sync}(\ell_1), \dots, \text{sync}(\ell_k), \text{acq}(X), \text{rel}(X), \text{rel}(Y) \\
 467
 \end{aligned}$$

468 *Concurrent trace.* We are now ready to describe the interleaving of the above threads in
 469 order to obtain the concurrent trace σ .

- 470 1. We execute the auxiliary trace t and all traces t_k , for $k \in [d]$ (i.e., the threads corresponding
 471 to the vectors of A_3) arbitrarily, as long as for every $k \in [d]$, every sequence of $\text{sync}(\ell_k)$
 472 events (a) starts with the $\text{sync}(\ell_k)$ event of t_k and proceeds with the $\text{sync}(\ell_k)$ event of t ,
 473 (b) strictly alternates in every two $\text{sync}(\ell_k)$ events between t and t_k , and (c) ends with
 474 the last $\text{sync}(\ell_k)$ event of t_k .
- 475 2. We execute all t_x^1 and t_y^1 (i.e., the first parts of all threads that correspond to the vectors
 476 in A_1 and A_2) arbitrarily, but after all traces t_k , for $k \in [d]$.
- 477 3. We execute all t_x^2 (i.e., the second parts of all traces that correspond to the vectors in
 478 A_1) arbitrarily, but before the segment $\text{acq}(X), \text{rel}(X), \text{rel}(Y)$ of t .
- 479 4. We execute all t_y^2 (i.e., the second parts of all traces that correspond to the vectors in
 480 A_2) arbitrarily, but after the segment $\text{acq}(X), \text{rel}(X), \text{rel}(Y)$ of t .

481 We refer to Appendix C for the correctness of the reduction and thus the proof of Theorem 5.

482 5 Violations of the Locking Discipline

483 5.1 Lock-Cover Races

484 We start with a simple reduction from OV to detecting lock-cover races. Given an OV instance
 485 $\text{OV}(n, d)$ on two vector sets A_1, A_2 , we create a trace σ as follows. We have a single variable
 486 x and two threads t_1, t_2 . We associate with each vector of the set A_i a write access event
 487 $e = \langle t_i, \mathbf{w}(x) \rangle$. Moreover, each such event holds up to d locks, so that e holds the k^{th} lock
 488 iff k^{th} coordinate of the vector corresponding to the event is 1. The trace σ is formed by
 489 ordering the sequence of events corresponding to vectors of A_1 of OV first, in a fixed arbitrary
 490 order, followed by the sequence of events corresponding to A_2 , again in arbitrary order. We
 491 refer to Appendix D for the correctness, which concludes the proof of Theorem 6.

492 5.2 Lock-Set Races

493 We now turn our attention to lock-set races. We first prove Theorem 9, i.e., that determining
 494 whether a trace σ has a lock-set race on a specific variable x can be performed in linear time.

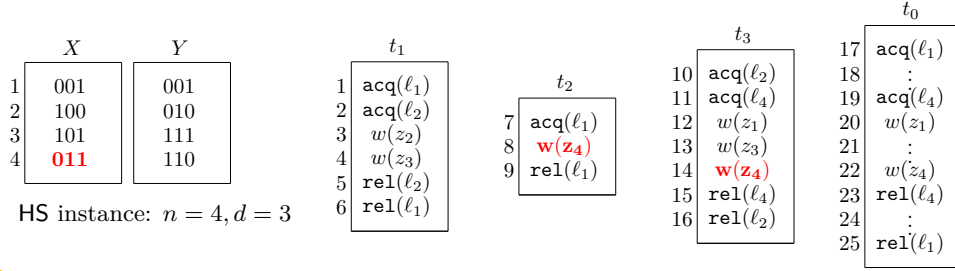
495 **A linear-time algorithm per variable.** Verifying that there are two conflicting events on
 496 x is straightforward by a single pass of σ . The more involved part is in computing the lock-set
 497 of x , i.e., the set $\bigcap_{e \in \text{Accesses}_\sigma(x)} \text{locksHeld}_\sigma(e)$, in linear time. Indeed, each intersection alone
 498 requires $\Theta(\mathcal{L})$ time, resulting to $\Theta(\mathcal{N} \cdot \mathcal{L})$ time overall.

499 Here we show that a somewhat more involved algorithm achieves the task. The algorithm
 500 performs a single pass of σ , while maintaining three simple sets A, B , and C . While processing
 501 an event e , the sets are updated to maintain the invariant

$$502 \quad A = \text{locksHeld}_\sigma(e) \quad B = \text{Locks}_\sigma \cap \bigcap_{e' \in \text{Accesses}_\sigma(x), e' \leq_{\text{tr}}^\sigma e} \text{locksHeld}_\sigma(e') \quad C = \bar{A} \cap B \quad (2)$$

503

23:14 Dynamic Data-Race Detection through the Fine-Grained Lens



■ **Figure 5** Reducing HS to detecting a lock-set race on trace σ with d threads. Thread t_k uses lock l_i if $y_i[k] = 0$, and $w(z_j)$ if $x_j[k] = 1$. Vector x_4 hits all vectors in Y , implying a lock-set race on z_4 .

504 The sets are initialized as $A = \emptyset$, $B = C = \text{Locks}_\sigma$. Then the algorithm performs a pass
505 over σ and processes each event e according to the description of Algorithm 3.

■ **Algorithm 3** Computing lock-set of variable x

| | | | |
|--|--|---|---|
| 1 acquire (t, ℓ): 2 $A \leftarrow A \cup \{\ell\}$ 3 if $\ell \in B$ then 4 $C \leftarrow C \setminus \{\ell\}$ | 5 release (t, ℓ): 6 $A \leftarrow A \setminus \{\ell\}$ 7 if $\ell \in B$ then 8 $C \leftarrow C \cup \{\ell\}$ | 9 read (t, y): 10 if $x = y$ then 11 $B \leftarrow B \setminus C$ 12 $C \leftarrow \emptyset$ | 13 write (t, y): 14 if $x = y$ then 15 $B \leftarrow B \setminus C$ 16 $C \leftarrow \emptyset$ |
|--|--|---|---|

506 The correctness of Algorithm 3 follows by proving the invariant in Equation (2). We refer to
507 Appendix D for the details, which concludes the proof of Theorem 9.

508 **Short witnesses for lock-set races.** Besides the advantage of a faster algorithm, The-
509 orem 9 implies that lock-set races have short witnesses that can be verified in linear time.
510 This allows us to prove that detecting a lock-set race is in $\text{NTIME}[\mathcal{N}] \cap \text{coNTIME}[\mathcal{N}]$, and
511 we can thus use [9, Corollary 2] to prove Theorem 7.

512 **Proof of Theorem 7.** First we argue that the problem is in $\text{NTIME}[n]$. Indeed, the certificate
513 for the existence of a lock-set race is simply the variable x on which there is a lock-set race.
514 By Theorem 9, verifying that we indeed have a lock-set race on x takes $O(\mathcal{N})$ time.

515 Now we argue that the problem is in $\text{coNTIME}[n]$, by giving a certificate to verify in linear
516 time that σ does not have a race of the required form. The certificate has size $O(|\text{Vars}_\sigma|)$,
517 and specifies for every variable, either the lock that is held by all access events of the variable,
518 or a claim that there exist no two conflicting events on that variable. The certificate can be
519 easily verified by one pass over σ . ◀

520 **Lock-set races are Hitting-Set hard.** Finally we prove Theorem 8, i.e., that determining
521 a single lock-set race is HS-hard, and thus also carries a conditional quadratic lower bound.
522 We establish a fine-grained reduction from HS. Given a HS instance $\text{HS}(n, d)$ on two vector sets
523 X, Y , we create a trace σ using $d + 1$ threads $\{t_j\}_{j \in \{0\} \cup [d]}$, n locks $\{\ell_i\}_{i \in [n]}$, and n variables
524 $\{z_i\}_{i \in [n]}$. Thread t_0 that executes $\text{acq}(\ell_1), \dots, \text{acq}(\ell_n), \mathbf{w}(z_1), \dots, \mathbf{w}(z_n), \text{rel}(\ell_n), \dots, \text{rel}(\ell_1)$.
525 Each of the threads t_j , for $j \in [d]$, has a single nested critical section consisting of the locks
526 $\ell_i \in [n]$ such that the i^{th} vector of Y has its j^{th} coordinate 0, i.e. $y_i[j] = 0$. The events in
527 the critical section are all write events of all variables $z_k \in [n]$ with $x_k[j] = 1$. The trace
528 orders all events of each thread t_d consecutively, and all the events overall in increasing order
529 of d . See Figure 5 for an illustration. We refer to Appendix D for the correctness, which
530 concludes the proof of Theorem 8.

531 **6 Conclusion**

532 In this work we have taken a fine-grained view of the complexity of popular notions of
533 dynamic data races. We have established a range of lower bounds on the complexity of
534 detecting HB races, sync-preserving races, as well as races based on the locking discipline
535 (lock-cover/lock-set races). Moreover, we have characterized cases where lower bounds based
536 on SETH are not possible under NSETH. Finally, we have proven new upper bounds for
537 detecting HB and lock-set races. To our knowledge, this is the first work that characterizes
538 the complexity of well-established dynamic race-detection techniques, allowing for a rigorous
539 characterization of their trade-offs between expressiveness and running time.

- 541 **1** Helgrind: a thread error detector. <https://valgrind.org/docs/manual/hg-manual.html>.
542 Accessed: 2021-04-30.
- 543 **2** Intel Inspector. [https://software.intel.com/content/www/us/en/develop/tools/oneapi/
544 components/inspector.html](https://software.intel.com/content/www/us/en/develop/tools/oneapi/components/inspector.html). Accessed: 2021-04-30.
- 545 **3** Amir Abboud, Virginia Vassilevska Williams, and Joshua Wang. Approximation and fixed
546 parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proceedings
547 of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16,
548 page 377–391, USA, 2016. Society for Industrial and Applied Mathematics.
- 549 **4** Utpal Banerjee, Brian Bliss, Zhiqiang Ma, and Paul Petersen. A theory of data race detection.
550 In *Proceedings of the 2006 Workshop on Parallel and Distributed Systems: Testing and
551 Debugging*, PADTAD '06, pages 69–78, New York, NY, USA, 2006. ACM. URL: [http:
552 //doi.acm.org/10.1145/1147403.1147416](http://doi.acm.org/10.1145/1147403.1147416), doi:10.1145/1147403.1147416.
- 553 **5** Hans-J. Boehm. How to miscompile programs with “benign” data races. In *Proceedings of the
554 3rd USENIX Conference on Hot Topic in Parallelism*, HotPar'11, page 3, USA, 2011. USENIX
555 Association.
- 556 **6** Hans-J. Boehm. Position paper: Nondeterminism is unavoidable, but data races are pure
557 evil. In *Proceedings of the 2012 ACM Workshop on Relaxing Synchronization for Multicore
558 and Manycore Scalability*, RACES '12, page 9–14, New York, NY, USA, 2012. Association for
559 Computing Machinery. URL: <https://doi.org/10.1145/2414729.2414732>, doi:10.1145/
560 2414729.2414732.
- 561 **7** Hans-J. Boehm and Sarita V. Adve. Foundations of the C++ Concurrency Memory Model.
562 In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design
563 and Implementation*, PLDI '08, page 68–78, New York, NY, USA, 2008. Association for
564 Computing Machinery. URL: <https://doi.org/10.1145/1375581.1375591>, doi:10.1145/
565 1375581.1375591.
- 566 **8** Karl Bringmann. Fine-Grained Complexity Theory (Tutorial). In Rolf Niedermeier and
567 Christophe Paul, editors, *36th International Symposium on Theoretical Aspects of Com-
568 puter Science (STACS 2019)*, volume 126 of *Leibniz International Proceedings in In-
569 formatics (LIPIcs)*, pages 4:1–4:7, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-
570 Zentrum fuer Informatik. URL: [http://drops.dagstuhl.de/opus/
571 volltexte/2019/10243](http://drops.dagstuhl.de/opus/volltexte/2019/10243), doi:10.4230/LIPIcs.STACS.2019.4.
- 572 **9** Marco L Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and
573 Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and
574 consequences for non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations
575 in Theoretical Computer Science*, pages 261–270, 2016.
- 576 **10** Bernadette Charron-Bost. Concerning the size of logical clocks in distributed systems. *Informa-
577 tion Processing Letters*, 39(1):11 – 16, 1991. URL: [http://www.sciencedirect.com/science/
578 article/pii/002001909190055M](http://www.sciencedirect.com/science/article/pii/002001909190055M), doi:[https://doi.org/10.1016/0020-0190\(91\)90055-M](https://doi.org/10.1016/0020-0190(91)90055-M).
- 579 **11** Peter Chini, Jonathan Kolberg, Andreas Krebs, Roland Meyer, and Prakash Saivasan. On
580 the Complexity of Bounded Context Switching. In Kirk Pruhs and Christian Sohler, editors,
581 *25th Annual European Symposium on Algorithms (ESA 2017)*, volume 87 of *Leibniz Inter-
582 national Proceedings in Informatics (LIPIcs)*, pages 27:1–27:15, Dagstuhl, Germany, 2017.
583 Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: [http://drops.dagstuhl.de/opus/
584 volltexte/2017/7873](http://drops.dagstuhl.de/opus/volltexte/2017/7873), doi:10.4230/LIPIcs.ESA.2017.27.
- 585 **12** Peter Chini, Roland Meyer, and Prakash Saivasan. Fine-grained complexity of safety verifica-
586 tion. In Dirk Beyer and Marieke Huisman, editors, *Tools and Algorithms for the Construction
587 and Analysis of Systems*, pages 20–37, Cham, 2018. Springer International Publishing.

- 588 13 Peter Chini and Prakash Saivasan. A Framework for Consistency Algorithms. In Nitin
589 Saxena and Sunil Simon, editors, *40th IARCS Annual Conference on Foundations of Soft-*
590 *ware Technology and Theoretical Computer Science (FSTTCS 2020)*, volume 182 of *Leibniz*
591 *International Proceedings in Informatics (LIPIcs)*, pages 42:1–42:17, Dagstuhl, Germany, 2020.
592 Schloss Dagstuhl–Leibniz-Zentrum für Informatik. URL: [https://drops.dagstuhl.de/opus/](https://drops.dagstuhl.de/opus/volltexte/2020/13283)
593 [volltexte/2020/13283](https://drops.dagstuhl.de/opus/volltexte/2020/13283), doi:10.4230/LIPIcs.FSTTCS.2020.42.
- 594 14 Anne Dinning and Edith Schonberg. Detecting access anomalies in programs with critical
595 sections. In *Proceedings of the 1991 ACM/ONR Workshop on Parallel and Distributed*
596 *Debugging*, PADD '91, pages 85–96, New York, NY, USA, 1991. ACM. URL: [http://doi.](http://doi.acm.org/10.1145/122759.122767)
597 [acm.org/10.1145/122759.122767](http://doi.acm.org/10.1145/122759.122767), doi:10.1145/122759.122767.
- 598 15 Tayfun Elmas, Shaz Qadeer, and Serdar Tasiran. Goldilocks: A race and transaction-aware java
599 runtime. In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language*
600 *Design and Implementation*, PLDI '07, pages 245–255, New York, NY, USA, 2007. ACM.
601 URL: <http://doi.acm.org/10.1145/1250734.1250762>, doi:10.1145/1250734.1250762.
- 602 16 Cormac Flanagan and Stephen N. Freund. Fasttrack: Efficient and precise dynamic race
603 detection. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language*
604 *Design and Implementation*, PLDI '09, pages 121–133, New York, NY, USA, 2009. ACM.
605 URL: <http://doi.acm.org/10.1145/1542476.1542490>, doi:10.1145/1542476.1542490.
- 606 17 Jiawei Gao, Russell Impagliazzo, Antonina Kolokolova, and Ryan Williams. Completeness
607 for first-order properties on sparse structures with algorithmic applications. *ACM Trans.*
608 *Algorithms*, 15(2), December 2018. URL: <https://doi.org/10.1145/3196275>, doi:10.1145/
609 3196275.
- 610 18 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *Journal of Computer*
611 *and System Sciences*, 62(2):367–375, 2001.
- 612 19 Ayal Itzkovitz, Assaf Schuster, and Oren Zeev-Ben-Mordehai. Toward integration of data race
613 detection in dsm systems. *J. Parallel Distrib. Comput.*, 59(2):180–203, November 1999. URL:
614 <http://dx.doi.org/10.1006/jpdc.1999.1574>, doi:10.1006/jpdc.1999.1574.
- 615 20 Baris Kasikci, Cristian Zamfir, and George Candea. Racemob: Crowdsourced data race detec-
616 tion. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*,
617 SOSP '13, page 406–422, New York, NY, USA, 2013. Association for Computing Machinery.
618 URL: <https://doi.org/10.1145/2517349.2522736>, doi:10.1145/2517349.2522736.
- 619 21 Dileep Kini, Umang Mathur, and Mahesh Viswanathan. Dynamic race prediction in linear
620 time. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language*
621 *Design and Implementation*, PLDI 2017, pages 157–170, New York, NY, USA, 2017. ACM.
622 URL: <http://doi.acm.org/10.1145/3062341.3062374>, doi:10.1145/3062341.3062374.
- 623 22 Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun.*
624 *ACM*, 21(7):558–565, July 1978. URL: <http://doi.acm.org/10.1145/359545.359563>, doi:
625 10.1145/359545.359563.
- 626 23 Umang Mathur, Dileep Kini, and Mahesh Viswanathan. What happens-after the first race?
627 enhancing the predictive power of happens-before based dynamic race detection. *Proc. ACM*
628 *Program. Lang.*, 2(OOPSLA):145:1–145:29, October 2018. URL: [http://doi.acm.org/10.](http://doi.acm.org/10.1145/3276515)
629 [1145/3276515](http://doi.acm.org/10.1145/3276515), doi:10.1145/3276515.
- 630 24 Umang Mathur, Andreas Pavlogiannis, and Mahesh Viswanathan. The complexity of dynamic
631 data race prediction. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic*
632 *in Computer Science*, LICS '20, page 713–727, New York, NY, USA, 2020. Association for
633 Computing Machinery. URL: <https://doi.org/10.1145/3373718.3394783>, doi:10.1145/
634 3373718.3394783.
- 635 25 Umang Mathur, Andreas Pavlogiannis, and Mahesh Viswanathan. Optimal prediction of
636 synchronization-preserving races. *Proc. ACM Program. Lang.*, 5(POPL), January 2021. URL:

- 637 <https://doi.org/10.1145/3434317>, doi:10.1145/3434317.
- 638 **26** Satish Narayanasamy, Zhenghao Wang, Jordan Tigani, Andrew Edwards, and Brad Calder.
639 Automatically classifying benign and harmful data races using replay analysis. In *Proceedings*
640 *of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation*,
641 PLDI '07, page 22–31, New York, NY, USA, 2007. Association for Computing Machinery.
642 URL: <https://doi.org/10.1145/1250734.1250738>, doi:10.1145/1250734.1250738.
- 643 **27** Robert O'Callahan and Jong-Deok Choi. Hybrid dynamic data race detection. *SIGPLAN*
644 *Not.*, 38(10):167–178, June 2003. URL: <http://doi.acm.org/10.1145/966049.781528>, doi:
645 10.1145/966049.781528.
- 646 **28** Andreas Pavlogiannis. Fast, sound, and effectively complete dynamic race prediction. *Proc.*
647 *ACM Program. Lang.*, 4(POPL), December 2019. URL: <https://doi.org/10.1145/3371085>,
648 doi:10.1145/3371085.
- 649 **29** Eli Pozniansky and Assaf Schuster. Efficient on-the-fly data race detection in multithreaded
650 c++ programs. *SIGPLAN Not.*, 38(10):179–190, June 2003. URL: [http://doi.acm.org/10.](http://doi.acm.org/10.1145/966049.781529)
651 [1145/966049.781529](http://doi.acm.org/10.1145/966049.781529), doi:10.1145/966049.781529.
- 652 **30** Jake Roemer, Kaan Genç, and Michael D. Bond. High-coverage, unbounded sound predictive
653 race detection. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming*
654 *Language Design and Implementation*, PLDI 2018, pages 374–389, New York, NY, USA,
655 2018. ACM. URL: <http://doi.acm.org/10.1145/3192366.3192385>, doi:10.1145/3192366.
656 3192385.
- 657 **31** Grigore Rosu. RV-Predict, Runtime Verification. [https://runtimeverification.com/](https://runtimeverification.com/predict/)
658 [predict/](https://runtimeverification.com/predict/), 2018. Accessed: 2018-04-01.
- 659 **32** Mahmoud Said, Chao Wang, Zijiang Yang, and Karem Sakallah. Generating data race
660 witnesses by an smt-based analysis. In *Proceedings of the Third International Conference on*
661 *NASA Formal Methods*, NFM'11, pages 313–327, Berlin, Heidelberg, 2011. Springer-Verlag.
662 URL: <http://dl.acm.org/citation.cfm?id=1986308.1986334>.
- 663 **33** Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro, and Thomas Anderson.
664 Eraser: A dynamic data race detector for multithreaded programs. *ACM Trans. Comput.*
665 *Syst.*, 15(4):391–411, November 1997. URL: <http://doi.acm.org/10.1145/265924.265927>,
666 doi:10.1145/265924.265927.
- 667 **34** Koushik Sen, Grigore Roşu, and Gul Agha. Detecting errors in multithreaded programs
668 by generalized predictive analysis of executions. In Martin Steffen and Gianluigi Zavattaro,
669 editors, *Formal Methods for Open Object-Based Distributed Systems*, pages 211–226, Berlin,
670 Heidelberg, 2005. Springer Berlin Heidelberg.
- 671 **35** Konstantin Serebryany and Timur Iskhodzhanov. ThreadSanitizer: Data Race Detection in
672 Practice. WBIA '09, 2009.
- 673 **36** Jaroslav Ševčík and David Aspinall. On validity of program transformations in the java
674 memory model. In Jan Vitek, editor, *ECOOP 2008 – Object-Oriented Programming*, pages
675 27–51, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- 676 **37** Yannis Smaragdakis, Jacob Evans, Caitlin Sadowski, Jaeheon Yi, and Cormac Flanagan. Sound
677 predictive race detection in polynomial time. In *Proceedings of the 39th Annual ACM SIGPLAN-*
678 *SIGACT Symposium on Principles of Programming Languages*, POPL '12, pages 387–400,
679 New York, NY, USA, 2012. ACM. URL: <http://doi.acm.org/10.1145/2103656.2103702>,
680 doi:10.1145/2103656.2103702.
- 681 **38** Martin Sulzmann and Kai Stadtmüller. Efficient, near complete, and often sound hybrid
682 dynamic data race prediction. In *Proceedings of the 17th International Conference on Managed*
683 *Programming Languages and Runtimes*, MPLR 2020, page 30–51, New York, NY, USA, 2020.
684 Association for Computing Machinery. URL: <https://doi.org/10.1145/3426182.3426185>,
685 doi:10.1145/3426182.3426185.

- 686 39 Christoph von Praun. *Race Detection Techniques*, pages 1697–1706. Springer US, Bo-
687 ston, MA, 2011. URL: [https://doi.org/10.1007/](https://doi.org/10.1007/978-0-387-09766-4_38)
688 [978-0-387-09766-4_38](https://doi.org/10.1007/978-0-387-09766-4_38).
- 689 40 Jaroslav Ševčík. Safe optimisations for shared-memory concurrent programs. *SIGPLAN*
690 *Not.*, 46(6):306–316, June 2011. URL: <https://doi.org/10.1145/1993316.1993534>, doi:
691 [10.1145/1993316.1993534](https://doi.org/10.1145/1993316.1993534).
- 692 41 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications.
693 *Theoretical Computer Science*, 348(2-3):357–365, 2005.
- 694 42 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity.
695 In *Proceedings of the ICM*, volume 3, pages 3431–3472. World Scientific, 2018.
- 696 43 M. Zhivich and R. K. Cunningham. The real cost of software errors. *IEEE Security and*
697 *Privacy*, 7(2):87–90, March 2009. URL: <https://doi.org/10.1109/MSP.2009.56>, doi:10.
698 [1109/MSP.2009.56](https://doi.org/10.1109/MSP.2009.56).

699 **A** Fine-Grained Complexity and Popular Hypotheses

700 In this section we present notions of fine-grained complexity theory that are relevant to our
701 work. We refer to the survey [42] for a detailed exposition on the topic.

702 This theory relates the computational complexity of problems under the following, more
703 refined, notion of reduction than the standard ones used in traditional complexity theory.
704 Informally, the definition says that if there is an algorithm for some problem B faster than its
705 assumed lower bound, then such a reduction from some problem A to B gives an algorithm
706 for A that is faster than its conjectured lower bound.

Fine-grained Reductions. Assume that A and B are computational problems and $a(n)$
and $b(n)$ are their conjectured running time lower bounds, respectively. Then we say A
(a, b)-reduces to B, denoted by $A \underset{(a)}{\preceq}_{(b)} B$, if for every $\epsilon > 0$, there exists $\delta > 0$, and an
algorithm R for A that runs in time $a(n)^{(1-\delta)}$ on inputs of length n , making q calls to an
oracle for B with query lengths n_1, \dots, n_q , where,

$$\sum_1^q (b(n_i))^{(1-\epsilon)} \leq (a(n))^{(1-\delta)}.$$

707 Problems that can be reduced to each other such that the lower bounds for each problem are
708 the same in both reductions, i.e., $A \underset{(a)}{\preceq}_{(b)} B$ and $B \underset{(b)}{\preceq}_{(a)} A$, are intuitively thought to have
709 the same underlying ‘reason’ for hardness, and are said to be fine-grained equivalent.

710 A reduction $A \underset{(a)}{\preceq}_{(b)} B$ would be interesting for B if $a(n)$ was a proven or well-believed
711 conjectured lower bound on A, thus implying a believable lower bound on B. One such
712 well-believed conjecture in complexity theory is SETH [18] for the classic CNF-SAT problem,
713 originally defined for deterministic algorithms, but now widely believed for randomized
714 algorithms as well.

715 ► **Hypothesis 1** (Strong Exponential Time Hypothesis (SETH)). *For every $\epsilon > 0$ there exists*
716 *an integer $k \geq 3$ such that CNF-SAT on formulas with clause size at most k and n variables*
717 *cannot be solved in $O(2^{(1-\epsilon)n})$ time even by a randomized algorithm.*

718 SETH implies a lower bound conjecture, denoted by OVH, on the Orthogonal Vectors problem
719 OV, as shown by a reduction from CNF-SAT to k-OV [41]. Thus, a conditional lower bound
720 under OVH implies one under SETH as well, leading to numerous conditional lower bound
721 results under OVH [See [42] for a detailed literature review]. This paper will also prove such
722 results on several data race detection problems, hence we now state k-OV and OVH formally.

723 An instance of k-OV is an integer $d = \omega(\log n)$ and k sets $A_i \subseteq \{0, 1\}^d$, $i \in [n]$ such that
724 $|A_i| = n$, and denoted by $OV(n, d)$.

725 ► **Problem 2** (Orthogonal Vectors (k-OV)). *Given an instance $OV(n, d, k)$, the k-OV problem*
726 *is to decide if there are k vectors $a_i \in A_i$ for all $i \in [n]$ such that the sum of their point wise*
727 *product is zero, i.e., $\sum_{j=1}^d \prod_{i=1}^k a_i[j] = 0$.*

728 For ease of exposition, we denote $OV(n, d, 2)$ and 2-OV by $OV(n, d)$ and OV respectively.

729 ► **Hypothesis 2** (Orthogonal Vectors Hypothesis (OVH)). *No randomized algorithm can solve*
730 *k-OV for an instance $OV(n, d, k)$ in time $O(n^{(k-\epsilon)} \cdot \text{poly}(d))$ for any constant $\epsilon > 0$.*

731 There is an impossibility result from [9] that proves that a reduction under SETH, and hence
732 under OVH, is not possible unless the following NSETH conjecture is false.

733 ► **Hypothesis 3** (Non-deterministic SETH (NSETH)). *For every $\epsilon > 0$, there exists a k so*
734 *that k-TAUT is not in $NTIME[2^{n(1-\epsilon)}]$, where k-TAUT is the language of all k-DNF formulas*
735 *which are tautologies.*

736 The impossibility result [9, Corollary 2] is as follows.

737 ► **Theorem 13.** *If NSETH holds and a problem $C \in NTIME[\mathbb{T}_C] \cap coNTIME[\mathbb{T}_C]$, then for*
 738 *any problem B that is SETH-hard under deterministic reductions with time T_B , and $\gamma > 0$,*
 739 *we cannot have a fine-grained reduction $B \underset{(\mathbb{T}_B)}{\preceq_{(c)}} C$ where $c = T_C^{(1+\gamma)}$.*

740 We show some of our problems satisfy the conditions of Theorem 13, and hence show lower
 741 bounds for these conditioned on one of two other hypotheses called HSH and FOPH($\forall\exists\exists$),
 742 described below.

743 An instance of the hitting set problem, denoted by HS, is an integer $d = \omega(\log n)$ and sets
 744 $X, Y \subseteq \{0, 1\}^d$, $i \in [n]$ such that $|X| = |Y| = n$, and denoted by HS(n, d).

745 ► **Problem 3 (Hitting Sets (HS)).** *Given an instance HS(n, d), the HS problem is to decide if*
 746 *there is a vector $x \in X$ such that for all $y \in Y$ we have $x \cdot y \neq 0$, or informally, some vector*
 747 *in X hits all vectors in Y .*

748 ► **Hypothesis 4 (Hitting Sets Hypothesis (HSH)).** *No randomized algorithm can solve HS for*
 749 *an instance HS(n, d) in time $O(n^{(2-\epsilon)} \cdot \text{poly}(d))$ for any constant $\epsilon > 0$.*

750 HSH implies OVH, but the reverse direction is not known.

751 Finally we consider a subclass of first order formula over structures of size n and with m
 752 relational tuples [17].

753 ► **Problem 4 (FO($\forall\exists\exists$)).** *Decide if a given a first-order formula quantified by $\forall\exists\exists$ is has a*
 754 *model on a structure of size n with m relational tuples.*

755 It is known that FO($\forall\exists\exists$) can be solved in $O(m^{3/2})$ time using ideas from triangle detection
 756 algorithms [17]. For dense structures ($m = \Theta(n^2)$), this yields the bound $O(n^3)$. Although
 757 sub-cubic algorithms might be possible, achieving a truly quadratic bound seems unlikely or
 758 at least highly non-trivial.

759 B Proofs of Section 3

760 B.1 Proofs from Section 3.1

761 ► **Lemma 10.** *Let $e_1 \leq_{tr}^\sigma e_2$ be events in σ such that $\text{tid}(e_1) \neq \text{tid}(e_2)$. We have, $e_1 \leq_{HB}^\sigma$
 762 $e_2 \iff \neg(\text{AcqLS}_{e_2}^\sigma \sqsubseteq \text{RelLS}_{e_1}^\sigma)$*

763 **Proof.** (\Rightarrow) Let $e_1 \leq_{HB}^\sigma e_2$. Using the definition of \leq_{HB}^σ , there must be a sequence of events
 764 $f_1, f_2 \dots f_k$ with $k > 1$, $f_1 = e_1$, $f_k = e_2$, and for every $1 \leq i < k$, $f_i \leq_{tr}^\sigma f_{i+1}$ and either
 765 $f_i \leq_{TO}^\sigma f_{i+1}$ or there is a lock ℓ , such that $f_i \in \text{Releases}_\sigma(\ell)$ and $f_{i+1} \in \text{Acquires}_\sigma(\ell)$. Let j
 766 be the smallest index i such that $\text{tid}(f_i) \neq \text{tid}(f_{i+1})$; such an index exists as $\text{tid}(e_1) \neq \text{tid}(e_2)$.
 767 Observe that there must be a lock ℓ for which $\text{op}(f_j) = \text{rel}(\ell)$ and $\text{op}(f_{j+1}) = \text{acq}(\ell)$.
 768 Observe that $\text{pos}_\sigma(f_j) < \text{pos}_\sigma(f_{j+1})$, $\text{RelLS}_{e_1}^\sigma(\ell) \leq \text{pos}_\sigma(f_j)$ and $\text{pos}_\sigma(f_{j+1}) \leq \text{AcqLS}_{e_2}^\sigma$,
 769 giving us $\text{RelLS}_{e_1}^\sigma(\ell) < \text{AcqLS}_{e_2}^\sigma(\ell)$.

770 (\Leftarrow) Let ℓ be a lock such that $\text{RelLS}_{e_1}^\sigma(\ell) < \text{AcqLS}_{e_2}^\sigma(\ell)$. Then, there is a release event f
 771 and an acquire event g on lock ℓ such that $\text{pos}_\sigma(f) < \text{pos}_\sigma(g)$, $e_1 \leq_{HB}^\sigma f$ and $g \leq_{HB}^\sigma e_2$. This
 772 means $f \leq_{HB}^\sigma g$ and thus $e_1 \leq_{HB}^\sigma e_2$. ◀

773 For the sake of completeness, we present the computation of release lockstamps. As with
 774 Algorithm 1, we maintain the following variables. For each thread t and lock ℓ , we will
 775 maintain variables \mathbb{C}_t and \mathbb{L}_ℓ that take values from the space of all lockstamps. We also
 776 additionally maintain an integer variable \mathbf{p}_ℓ for each lock ℓ that stores the index (or relative
 777 position) of the earliest (according to the trace order \leq_{tr}^σ) release event of lock ℓ in the trace.

23:22 Dynamic Data-Race Detection through the Fine-Grained Lens

778 Initially, we set each \mathbb{C}_t and \mathbb{L}_m to $\lambda \ell \cdot \infty$, for each thread t and lock m . Further, for each
 779 lock m , we set \mathbf{p}_m to $n_m + 1$, where n_m is the number of release events of m in the trace;
 780 this can be obtained in a linear scan (or by reading the value of \mathbf{p}_m at the end of a run of
 781 Algorithm 1). We traverse the events according to the total trace order and perform updates
 782 to the data structures as described in Algorithm 4, by invoking the appropriate *handler*
 783 based on the thread and operation of the event $e = \langle t, op \rangle$ being visited. At the end of each
 784 handler, we assign the lockstamp RelLS_e^σ to the event e .

■ **Algorithm 4** *Assigning release lockstamps to events in the trace*

| | | |
|---|---|---|
| 1 acquire (t, ℓ): 2 $\mathbb{L}_\ell \leftarrow \mathbb{C}_t$ 3 $\text{RelLS}_e^\sigma \leftarrow \mathbb{C}_t$ | 4 release (t, ℓ): 5 $\mathbf{p}_\ell \leftarrow \mathbf{p}_\ell - 1$ 6 $\mathbb{C}_t \leftarrow \mathbb{C}_t[\ell \mapsto \mathbf{p}_\ell] \sqcap \mathbb{L}_\ell$ 7 $\text{RelLS}_e^\sigma \leftarrow \mathbb{C}_t$ | 8 read (t, x): 9 $\text{RelLS}_e^\sigma \leftarrow \mathbb{C}_t$ 10 write (t, x): 11 $\text{RelLS}_e^\sigma \leftarrow \mathbb{C}_t$ |
|---|---|---|

785 Let us now state the correctness of Algorithm 1 and Algorithm 4.

786 ► **Lemma 14.** *On input trace σ , Algorithm 1 and Algorithm 4 correctly compute the lock-*
 787 *stamps AcqLS_e^σ and RelLS_e^σ respectively for each event $e \in \text{Events}_\sigma$.*

788 **Proof Sketch.** We focus on the correctness proof of Algorithm 1; the proof for Algorithm 4
 789 is similar. The proof relies on the invariant maintained by Algorithm 1 the variables \mathbb{C}_t , \mathbb{L}_ℓ
 790 and \mathbf{p}_ℓ for each thread t and lock ℓ , which we state next. Let π be the prefix of the trace
 791 processed at any point in the algorithm. Let C_t^π , L_ℓ^π and p_ℓ^π be the values of the variables
 792 \mathbb{C}_t , \mathbb{L}_ℓ and \mathbf{p}_ℓ after processing the prefix π . Then, the following invariants are true:

- 793 ■ $C_t^\pi = \text{AcqLS}_{e_t^\pi}^\pi = \text{AcqLS}_{e_t^\pi}^\sigma$, where e_t^π is the last event in π performed by thread t
- 794 ■ $L_\ell^\pi = \text{AcqLS}_{e_\ell^\pi}^\pi = \text{AcqLS}_{e_\ell^\pi}^\sigma$, where e_ℓ^π is the last acquire event on lock ℓ in π .
- 795 ■ $p_\ell^\pi = \text{pos}_\ell^\pi(e_\ell^\pi)$, where e_ℓ^π is the last acquire event on lock ℓ in π .

796 These invariants can be proved using a straightforward induction, each time noting the
 797 definition of \leq_{HB}^σ . ◀

798 ► **Lemma 15.** *For a trace with \mathcal{N} events and \mathcal{L} locks, Algorithm 1 and Algorithm 4 both*
 799 *take $O(\mathcal{T} \cdot \mathcal{L})$ time.*

800 **Proof.** We focus on Algorithm 1; the analysis for Algorithm 4 is similar. At each acquire
 801 event, the algorithm spends $O(1)$ time for updating \mathbf{p}_ℓ , $O(\mathcal{L})$ time for doing the \sqcap operation,
 802 and $O(\mathcal{L})$ time for the copy operation ($\text{AcqLS}_e^\sigma \leftarrow \mathbb{C}_t$). For a release event, we spend $O(\mathcal{L})$
 803 for the two copy operations. At read and write events, we spend $O(\mathcal{L})$ for copy operations.
 804 This gives a total time of $O(\mathcal{N} \cdot \mathcal{L})$. ◀

805 ► **Lemma 11.** *A trace σ has an HB-race iff there is pair of consecutive conflicting events in*
 806 *σ that is an HB-race. Moreover, σ has $O(\mathcal{N})$ many consecutive conflicting pairs of events.*

807 **Proof.** We first prove that if there is an HB-race in σ , then there is a pair of consecutive
 808 conflicting events that is in HB-race. Consider the first HB-race, i.e., an HB-race (e_1, e_2)
 809 such that for every other HB-race (e'_1, e'_2) , either $e_2 \leq_{\text{tr}}^\sigma e'_2$ or $e_2 = e'_2$ and $e'_1 \leq_{\text{tr}}^\sigma e_1$. We
 810 remark that such a race (e_1, e_2) exists if σ has any HB-race. We now show that (e_1, e_2)
 811 are a consecutive conflicting pair (on variable x). Assume on the contrary that there is an
 812 event $f \in \text{Writes}_\sigma(x)$ such that $e_1 <_{\text{tr}}^\sigma f <_{\text{tr}}^\sigma e_2$. If either (e_1, f) or (f, e_2) is an HB-race, then
 813 this contradicts our assumption that (e_1, e_2) is the first HB-race in σ . Thus, $e_1 \leq_{\text{HB}}^\sigma f$ and
 814 $f \leq_{\text{HB}}^\sigma e_2$, which gives $e_1 \leq_{\text{HB}}^\sigma e_2$, another contradiction.

815 We now turn our attention to the number of consecutive conflicting events in σ . For every
 816 read or write event e_2 , there is at most one write event e_1 such that (e_1, e_2) is a consecutive
 817 conflicting pair (namely the latest conflicting write event before e_2) Further, for every read
 818 event e_1 , there is at most one write event e_2 such that (e_1, e_2) is a consecutive conflicting
 819 pair (namely the earliest conflicting write event after e_1). This gives at most $2\mathcal{N}$ consecutive
 820 conflicting pairs of events. ◀

821 Let us now state the correctness of Algorithm 2.

822 ▶ **Lemma 16.** *For a trace σ , Algorithm 2 reports a race iff σ has an HB-race.*

823 **Proof Sketch.** The proof relies on the following straightforward invariants; we skip their
 824 proofs as they are straightforward. In the following, e_x^π is the last event with $\text{op}(e_x^\pi) = \mathbf{w}(x)$
 825 in a trace π .

- 826 ■ After processing the prefix π of σ , $t_x^w = \text{tid}(e_x^\pi)$ and $\mathbb{W}_x = e_x^\pi$.
- 827 ■ After processing the prefix π of σ , the set \mathbb{S}_x is $\{(t, L) \mid \exists e \in \text{Reads}_\pi(x), e_x^\pi \leq_{\text{tr}}^\pi e, \text{tid}(e) =$
 828 $t, \text{RelLS}_e^\sigma = L\}$.

829 The rest of the proof follows from Lemma 15 and Lemma 14. ◀

830 Let us now characterize the time complexity of Algorithm 2.

831 ▶ **Lemma 17.** *On an input trace with \mathcal{N} events and \mathcal{L} locks, Algorithm 2 runs in time*
 832 *$O(\mathcal{N} \cdot \mathcal{L})$.*

833 **Proof Sketch.** Each pair (t, L) of thread identifier and lockstamp is added atmost once in
 834 some set \mathbb{S}_x (for some x). Also, each such pair is also compared against another timestamp
 835 atmost once. Each comparison of timestamps take $O(\mathcal{L})$ time. This gives a total time of
 836 $O(\mathcal{N} \cdot \mathcal{L})$. ◀

837 ▶ **Theorem 4.** *Deciding whether σ has an HB race can be done in time $O(\mathcal{N} \cdot \min(\mathcal{T}, \mathcal{L}))$.*

838 **Proof.** We focus on proving that there is an $O(\mathcal{N} \cdot \mathcal{L})$ time algorithm, as the standard vector-
 839 clock algorithm [19] for checking for an HB-race runs in $O(\mathcal{N} \cdot \mathcal{T})$ time. Our algorithm's
 840 correctness is stated in Lemma 16 and its total running time is $O(\mathcal{N} \cdot \mathcal{L})$ (Lemma 17 and
 841 Lemma 15). ◀

842 B.2 Proofs from Section 3.2

843 ▶ **Theorem 1.** *For any $\epsilon > 0$, there is no algorithm that detects even a single HB race that*
 844 *involves a read in time $O(\mathcal{N}^{2-\epsilon})$, unless the OV hypothesis fails.*

845 **Proof.** Consider a pair of events $\mathbf{w}(z)$ from the d threads $t(x, i), i \in [d]$, and $\mathbf{r}(z) \in t_y$ for
 846 some x, i, y . We have $\mathbf{w}(z) \leq_{\text{HB}}^\sigma \mathbf{r}(z)$ iff there is some path from $\mathbf{w}(z)$ to $\mathbf{r}(z)$ in $\mathbb{G}(\leq_{\text{HB}}^\sigma)$. As
 847 $\mathbf{w}(z)$ and $\mathbf{r}(z)$ are in different threads, such a path can only be through lock events in a
 848 sequence of threads such that the first and last threads are $t(x, i)$ for some $i \in [d]$ and t_y , and
 849 every consecutive pair of threads in the sequence holds a common lock. Now all the locks in
 850 t_y are $l(y, i)$ for all i where $y[i] = 1$. Consider the lock corresponding to any $i \in [d]$. The
 851 only thread $t(x', i)$ that also holds this lock corresponds to the last x' such that $x'[i] = 1$.
 852 The only other lock held by $t(x', i)$ is l_i . If $\mathbf{w}(z)$ is in $t(x', i)$, we are done. Otherwise the
 853 only common lock between these threads $t(x', i)$ and those of $\mathbf{w}(z)$ can be one of the l_i . The
 854 threads of $\mathbf{w}(z)$ contain all l_i where $x[i] = 1$. Hence, for there to be a common lock between

855 these threads, there must be at least one i such that $x'[i] = 1$ and $x[i] = 1$. As this thread
856 also has the lock $l(y, i)$, $y[i]$ is also 1.

857 Thus, there is a path from $w(z)$ to $r(z)$ if and only if there is at least one $i \in [d]$ such that
858 $x[i] = y[i] = 1$, hence x and y are not orthogonal. A pair of orthogonal vectors of **OV** thus
859 corresponds to a write-read **HB**-race in the reduced trace.

860 Finally we turn our attention to the complexity. In time $O(n \cdot d)$, we have reduced an **OV**
861 instance to determining whether there is a write-read **HB** race in a trace of $\mathcal{N} = O(nd)$ events.
862 If there was a sub-quadratic i.e. $O((n \cdot d)^{(2-\epsilon)}) = n^{(2-\epsilon)} \cdot \text{poly}(d)$ algorithm for detecting a
863 write-read **HB** race, then this would also solve **OV** in $n^{(2-\epsilon)} \cdot \text{poly}(d)$ time, refuting the **OV**
864 hypothesis. ◀

865 ▶ **Lemma 12.** $FO(\forall\exists\exists)$ reduces to **MCONN** on a graph G with $O(n)$ nodes in $O(n^2)$ time.

866 **Proof.** For intuition, assume the first order property is on an undirected graph with n
867 variables and m edges. Let the property be specified in quantified 3-DNF form with a
868 constant number of predicates, i.e., $\phi = \forall x \exists y \exists z (\psi_1 \vee \psi_2 \vee \dots \vee \psi_k)$, where x, y, z represent
869 nodes of the graph, and each ψ_i is a conjunction of 3 variables representing edges of the
870 graph, for example $e(x, y) \wedge \neg e(y, z) \wedge e(x, z)$. The property is then true if and only if some
871 predicate is satisfied, which is true if all of its variables are satisfied ($e(x, y)$ is satisfied when
872 edge (x, y) is in the graph). Denote the graph on which ϕ is defined by $H(I, J)$, where I and
873 J are respectively the sets of nodes and edges of H .

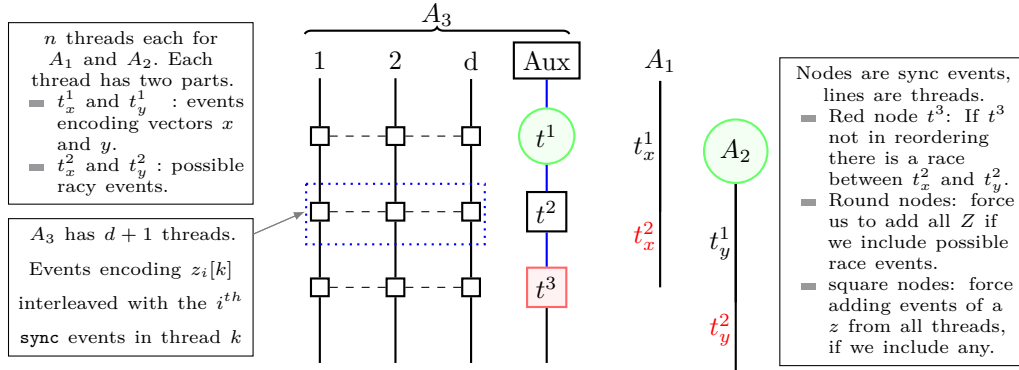
874 The instance of **MCONN** is constructed given H and ϕ as follows. Construct a $(2k+2)$ -partite
875 graph $G(V, E)$ by first creating $2k+2$ copies of I . Denote these copies by S, Y_i, Z_i, T , $i \in [k]$,
876 and the copy of each node $x \in I$ in any part, say S , by $x(S)$. $\psi_i = (e_1 \wedge e_2 \wedge e_3)$ is encoded by
877 connecting the sets (S, Y_i) to represent e_1 , (Y_i, Z_i) for e_2 and (Z_i, T) for e_3 as follows. If e_i is of
878 the form $e(x, y)$ (and not its negation), then draw a copy of H between its corresponding sets,
879 say S and Y_i without loss of generality. That is, for every x, y , $(x, y) \in J \Leftrightarrow (x(S), y(Y_i)) \in E$.
880 If on the other hand e_i is of the form $\neg e(x, y)$ then connect a copy of the complement of H ,
881 i.e., $(x, y) \notin J \Leftrightarrow (x(S), y(Y_i)) \in E$.

882 Finally define $|I|$ pairs $(x(S), x(T))$ as the (s, t) pairs for **MCONN**.

883 We now prove this reduction is correct. First, assume ϕ is true. Then for every node x , there
884 exist nodes y, z such that some predicate is true. If ψ_i is the predicate that is satisfied for
885 some node u , then there is a path between $u(S)$ and $u(T)$ through the parts S, Y_i, Z_i and
886 T as follows. As the first variable is satisfied, then if it is $e(x, y)$, then $(x, y) \in J$, and $x(S)$
887 is connected to $y(Y_i)$, and if it is $\neg e(x, y)$, then $(x, y) \notin J$ and again $x(S)$ is connected to
888 $y(Y_i)$. Similarly, $y(Y_i)$ is connected to $z(Z_i)$, and $z(Z_i)$ to $x(T)$. These edges form a 3 length
889 path between $x(S)$ and $x(T)$.

890 Now consider the reverse case, and assume the **MCONN** problem is true, that is, there is
891 a path between every $(x(S), x(T))$ pair. Note that the construction of edges in G is such
892 that any path from $x(S)$ to $x(T)$ has to be a 3 length path, connecting the copy of x in S
893 to its copy in some Y_i , from this Y_i to its corresponding Z_i , and from Z_i to T . Also, this
894 path exists only if all variables of the corresponding ψ_i are true. Hence, as there is a path
895 between every pair $(x(S), x(T))$, and one pair is defined for every variable x , some predicate
896 is satisfied for every x . Thus ϕ is also true.

897 Finally, the time of the reduction is equal to the size of G . This is $2k+2 = O(1)$ graphs, each
898 of which is either H or its complement. Hence $|G| = O(m + n + (n^2 - m) + n) = O(n^2)$. ◀



■ **Figure 6** Intuition for the reduction from 3-OV to sync-preserving race detection. Thread *Aux* allows forming a trace such that 3-OV has a solution iff there is a sync-preserving race.

899 ► **Theorem 3.** For any $\epsilon > 0$, if there is an algorithm for detecting any HB race in time
 900 $O(\mathcal{N}^{1+\epsilon})$, then there is an algorithm for FO($\forall\exists\exists$) formulas in time $O(m^{1+\epsilon})$.

901 **Proof.** We first reduce the instance of FO($\forall\exists\exists$) to MCONN as in the proof of Lemma 12.
 902 Let $G(V, E)$ be the multi-partite graph for MCONN and S, T the first and last parts of nodes
 903 of G . We add a sufficient number of nodes, referred as dummy nodes, to make G sparse. Let
 904 every node x of $V \setminus T$ correspond to a distinct thread t_x and form one write access event to a
 905 distinct variable v_x in the thread. Let each node t in T also correspond to a write access
 906 event of the variable corresponding to the copy of t in S , and be in a new thread. Define
 907 $|E|$ locks, and for every edge $(a, b) \in E$, let the events corresponding to v_a and v_b hold the
 908 lock $l_{(a,b)}$ corresponding to (a, b) . The trace σ for first lists all threads corresponding to the
 909 dummy nodes in some fixed arbitrary order, then the threads corresponding to nodes in S ,
 910 followed by those in each Y_i , followed by those in each Z_i , in a fixed arbitrary order, and
 911 finally those in T .

912 This reduction is seen to be correct by observing that G was modified to be the transitive
 913 reduction graph of σ , and the only HB-race events can be the pairs of write events corres-
 914 ponding to the pairs of nodes given as input to MCONN. Thus, each pair of events does not
 915 form an HB-race if and only if G has a path between its corresponding pair of nodes.

916 To analyze the time of the reduction, first we see that the size of σ is the size of G , with
 917 dummy nodes added to have $n = O(n^2)$, and hence $O(n^2)$. There are $O(n^2)$ variables,
 918 locks and threads in σ . If deciding if the given trace has an HB-race has an $O((n^2)^{1+\epsilon})$
 919 time algorithm, then FO($\forall\exists\exists$) can be solved in $O(n^{2+\epsilon'})$ time, which is $O(m^{1+\epsilon'})$ time for
 920 properties on dense structures. ◀

921 C Proofs of Section 4

922 ► **Theorem 5.** For any $\epsilon > 0$, there is no algorithm that detects even a single sync-preserving
 923 race in time $O(\mathcal{N}^{3-\epsilon})$, unless the 3-OV hypothesis fails. Moreover, the statement holds even
 924 for traces over a single variable.

925 **Proof.** Consider any sync-preserving correct reordering ρ of σ that exposes a data race
 926 $(w(z), r(z))$ on the local traces t_x and t_y . The following statements are straightforward to
 927 verify based on the definition of sync-preserving correct reorderings.

928 1. For every $k \in [d]$, the first $\text{sync}(\ell_k)$ event the trace t_k is also in ρ .

929 2. The auxiliary trace t cannot have an open critical section in ρ . This implies that for every
 930 trace t_k with $k \in [d]$, the last event of t_k in ρ cannot be its i^{th} $\text{sync}(\ell_k)$ event, where i is
 931 even. Moreover, the number of $\text{sync}(\ell_k)$ events in ρ is the same for every trace t_k with
 932 $k \in [d]$.

933 First, consider that the 3-OV instance has a solution, i.e., there exist $x \in A_1$, $y \in A_2$ and
 934 $z \in A_3$ such that x, y, z are orthogonal, and we argue that σ has a data race that is also
 935 sync-preserving. We construct a sync-preserving correct reordering ρ of σ that exposes
 936 the data race. We only specify the local traces that exist in ρ , as their interleaving that
 937 constructs ρ will be identical to the one in σ (in other words, we only specify the prefix up
 938 to which every local trace of σ is executed in ρ). We execute the traces t_x and t_y all the
 939 way before the corresponding $w(z)$ and $r(z)$ events (hence we are exposing a race between
 940 these two events). For every $k \in [d]$ if $z[k] = 0$ or $x[k] = 0$, we execute t_k up to the $(2 \cdot i)^{\text{th}}$
 941 $\text{sync}(\ell_k)$ event, where i is such that z is the i^{th} vector of A_3 . On the other hand, if $y[k] = 0$,
 942 we execute t_k up to the first $\text{rel}(\ell_k)$ event that appears after the $(2 \cdot i)^{\text{th}}$ $\text{sync}(\ell_k)$ event in
 943 t^k . Finally, we execute t^k until its $(i - 1)^{\text{th}}$ $\text{rel}(X)$ event.

944 It is easy to verify that ρ is a valid correct reordering. Indeed, we have two open critical
 945 sections in the threads t_x and t_y , on the locks X and Y respectively. Moreover, for every
 946 $k \in [d]$, we have the following.

- 947 1. If $z[k] = 0$, there are no other open critical sections.
- 948 2. If $z[k] = 1$ and $x[k] = 0$, there is one open critical section in the thread $z[k]$ on lock l_k .
- 949 3. If $z[k] = x[k] = 1$ and $y[k] = 0$, there is one open critical section in the thread $z[k]$ on
 950 lock l'_k .

951 We now consider the opposite direction, i.e., assume that there is a sync-preserving race in σ ,
 952 and we argue that there exist $x \in A_1$, $y \in A_2$ and $z \in A_3$ such that x, y, z are orthogonal.
 953 Consider any sync-preserving correct reordering ρ that exposes a race on the access events
 954 of two local traces t_x and t_y . Because of Item 1 above, every trace t_k is at least partially
 955 present in ρ . Because of Item 1 above, every such trace executes the same number of $\text{sync}(\ell_k)$
 956 events in ρ , and this number is odd. We argue that the triplet x, y, z is orthogonal, where z
 957 is the i^{th} vector of A_3 such that each t_k executes $2 \cdot (i - 1) + 1$ $\text{sync}(\ell_k)$ events in ρ . Indeed,
 958 consider any $k \in [d]$ and assume that $x[k] = y[k] = 1$. If $z[k] = 1$, then we have a $\text{acq}(\ell)$
 959 event in t_k that immediately precedes its last $\text{sync}(\ell_k)$ event. Since $x[k] = 1$, the trace t_x
 960 also has an $\text{acq}(\ell_k)$ event. Since the $\text{acq}(\ell_k)$ event of t_x is after the $\text{acq}(\ell_k)$ event of t_k in
 961 σ , the matching $\text{rel}(\ell_k)$ of t_k must also be in ρ . This implies that the $\text{acq}(l'_k)$ event of t_k
 962 that immediately succeeds its last $\text{rel}(\ell_k)$ event is also in ρ . Since $y[k] = 1$, the trace t_y
 963 also has an $\text{acq}(l'_k)$ event. Since the $\text{acq}(l'_k)$ event of t_y is after the $\text{acq}(l'_k)$ event of t_k in
 964 σ , the matching $\text{rel}(l'_k)$ of t_k must also be in ρ . However, we now have another $\text{sync}(\ell_k)$
 965 event of t_k in ρ , in particular, the $\text{sync}(\ell_k)$ event that immediately precedes its last $\text{rel}(l'_k)$
 966 event. But this results in an even number of $\text{sync}(\ell_k)$ events of t_x being present in ρ , which
 967 contradicts our observation in Item 2. Thus, if $x[k] = y[k] = 1$, we necessarily have that
 968 $z[k] = 0$, and the triplet x, y, z is orthogonal.

969 The desired result follows. ◀

970 **D** Proofs of Section 5

971 **► Theorem 6.** *For any $\epsilon > 0$, any $\mathcal{T} \geq 2$ and any $\mathcal{L} = \omega(\log \mathcal{N})$, there is no algorithm that
 972 detects even a single lock-cover race in time $O(\mathcal{N}^{2-\epsilon})$, unless the OV hypothesis fails.*

973 **Proof of Theorem 6.** To see why the reduction is correct, observe that if there a solution
 974 to OV, that is, a pair of vectors x, y such that for all $k \in [d]$, $x[k] = 0$ or $y[k] = 0$, implies
 975 that the corresponding events in σ , say e_x and e_y , have for each lock $k \in [d]$ either e_x does
 976 not hold the lock or e_y does not. As they have distinct thread ids too, e_x and e_y form
 977 two conflicting events with $\text{locksHeld}_\sigma(e_1) \cap \text{locksHeld}_\sigma(e_2) = \emptyset$. Similarly, a lock-cover race
 978 implies that for every lock, one of the events in race do not hold the lock, hence have their
 979 corresponding coordinate in OV 0. The events are thus orthogonal to each other.

980 Regarding the complexity, we have used $O(n \cdot d)$ time to construct a trace σ with $\mathcal{N} = O(n \cdot d)$
 981 events. If we can detect a lock-cover race in σ in $O(\mathcal{N}^{2-\epsilon})$ time, then OV can be solved in
 982 $O(n^{2-\epsilon} \cdot \text{poly}(d))$ time, contradicting the OV hypothesis. ◀

983 ▶ **Theorem 9.** *Deciding whether a trace σ has a lock-set race on a variable x can be*
 984 *performed in $O(\mathcal{N})$ time. Thus, deciding whether σ has a lock-set race can be performed in*
 985 *$O(\mathcal{N} \cdot \min(\mathcal{L}, \mathcal{V}))$ time.*

986 **Proof.** We first argue that the algorithm maintains the invariant stated in Equation (2).
 987 The invariant for A is trivial to verify. Moreover, it is easy to see that, assuming that the
 988 invariant holds before processing an $\text{acq}(\ell)$ or $\text{rel}(\ell)$ event, it also holds after processing that
 989 event. Indeed, for an event $\text{acq}(\ell)$, we have $\ell \in A$, and to maintain $C = \overline{A} \cap B$, we remove ℓ
 990 from C if $\ell \in B$. Similarly for an event $\text{rel}(\ell)$. To see that the invariant is maintained after
 991 processing an access event $\text{w}(x)/\text{r}(x)$, note that we have

$$992 \quad B \setminus C = B \cap \overline{C} = B \cap (\overline{B \cap A}) = B \cap (\overline{B} \cup A) = B \cap A$$

993 and thus updating $B \leftarrow B \setminus C$ yields

$$994 \quad \text{Locks}_\sigma \cap \bigcap_{\substack{e' \in \text{Accesses}_\sigma(x) \\ e' <_{tr}^\sigma e}} \text{locksHeld}_\sigma(e') \cap \text{locksHeld}_{tr}(e) = \text{Locks}_\sigma \cap \bigcap_{\substack{e' \in \text{Accesses}_\sigma(x) \\ e' \leq_{tr}^\sigma e}} \text{locksHeld}_\sigma(e')$$

995 Finally, at this point we have $\overline{A} \cap B = \overline{A} \cap B \cap A = \emptyset$, thus the invariant also holds for C .

996 We now turn our attention to complexity. Using a bit-set representation of the sets A , B
 997 and C , it is clear that each of the operations except $\text{w}(x)/\text{r}(x)$ take constant time per event.
 998 Each $\text{w}(x)/\text{r}(x)$ operation takes $O(|C|)$ time. Note, however, that because of the previous
 999 invariant, every lock is removed from B at most once, hence the total time for performing all
 1000 set differences $B \leftarrow B \setminus C$ is $O(\mathcal{N} + \mathcal{L}) = O(\mathcal{N})$. Thus the total time is $O(\mathcal{N})$. The desired
 1001 result follows. ◀

1002 ▶ **Theorem 8.** *For any $\epsilon > 0$ and any $\mathcal{T} = \omega(\log n)$, there is no algorithm that detects even*
 1003 *a single lock-cover race in time $O(\mathcal{N}^{2-\epsilon})$, unless the HS hypothesis fails.*

1004 **Proof.** First, assume there is a solution to HS, i.e., $\exists x_k \in X \forall y_i \in Y \exists j \in [d] x_k[j] = y_i[j] = 1$.
 1005 Then for the variable z_k , for every lock ℓ_i , there is a thread t_j that contains $\text{w}(z_k)$ (as $x_k[j] = 1$)
 1006 but does not contain lock ℓ_i (as $y_i[j] = 1$). Thus $\bigcap_{e \in \text{Accesses}_\sigma(z_k)} \text{locksHeld}_\sigma(e) = \emptyset$, and we
 1007 have a lock-set race on variable z_k as there are at least two $\text{w}(z_k)$ conflicting events, one in
 1008 the thread t_j and the other in thread t_0 . For the opposite direction, assume that HS does
 1009 not have a solution, i.e., $\forall x_k \in X \exists y_i \in Y \forall j \in [d] (x_k[j] = 0 \text{ or } y_i[j] = 0)$. Then for each
 1010 variable z_k , there is some lock ℓ_i such that every thread that contains a write event $\text{w}(z_k)$
 1011 (thus $x_k[j] = 1$) also contains the lock ℓ_i (as necessarily $y_i[j] = 0$). Hence, for every variable
 1012 z_k , some lock ℓ_i is held by all its access events. Thus σ does not have a lock-set race.

23:28 Dynamic Data-Race Detection through the Fine-Grained Lens

1013 Regarding the complexity, we have created a trace σ with $\mathcal{N} = O(n \cdot d)$ events in $O(n \cdot d)$ time.
1014 Thus, any $O(\mathcal{N}^{(2-\epsilon)})$ time algorithm for HS implies an $O(n^{(2-\epsilon)} \cdot \text{poly}(d))$ time algorithm
1015 for HS, contradicting the HS hypothesis. ◀