Efficient Data Race Detection of Async-Finish Programs Using Vector Clocks

Shivam Kumar shivamkm07@gmail.com Indian Institute of Technology Kanpur India Anupam Agrawal anupamag@cse.iitk.ac.in Indian Institute of Technology Kanpur India Swarnendu Biswas swarnendu@cse.iitk.ac.in Indian Institute of Technology Kanpur India

Abstract

Existing data race detectors for task-based programs incur significant run time and space overheads. The overheads arise because of frequent lookups in fine-grained tree data structures to check whether two accesses can happen in parallel.

This work shows how to efficiently apply vector clocks for dynamic race detection of async-finish programs with locks. Our proposed technique, *FastRacer*, builds on the FastTrack algorithm with per-task and per-variable optimizations to reduce the size of vector clocks. FastRacer exploits the structured parallelism of async-finish programs to use a coarsegrained encoding of the dynamic task inheritance relations to limit the metadata in the presence of many concurrent readers. Our evaluation shows that FastRacer improves time and space overheads over FastTrack and is competitive with the state-of-the-art race detectors for async-finish programs.

CCS Concepts: • Software and its engineering \rightarrow Software testing and debugging; Runtime environments; Multiprocessing / multiprogramming / multitasking.

Keywords: Concurrency bugs, data races, happens-before, dynamic program analysis, task parallelism, async-finish

ACM Reference Format:

Shivam Kumar, Anupam Agrawal, and Swarnendu Biswas. 2022. Efficient Data Race Detection of Async-Finish Programs Using Vector Clocks . In *The 13th International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM'22), April 2–6, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3528425.3529101

1 Introduction

The task-based programming abstraction helps write efficient and portable parallel code without having to think of low-level

PMAM'22, April 2–6, 2022, Seoul, Republic of Korea © 2022 Association for Computing Machinery. ACM ISBN 978-1-4503-9339-3/22/04...\$15.00 https://doi.org/10.1145/3528425.3529101 threads. Tasks execute in parallel as hardware-agnostic logical units of work, and programmers only specify the dependencies among the tasks. An accompanying runtime schedules tasks to threads and provides performance features like workstealing. Cilk [5, 13], X10 [8], Habanero-Java [7], and Java Fork-Join [17] are popular task-based frameworks.

Task-based programs are susceptible to concurrency errors such as atomicity violations and data races [11, 30, 31]. A data race occurs when two accesses, with at least one write, from different tasks are incorrectly synchronized. Data races in shared-memory programs often indicate the presence of other concurrency errors, and can affect an execution by crashing or corrupting data. Data races are hard to detect and fix since they may occur nondeterministically under specific thread interleavings, program inputs, and execution environments. Data races have led to several real-world disasters [18, 27, 29]; such high-profile failures are a testament that data races are present even in well-tested code.

The problem. There exists prior work to detect data races in task-based programs [1, 9, 11, 22, 30, 31, 37, 39]. Most analysis utilize the series-parallel structure of execution of task-based programs to check whether accesses can potentially execute in parallel (called *series-parallel maintenance*) [1, 11, 31, 37, 39]. Prior techniques are either serial (e.g., [9, 11, 30]) or are difficult to parallelize (e.g., [1]), detect races only in a given schedule (e.g., [12]), continue to incur high runtime overheads (e.g., [22, 31, 39]), require tight coupling with the runtime scheduler for good performance (e.g., [37]), or do not support lock-based synchronization (e.g., [1, 31, 37]).

Our approach. We focus on efficient detection of *per-input* apparent data races in task-based programs with async-finish semantics.¹ For a given application and an input, per-input races include races observed in the current schedule as well as other schedules with possibly permuted memory operations, ignoring schedule-sensitive branches [39]. While prior work has *ignored* how to optimize vector clocks for efficient race detection of task-based programs, we *argue* that analyses based on vector clocks are generic, inherently parallel, and have better data locality than tree-based data structures

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

¹Apparent data races occur because of the usage of parallel task constructs and ignore the per-schedule dynamic interleavings [25]. Feasible data races consider the nondeterministic timing variations during execution.

for series-parallel (SP) maintenance. However, a naïve application of vector-clock-based analysis to tasks results in prohibitive memory and run-time overhead [31, 39].

This paper presents *FastRacer*, a vector-clock-based race detector for async-finish programs. FastRacer avoids the redundancy in *per-task* vector clocks by using auxiliary data structures to maintain space- and time-efficient lossless clock representations correctly. FastRacer exploits the structured parallelism in async-finish programs to optimize the space requirement of *per-variable* metadata in the presence of many concurrent readers. Prior work has shown that a careful selection of only two "concurrent read" accesses is sufficient for detecting read-write data races for async-finish programs [31, 39]. FastRacer uses coarse-grained encoding of dynamic task inheritance relationships to identify the two accesses (for both reads and writes) necessary for race detection, and uses vector clocks to check whether two accesses can race.

We evaluate the performance and correctness of FastRacer on C++ applications that use Intel TBB for task parallelism and compare with prior work [12, 37, 39]. Our evaluation shows that the run time and memory overhead of FastRacer is substantially lower compared to prior race detectors.

Contributions. This paper makes the following contributions:

- To the best of our knowledge, this work is the *first* to show the viability of using vector clocks for efficient dynamic analysis of task-based programs;
- a race detector called FastRacer that detects per-input apparent races in async-finish programs with locks,
- publicly available implementations of FastRacer and related techniques, and an evaluation that shows Fast-Racer significantly outperforms prior work.

2 Background and Motivation

This section reviews data race detection of multithreaded programs using vector clocks. We also discuss closely related prior work on race detection of async-finish programs.

2.1 Race Detection with Vector Clocks

Many race detectors for multithreaded programs use vector clocks to track happens-before (HB) relations in an execution [6, 12, 35]. Each thread maintains a scalar clock that is incremented at *synchronization release* operations (e.g., lock release, monitor wait, thread fork and join, and volatile write). Each thread T also maintains a vector clock **C** of size *n*, where there are *n* threads in the application. The clock entry $C_T(U)$ records the clock of thread U when thread T last synchronized with U. A dynamic analysis updates per-variable vector clock metadata whenever a thread accesses a shared data or a lock variable. Vector clock operations require $\mathcal{O}(n)$ time and storage to monitor an execution with *n* threads.

Alg	Algorithm 1 FastTrack analysis at synchronization operations						
1:	procedure Spawn	⊳ Thread T spawns U					
2:	U.vc \leftarrow T.vc \cup {T.epoch}	-					
3:	T.epoch \leftarrow T.epoch + 1	▷ Increment T's scalar clock					
4:	$\textbf{U.epoch} \leftarrow \textbf{U.epoch} + 1$						
5:	procedure JOIN	\triangleright Thread T joins with U					
6:	for all $<$ t,c> in U.vc do	5					
7:	$T.vc[t] \leftarrow max(U.vc[t])$], T.vc[t])					
8:	procedure ACQUIRE	\triangleright T acquires lock L					
9:	for all <t,c> in L.vc do</t,c>	-					
10:	$T.vc[t] \leftarrow max(L.vc[t])$], T.vc[t])					
11:	procedure RELEASE	⊳ T releases lock L					
12:	$L.vc \leftarrow T.vc$						
13:	T.epoch \leftarrow T.epoch + 1						
14:	function CHECKHB(c@u,7	$() \qquad \triangleright \text{ Check HB between}$					
	epoch						
15:	return c@u <u>≺</u> T.vc						

FastTrack. The FastTrack algorithm tracks a single last writer and, in many cases, a single last reader [12]. The total order on writes in a data-race-free program allows FastTrack to store only the last write information. FastTrack stores the write metadata as an epoch c@T, which is a tuple consisting of the writer thread identifier T and the value of T's clock (say c) at the time of the write. The read metadata alternates between epoch and vector clock forms. An epoch representation suffices when there is a single reader or the current read happens after all previous reads. When there are concurrent readers, the read metadata is a vector clock (denoted by vc). Algorithm 1 shows the pseudocode for the dynamic analysis performed by FastTrack at synchronization operations. Before each access to a shared variable x by a thread T, FastTrack checks whether the current access by T happens after the previous write and all previous reads to x (CheckHB, Algorithm 1). A data race is reported if the current access by T is concurrent with the last accesses. The shared data and lock variable metadata are updated upon a thread access. FastTrack is popularly used as the basis for dynamically sound and precise² data race detection of multithreaded programs [6, 35].

2.2 Race Detection for Async-Finish Programs

Frameworks like X10 [8] and Habanero Java [7] support structured task parallelism with async-finish semantics. The statement "async {t}" creates a new child task t that can run in series or in parallel with its parent task. The statement "finish {t}" causes the current task to wait for all the recursively-created tasks within the block t. The async-finish model is terminally-strict, which means each join edge goes from the last instruction of a task to *any* of its ancestors in the inheritance tree [30]. In the following, we discuss dynamic data race detection techniques for async-finish programs.

SPD3. SPD3 [31] uses a *dynamic program structure tree* (DPST) to capture the semantics of an async-finish program.

²Sound means no false negatives, and precise means no false positives.



(a) An async-finish program with tasks T1–T5, shared variables var1 and var2, and a lock variable L1.



(b) The DPST for the program in Figure 1a.

Figure 1. An async-finish program and its DPST.

A DPST consists of step, finish, and async nodes. A step node represents the maximal sequence of instructions without any task management. An async node represents the spawning of a child task by a parent task. The descendants of an async node execute asynchronously with the remainder of the parent task. A finish node is created when a parent task spawns a child task and waits for the child, and its descendants, to complete. A finish node is thus the parent of all async, finish, and step nodes executed by its children or their descendants. Figure 1 shows an async-finish program and the corresponding DPST. All executions of a data-race-free async-finish program with the same input result in the same DPST [31].

The operational semantics of async-finish programs imply a left-to-right computation order of sibling nodes belonging to a common parent task. Thus, a DPST node's children are also ordered left-to-right to reflect the computation order in their parent. On a variable access, SPD3 searches for the *lowest common ancestor* (LCA) of the current access (i.e., a step node) and the last access stored in the variable's metadata. SPD3 reports a race if the left child of the LCA, which is an ancestor of the step node representing the last access, is an async node that indicates *concurrent* execution of the last access and the current task. The DPST allows SPD3 to maintain one metadata location for writes and two locations for reads in shadow memory.

PTRacer. PTRacer extends SPD3 by detecting apparent races in async-finish programs with locks [39]. PTRacer maintains two metadata locations each for reads and writes to a shared variable. The metadata per variable is proportional to the number of different locksets (i.e., set of locks held by the tasks at any time) with which the variable is accessed, which is reasonable in practice because variables are usually accessed with similar locking patterns.

PTRacer selects two "last read" ("last write") accesses from multiple parallel accesses with the same lockset to maintain constant metadata, such that any future write which can race with any one of the parallel reads (writes) will race with either one of the two chosen "last readers" ("last writers"). PTRacer makes these choices by selecting step nodes with the highest LCA among all parallel step nodes. PTRacer detects *all* races for a given input even in the presence of lock-based synchronization. Consider the shared variable var1 which is updated by the parallel tasks T2, T4, and T5 in Figure 1a. The step nodes corresponding to these accesses are S21, S41, and S51, respectively. Since, S21 and S51 have the highest LCA in the corresponding DPST, PTRacer will store these two accesses and discard the access information for S41.

The race analysis in PTRacer is similar to SPD3. PTRacer will report a race on var2 for the example in Figure 1 because the left child of the LCA of step nodes S21 and S31 is an async node. SPD3 will report false races on the variable var1.

PTRacer uses the DPST to maintain a constant amount of per-variable metadata independent of the number of tasks executing the program. Furthermore, PTRacer performs frequent lookups in the DPST to check whether two accesses can happen in parallel. However, the DPST can be deep for programs with a recursive pattern of task creation and large because of many step nodes. These lead to high run time and memory overheads (Section 5.2). PTRacer uses an array-based representation of the DPST and caches LCA lookups to improve the performance of LCA. However, the DPST and the LCA computation continue to be a significant bottleneck for several benchmarks. Thus, there is a need for *more efficient* techniques to help detect data races in async-finish programs.

3 FastRacer: Efficient Dynamic Data Race Detection for Async-Finish Programs

The thesis of this work is that vector clocks can provide better data locality for series-parallel (SP) maintenance in task-based programs compared to tree-based data structures used in prior work (e.g., [31, 37, 39]). We present *FastRacer*, a novel algorithm that integrates the benefits of vector-clockbased analysis *and* exploits structured parallelism of asyncfinish programs to limit the amount of per-variable metadata.

Task-based programs create more, often orders of magnitude, parallel tasks than threads in multithreaded programs. Race detectors like FastTrack use per-thread and per-variable vector clocks to capture the clock values of all the threads in the system. While this representation works fine for multi-threaded programs where the number of concurrent threads is comparatively small (~#cores), it is impractical for task-based programs and leads our FastTrack implementation to run out of memory on several benchmarks (Section 5.2). Storing only non-zero entries in a vector clock does not help since several concurrent tasks potentially access shared read-only variables. Furthermore, maintaining vector clocks proportional to the number of threads can detect only feasible data races and misses races among concurrent tasks [14]. In the following, we discuss novel ideas to solve these challenges.

Algorithm 2 FastRacer analysis at synchronization operations					
1:	procedure SPAWN > Task T spawns task U				
2:	if size(T.rw_vc) > THRESHOLD then				
3:	$T.ro_vc \leftarrow REF(\{T.ro_vc \cup T.rw_vc\}); T.rw_vc \leftarrow \emptyset;$				
4:	else				
5:	$U.ro_vc \leftarrow T.ro_vc; U.rw_vc \leftarrow T.rw_vc;$				
6:	$U.rw_vc \leftarrow U.rw_vc \cup \{T.epoch\}$				
7:	U.joined \leftarrow T.joined; U.lockset \leftarrow T.lockset;				
8:	$\bigcup IVC \leftarrow T.IVC \cup \{getClock(\bigcup epoch)\}$				
9:	$\textbf{T.epoch} \leftarrow \textbf{T.epoch} + 1 \textbf{; U.epoch} \leftarrow \textbf{U.epoch} + 1 \textbf{;}$				
10:	procedure JOIN ▷ Task T joins with U				
11:	T.joined \leftarrow T.joined \cup {getTaskId(U.epoch)}				
12:	function CHECKHB(c@u,T)				
13:	▷ Check HB between epoch c@u and T's access				
14:	return c@u \leq T.rw_vc or c@u \leq T.ro_vc or u \in				
	T.joined				

3.1 Adapting Vector Clocks for Task Parallelism

We analyzed the performance of FastTrack and found that operations on task vector clocks incur high time and space overhead. For example, the task join operation is a bottleneck because it requires comparing and merging *all* the clock values in the child and the parent tasks' vector clocks. Naïvely merging the vector clocks is not required since most vector clock entries remain unchanged during the lifetime of a task.

Tracking read-only clock entries. The first insight in Fast-Racer is that *most* vector clock entries for a task remain unchanged during the lifetime of the task, and the clock values continue to be the same as in the parent task. Thus, maintaining per-task copies of the clock entries is mostly redundant.

In FastRacer, a task vector clock is partitioned into a readonly (denoted by ro_vc) and a read-write portion (denoted by rw_vc). Child tasks in FastRacer maintain a reference to the ro_vc of their parents instead of maintaining redundant copies. During a spawn operation (Algorithm 2), FastRacer first checks if the size of rw_vc is greater than a threshold. If yes, then FastRacer merges ro_vc and rw_vc of the parent task into a new ro_vc for the child task and rw_vc of child task is kept empty. Otherwise, the child ro_vc points to the parent ro_vc and the parent's rw_vc is copied to the child's



Figure 2. The inheritance tree for program in Figure 1a.

rw_vc. Avoiding needless copies and redundant operations on the read-only portions of vector clocks helps reduce space overheads *and* improve performance. In the common case, most vector clock entries of a task remain unchanged, i.e., the size of rw_vc is small. Complete vector clock copies happen only when the size of rw_vc is greater than THRESHOLD.

Optimizing vector clock join. An access to a shared variable by a parent task after joining with a child task always happens after the child's accesses, since the accesses are synchronized by the join operation. The second insight is that FastRacer does not need to store the clock values of the child tasks *after* the join operation. Instead, tracking the set of all child tasks that have joined with the parent task suffice. Each task in FastRacer maintains the set of child tasks that have already joined with it in a joined data structure. No vector clock join occurs when a task T joins with task U, instead, the child task is added to the joined of the parent task (Algorithm 2). When a parent task spawns a new child task, the parent joined is copied to the child joined. The size of task vector clocks reduces significantly due to the joined optimization.

Vector clock caching. The vector clocks for a few tasks can be large even after the optimizations. FastRacer uses explicit attributes to cache recently used vector clock values to avoid the cost of looking up the map data structure representing vector clocks. FastRacer indexes into the vector clock map when the thread id is not among the most recently used.

3.2 Specializing for Async-Finish Programs

It can be expensive to maintain all concurrent readers of a shared variable in task-based programs. FastRacer uses coarse-grained tracking of task inheritance relationships to select relevant accesses from parallel readers/writers,³ which allows maintaining constant per-variable metadata per-lockset (i.e., the set of locks held by the task).

Maintaining constant per-variable metadata. FastRacer models the parent-child relationship among different tasks with a *task inheritance* tree. The nodes represent tasks, and edges represent the creation of child tasks by the parent. Figure 2 shows the inheritance tree for the program shown in

³When tasks use locks, two writes can happen in parallel but do not constitute a race if they are protected by the same lock.

Figure 1a (ignore the IVC labels for now). FastRacer uses the inheritance tree to efficiently select two accesses out of multiple concurrent accesses with the same lockset, such that any racy future access that races with any one of the concurrent accesses must be racy with either one of the two chosen last accesses. For parallel tasks accessing a variable with the same lockset, FastRacer stores the access history of the two tasks with the highest LCA in the inheritance tree. Consider the three accesses to var1 from tasks T2, T4, and T5 in Figure 1a. In the inheritance tree shown in Figure 2, task nodes T2 and T4 (or T5) have the highest LCA, and so FastRacer stores the access histories from T2 and T4 and discards T5. Any later access to var1, which is not parallel with both T2 and T4, will not be parallel with T5. So, discarding T5's access information is correct.⁴ The metadata stored per shared variable in Fast-Racer is proportional to the number of the different locksets with which the variable has been accessed.

The primary difference between a DPST and FastRacer's inheritance tree is in the *granularity* of the nodes. While a DPST decomposes a task into several unsynchronized regions represented by step nodes, an inheritance tree has just one node per task. The coarser modeling makes the inheritance tree much smaller and shallower than the DPST. Unlike a DPST, there is no left-to-right ordering in an inheritance tree. While PTRacer uses DPST to check the concurrency between two accesses and select accesses with the highest LCA, Fast-Racer only does the latter with the inheritance tree. FastRacer uses vector clocks for data race detection to compensate for the coarser modeling and loss of ordering between the nodes.

Inheritance vector clock. Instead of building an inheritance tree, FastRacer encodes the inheritance relations in a per-task array of clock values called *Inheritance Vector Clock* (IVC) for better performance. An IVC is an immutable vector clock that contains the clock values of all the reachable parents of a task T at the time of the creation of T. An IVC identifies the unique path in the inheritance tree from the root task to T, since a task spawn increments the scalar clock of the caller task. Whenever a parent task creates a child task, FastRacer copies the parent's IVC to the child and appends the parent's clock value at the end of child IVC. In Figure 2, the IVC of the parent task T3 is copied to the child task T4 and the current clock of T3 (assumed to be one) is appended.

Both IVC labeling and the Offset-Span (OS) labeling [22] schemes compute a unique label from the label of the immediate predecessor, and guarantee that the length of a task's label will always be proportional to the depth of the task in the inheritance tree. However, there are two differences. First, the IVC of a task, once created, is immutable. Unlike OS labeling, any further task join operations do not modify the IVC. Second, the Span part of OS labeling can only be assigned after

the task graph is complete, i.e., OS labels cannot be computed while building the task graph. Intuitively, IVC labels are more similar to the Offset part, with the constraint that they do not get updated at join operations. Whereas prior work uses OS labels for race detection [14, 22], FastRacer uses IVC labels to identify two accesses out of multiple parallel accesses having the highest LCA in the inheritance tree.

Computing tasks with the highest LCA. To compute the highest LCA among three parallel tasks, FastRacer iterates over the IVCs of all the tasks simultaneously and stops at the first point of difference. After that, FastRacer chooses the task with a different clock value and any one of the other two. If all the three clock values are different, any two out of three tasks can be chosen. If FastRacer reaches the end of an IVC for any task T, it indicates that T must be the parent of the other two in the inheritance tree, and so, FastRacer chooses the parent task T and any one from the other two. The tasks so chosen will have the highest LCA in the inheritance tree. The computation overhead depends on the sizes of the IVC, which is equal to the height of the inheritance tree. Since an inheritance tree is much shallower compared to a DPST, the computation is relatively efficient.

In Figure 1a, the accesses to var1 from three parallel tasks, T2, T4, and T5, are with the same lockset {L1}. The IVCs of the corresponding task nodes, shown in Figure 2, are IVC_{T2}=[1], IVC_{T4} = [2,1], and IVC_{T5} = [2,2]. To select two out of the three accesses, FastRacer iterates over the three IVCs simultaneously. The first point of difference is at the first entry itself: IVC_{T2}[0]=1, IVC_{T4}[0]=IVC_{T5}[0]=2. So, FastRacer selects a task with a different clock value, i.e., T2 and any one of other two, say T4 (or T5). For the example in Figure 2, tasks T2 and T4 (or T5) indeed have the highest LCA in the inheritance tree.

3.3 Detecting Data Races

FastRacer extends the FastTrack algorithm (Section 2.1) to detect races and update metadata.

Metadata. Figure 3 compares the metadata maintained in FastTrack and FastRacer. Given a task T, FastTrack uses a vector clock for storing the epoch values of *all* other tasks (line 10, Figure 3a). The attributes ro_vc, rw_vc, and joined in PerTaskMetadata in Figure 3b correspond to the optimizations introduced in Section 3.2. While FastTrack uses vector clocks to track lock operations, FastRacer uses the lockset mechanism [33]. Every task in FastRacer maintains and updates lockset at lock acquire and release operations.

The per-variable state in FastRacer is an array of lock metadata. Each lock metadata stores the access history of the variable with the distinct set of locks held before the access. Every per-lock metadata contains four access history entries, two each for reads and writes, which contain the epoch value and a reference to the IVC of the task at the time of access. Storing references to the IVCs suffice since they remain unchanged during the lifetime of a task.

⁴In async-finish semantics, a task must join with its ancestor, either immediate or recursive. In case a parent task calls join, all children tasks within this join scope, either immediate or recursive, join with it.

(a) Metadata maintained in FastTrack.

```
1 class PerVariableMetadata {
    epoch wr_md, rd_md; // Write and Read epoch
2
3
    map<taskid, clock> rd_vc; // Read vector clock
4 }
5 class PerLockMetadata {
    map<taskid, clock> vc;
6
7 }
8
  class PerTaskMetadata {
9
    epoch epoch;
    map<taskid, clock> vc;
10
11 }
```

(b) Metadata maintained in FastRacer.

```
1 class PerVariableMetadata {
    PerLockMetadata lock[];
2
3 }
4 class PerLockMetadata {
    set<lockID> lockset;
5
    epoch rd1, rd2; // Reader 1 and 2 metadata
6
    IVC* rd1_ivc, rd2_ivc; // IVC of Reader 1 and 2
7
    epoch wr1, wr2; // Writer 1 and 2 metadata
8
    IVC* wr1_ivc,wr2_ivc; // IVC of Writer 1 and 2
9
10 }
11 class PerTaskMetadata {
12
    epoch epoch;
    map<taskid, clock> ro_vc; // Rd-only vector clock
13
    map<taskid, clock> rw_vc; // Rd-wr vector clock
14
15
    set<taskid> joined; // All joined child tasks
    set<lockId> lockset; // Locks held by the task
16
17
    int IVC[];
18
 }
```

Figure 3. Comparison of metadata state maintained in Fast-Track and FastRacer.

Race checks. When a shared variable is accessed, Fast-Racer iterates over all the lock metadata corresponding to distinct locksets with which the shared memory variable has been accessed. An empty intersection of the lockset of the current access and the lock metadata implies potentially parallel accesses. If the two locksets are disjoint, FastRacer checks if the epoch values stored in the access history happens before the current access using vector clocks (CHECKHB, Algorithm 2). If there is no such relationship, it implies that the prior access is concurrent with the current access. Finally, FastRacer reports a data race if one of the two accesses is a write.

Before accessing a variable x, FastTrack compares the current task's vector clock with the epoch(s) stored in x's access history to determine if the current access happens *after* the past accesses (CHECKHB, Algorithm 1). FastRacer, apart from the vector clock entry check, also checks if the tasks present in x's access history belong to joined of the current task. If all the tasks are present in joined, FastRacer infers that the current access happens *after* prior accesses. Since the vector clock is spread across ro_vc, rw_vc, and joined, CHECKHB (Algorithm 2) checks the HB relation against all of them.

Metadata updates. FastRacer updates the read metadata corresponding to the current lockset if a read does not race with prior writes. FastRacer checks if any of the read epochs in the lock metadata corresponding to the current lockset happens before the current task's access. If yes, then Fast-Racer updates that access entry with the current task's epoch and IVC. Otherwise, there are three parallel reads, and Fast-Racer needs to select two with the highest LCA. FastRacer iterates over the IVC of all three access entries and stops either at the first point of difference or if one of the IVCs end. FastRacer stores the access history of the task corresponding to the selected IVC and any one of the other two. Using any one of the other two works since, in both the cases, the two chosen tasks will have the highest LCA in the inheritance tree (Section 3.2). Figure 4 shows an example of how FastRacer updates the read metadata. Assume tasks T2, T3, and T4 all read a shared variable x and task T5 writes x. After the reads from T2 and T3, the two readers stored for the variable x are T2 and T3, because both these tasks can run in parallel. Since task T4 is spawned by T1 after T1 synchronizes with T2, so the read by T4 happens after the read of T2. The read metadata entry of T2 is replaced by T4. Next, when T5 writes x, Fast-Racer checks the access with all previous reads and reports a read-write race between the accesses from T4 and T5. The steps performed by FastRacer on a write access are similar.

Synchronization operations. During a spawn operation (Algorithm 2), FastRacer checks if the size of rw_vc is greater than a threshold. If yes, FastRacer merges ro_vc and rw_vc of the parent task into a new ro_vc for the child task and rw_vc of the child is kept empty. Otherwise, the child ro_vc references the parent ro_vc, and the child rw_vc is copied from parent rw_vc. Thereafter, FastRacer copies the parent's joined and lockset to the child's joined and lockset, respectively. In case of a join operation, the joined of the parent task is updated to contain the id of the child task.

A task's vector clock is not updated in case of lock acquire and release operations. Instead, FastRacer uses PTRacer's mechanism to deal with lock operations. Each task maintains a lockset. When a task T acquires a lock L, the lockset of T is updated to contain lock L. In case of a lock release, L is removed from the lockset. When T accesses a shared variable x, the lockset is copied to the variable metadata. During a race check, FastRacer checks for the intersection of the locksets from the metadata history to infer a data race. In practice, variables are accessed with the same set of locks, and hence maintaining access metadata for different sets of unique locks is reasonable. Furthermore, metadata update operations depend on the size of the IVC, but we find that the maximum depth of the inheritance tree is small (\leq 30 for our benchmarks).



Figure 4. An async-finish program with five tasks T1–T5 and the corresponding inheritance tree.

3.4 Characterizing FastRacer

While FastTrack's race coverage is limited to the observed schedule, FastRacer can detect per-input apparent races like PTRacer. FastRacer can detect races in other schedules due to the following differences with FastTrack. First, the races reported by FastTrack can vary across schedules since many tasks can map to a single thread, and the exact sequence of tasks mapped may vary across schedules. FastRacer stores vector clocks per task, and is not impacted by the mapping of tasks to threads. Second, FastTrack tracks the HB relations to establish order among synchronization operations, which is sensitive to the order of lock operations and varies across schedules. FastRacer uses the lockset technique to track synchronization operations and stores two reads and two writes per lockset in per-variable metadata. The metadata structure enables FastRacer to detect races irrespective of the order of lock operations, so the number of races reported is the same across schedules. The companion report discusses the correctness of FastRacer [16].

4 Implementation

Our implementation extends the PTRacer artifact.⁵ A static compiler pass using LLVM 3.7 instruments load and store instructions in C++ programs, and inserts function calls to execute the appropriate race detection analysis. The implementation uses Intel TBB for task parallelism. The public implementation of PTRacer reports wrong race results for the benchmarks fluidanimate, kmeans, streamcluster, and sort (see Section 5.1). We found an implementation error was corrupting the DPST built by PTRacer, and there was a race while updating a global array used for LCA hashing. After fixing these issues, the race reports were the same across all the tools. Our modifications have minimal (<1%) impact on the performance of PTRacer, and we use our fixed version for the evaluation. Our prototype implementation of FastRacer extend the same static compiler pass to ensure all the prototypes do the same work. We have also reimplemented FastTrack. Our implementations are publicly available.⁶

⁵https://github.com/rutgers-apl/PTRacer

⁶https://github.com/prospar/fastracer-pmam-2022

Race detection for fork-join programs. Utterback et al. propose a parallel and asymptotically optimal algorithm called WSP-Order for race detection of fork-join programs [37]. The algorithm uses two order maintenance (OM) data structures to maintain two total orders of all strands in the computation. A strand is a sequence of instructions that contain no parallel primitives and executes sequentially. A strand x logically precedes strand y if and only if x precedes y in both orderings. These orderings are sufficient to determine SP relationships. The two OM data structures support constant-time operations like insert and query, and most concurrent updates do not need synchronization. However, large parts of the OM data structure are updated during relabel operations and hence require synchronization. Since relabel operations are serialized, the algorithm modifies a work-stealing task scheduler to prioritize the operations. Furthermore, workers blocked on an insert or a query operation help with the relabel instead of being idle. Note that the WSP-Order algorithm does not support lock-based synchronization and requires tight coupling with a work-stealing scheduler for good performance [37].

The public implementation of WSP-Order is called C-RACER.⁷ We reimplement C-RACER with Intel TBB in LLVM for a fair comparison. An important contribution in the C-RACER work is the parallelization of the relabel operations. We have not implemented task scheduler support for parallel relabel operations, relabels in our implementation are serial. The total time taken in the serial relabel operations is small in our experiments. The run time of the benchmarks we report for C-RACER *does not account* for the time taken for the relabel operations, which is a lower bound (Section 5.2).

5 Evaluation

This section compares FastRacer with the closest prior work, FastTrack [12], PTRacer [39], and C-RACER [37].

5.1 Experimental Setup

Benchmarks. We reuse twelve TBB-based applications used by PTRacer for our evaluation. These include four applications, blackscholes, fluidanimate, streamcluster, and swaptions, from the PARSEC benchmark suite [2], five geometry and graphics applications, convexHull, delRefine, delTriang, nearestNeigh, and rayCast, from the PBBS suite [36], and three applications, karatsuba, kmeans, and sort, from the Structured Parallel Programming book [21]. We left out the PARSEC application, bodytrack, because of a compilation error, and ignore the C-RACER benchmarks because they use Cilk-5 [37].

The benchmarks follow spawn-sync semantics where a child task joins with its immediate parent (async-finish semantics are more general) and do not use locks, so we were able to run C-RACER successfully for all the benchmarks.

⁷https://github.com/wustl-pctg/cracer

	UM	FT	CR	РТ	FR
blackscholes	0.62	9.64	10.38	54.78	48.03
fluidanimate	0.51	69.93	0.58	9.17	2.04
streamcluster	0.84	-	1.17	22.76	6.91
swaptions	0.19	-	0.98	17.34	3.87
convexHull	1.69	-	2.97	19.94	12.71
delRefine	3.71	-	6.25	125.37	125.39
delTriang	3.63	-	8.19	113.69	106.77
nearestneigh	2.19	-	5.60	106.86	125.26
rayCast	1.28	-	4.08	18.37	14.35
karatsuba	0.02	0.86	0.18	9.34	2.81
kmeans	0.11	-	2.52	15.82	5.99
sort	0.15	18.96	0.20	9.87	2.40

Table 1. Comparison of the memory overhead (in GBs).

Evaluation platform. The experiments execute on an Intel Xeon Gold 5218 system with one 16-core processor with hyperthreading disabled, 128 GB DDR4 primary memory, running Ubuntu Linux 20.04.3 LTS with kernel version 5.11.0.

5.2 Performance Results

Figure 5 reports the performance of C-RACER, PTRacer, and FastRacer for all the benchmarks (arranged alphabetically). Every bar averages ten trials and is normalized to the baseline, which runs the unmodified benchmarks without instrumentation. Smaller bars mean better run times. By default, the TBB scheduler creates c threads and multiplexes the tasks to the c threads, where c is the number of cores in the system. The results in Figure 5 are with 16 threads.

Our reimplementation of C-RACER incurs an overhead of 10.66X over the unmodified applications. The average performance slowdown incurred by PTRacer is 7.70X, while it is 6.41X for FastRacer. FastRacer shows substantial improvement for multiple benchmarks like convexHull, delTriang, fluidanimate, kmeans, and rayCast. The better performance of FastRacer is due to improved locality from vector clocks compared to cache-unfriendly tree traversals in C-RACER and PTRacer. On average, FastRacer outperforms both C-RACER and PTRacer by 1.66X and 1.20X, respectively.

FastTrack successfully executed with four benchmarks that required less memory, blackscholes, fluidanimate, karatsuba, and sort, and got killed on the other benchmarks. The overhead of FastTrack for the four benchmarks is 50.26X, and FastRacer outperforms FastTrack by 4.62X. This result shows that the optimizations proposed in FastRacer are effective in reducing both run time and space overheads.

Memory overhead. Table 1 compares the peak memory requirement of each benchmark with the four techniques, as reported by the Massif tool in Valgrind [24]. Column 2, *UM*, shows the memory requirement of the unmodified application, while *FT*, *CR*, *PT*, and *FR* stand for FastTrack, C-RACER, PTRacer, and FastRacer, respectively. The results show that

using IVCs for encoding task inheritance in FastRacer compared to the fine-grained DPST structure in PTRacer provides significant memory savings in maintaining per-task metadata, especially for fluidanimate, karatsuba, kmeans, sort, and swaptions. C-RACER usually requires the least memory since it maintains order maintenance structures, but has high performance overhead compared to FastRacer.

Scalability. Figure 6 shows the scalability plots of a few benchmarks as we vary the number of threads (in powers of two) used by the TBB scheduler. The *Unmodified* configuration shows that most applications scale well, excepting blackscholes, convexHull, and karatsuba. FastRacer scales better than PTRacer for delTriang, kmeans, and swaptions, for the range of thread counts we have evaluated. The scaling behavior of C-RACER, PTRacer, and FastRacer are very similar for the remaining benchmarks, which we omit for lack of space.

Platform sensitivity. We evaluate the sensitivity of the optimizations by re-running experiments on an Intel Xeon Silver 4114 system with two ten core processors with hyperthreading turned off, 128 GB primary memory, running Ubuntu Linux 18.04.6. We do not show the plots for lack of space. C-RACER, PTRacer, and FastRacer incur overheads of 11.60X, 9.44X, and 6.85X, respectively. FastRacer outperforms both C-RACER and PTRacer by 1.69X and 1.38X, respectively.

FastRacer outperforms prior work on the *same* set of benchmarks as used in PTRacer. Given the generic nature of the optimizations, we expect a similar qualitative trend for benchmarks where threads join with arbitrary ancestors (e.g., asyncfinish semantics). More importantly, FastRacer disproves the assumption made in *all* prior work that vector-clock-based analysis is not suited for task-based programs.

5.3 Data Races and Run-time Statistics

Table 2 summarizes the run-time statistics. The data is the average from 10 trials with a statistic-collecting configuration. Columns 2–4 in the table show the average number of tasks spawned by the benchmarks and the number of read and write accesses. Columns 5–7 show the number of data races reported by the different tools. To stress-test the correctness of our implementations, we introduced data races in a few benchmarks that already did not have known races. The suffix "-r" denotes benchmarks that have been modified to introduce races for evaluation. All the tools are expected to report the same number of races for a given application with a fixed input. C-RACER (CR), PTRacer (PT), and FastRacer (FR) report the same violations for all the benchmarks.

6 Related Work

Race detection for task-based programs. Mellor-Crummey exploited the structural property of fork-join programs to show that tracking two readers and a single writer per memory location is sufficient for sound data race detection [22].



Figure 5. Performance comparison of the different techniques normalized to the unmodified execution time of the benchmarks.



Figure 6. Scalability results of a few benchmarks on the Intel Gold platform described in Section 5.1.

	# Tasks	ACC ($\times 10^{6}$)		Da	Data Races		
	$(\times 10^{3})$	RDs	WRs	CR	РТ	FR	
blackscholes	0.20	90	50	21	21	21	
fluidanimate-r	1.60	26	0.7	40	40	40	
streamcluster-r	180	363	13	80	80	80	
swaptions	960	77	77	0	0	0	
convexHull	8.50	30	0	0	0	0	
delRefine	1000	15	0	0	0	0	
delTriang	790	30	20	0	0	0	
nearestNeigh	2800	51	8	0	0	0	
rayCast	1900	160	0	0	0	0	
karatsuba	1.98	3.4	0.8	0	0	0	
kmeans-r	35	570	10	75	75	75	
sort	0.70	11	0.06	1024	1024	1024	

Table 2. Run-time statistics across different benchmarks.

Since then, there has been much work to design race detection algorithms to utilize the serial-parallel (SP) structure of programs with constant space overhead for metadata [9, 11]. ESP-bags is an extension to SP-bags that supports the finish construct in async-finish programs [30]. However, these approaches constrain the program to execute serially in depthfirst order, which does not scale. TARDIS does not keep track of the SP relationships among program strands [19]. Instead, TARDIS maintains log-based access sets and lazily detects races by checking for overlapping intersections of access sets of logically parallel sub-computations at join points.

Race detection for multithreaded programs. Static analysis can potentially detect all *feasible* data races across all possible executions (i.e., no false negatives), but usually do not scale to large programs and suffer from false positives, which developers loathe [4, 23]. Dynamic analyses have the potential to be sound and precise for the *observed* executions. Many dynamic data race detection analyses track the happensbefore relation to infer data races [6, 35]. *Lockset* analysis reports data races when a locking discipline is violated, but can report false races [26, 33]. Hybrid techniques integrate both HB and lockset analysis [26], but continue to suffer from the disadvantages of both techniques. Other techniques sacrifice soundness for performance by sampling memory accesses [3, 6, 20] or require hardware support to speed up the race detection analysis [28, 38].

Improve race detection coverage. Many techniques perturb the execution in an attempt to break spurious HB relations [10, 34]. *Predictive* techniques aim to detect data races that can occur in other correct reorderings of memory accesses by observing one dynamic execution [15, 32].

7 Conclusion

Prior work has overlooked the possibility of using vector clocks for race detection of task-based programs. FastRacer introduces novel optimizations that reduce the metadata redundancy of task vector clocks. FastRacer also exploits the structured execution of async-finish programs to limit the pervariable metadata overhead by using a coarse-but-efficient encoding of task inheritances. FastRacer shows substantial performance improvements over FastTrack, and outperforms state-of-art approaches C-RACER and PTRacer. Our proposed vector clock optimizations allow for efficient dynamic data race detection for task-based async-finish programs.

Acknowledgments

We thank Michael D. Bond for feedback on the text, the PTRacer authors for help with their implementation, and Vivek Kumar for discussions and advice. This material is based upon work supported by IITK Initiation Grant and SERB under Grant SRG/2019/000384.

References

- Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Charles E. Leiserson. 2004. On-the-Fly Maintenance of Series-Parallel Relationships in Fork-Join Multithreaded Programs. In SPAA. 133–144.
- [2] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In *PACT*. 72–81.
- [3] Swarnendu Biswas, Man Cao, Minjia Zhang, Michael D. Bond, and Benjamin P. Wood. 2017. Lightweight Data Race Detection for Production Runs. In CC. 11–21.
- [4] Sam Blackshear, Nikos Gorogiannis, Peter W. O'Hearn, and Ilya Sergey. 2018. RacerD: Compositional Static Race Detection. *PACMPL* 2, OOPSLA, Article 144 (Oct. 2018).
- [5] Robert D. Blumofe, Christopher F. Joerg, Bradley C. Kuszmaul, Charles E. Leiserson, Keith H. Randall, and Yuli Zhou. 1995. Cilk: An Efficient Multithreaded Runtime System. In *PPoPP*. 207–216.
- [6] Michael D. Bond, Katherine E. Coons, and Kathryn S. McKinley. 2010. PACER: Proportional Detection of Data Races. In *PLDI*. 255–268.
- [7] Vincent Cavé, Jisheng Zhao, Jun Shirako, and Vivek Sarkar. 2011. Habanero-Java: the New Adventures of Old X10. In PPPJ. 51–61.
- [8] Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun, and Vivek Sarkar. 2005. X10: An Object-Oriented Approach to Non-Uniform Cluster Computing. In *OOPSLA*. 519–538.
- [9] Guang-Ien Cheng, Mingdong Feng, Charles E. Leiserson, Keith H. Randall, and Andrew F. Stark. 1998. Detecting Data Races in Cilk Programs That Use Locks. In SPAA. 298–309.
- [10] Mahdi Eslamimehr and Jens Palsberg. 2014. Race Directed Scheduling of Concurrent Programs. In *PPoPP*. 301–314.
- [11] Mingdong Feng and Charles E. Leiserson. 1997. Efficient Detection of Determinacy Races in Cilk Programs. In SPAA. 1–11.
- [12] Cormac Flanagan and Stephen N. Freund. 2009. FastTrack: Efficient and Precise Dynamic Race Detection. In *PLDI*. 121–133.
- [13] Matteo Frigo, Charles E. Leiserson, and Keith H. Randall. 1998. The Implementation of the Cilk-5 Multithreaded Language. In *PLDI*. 212– 223.
- [14] Yizi Gu and John Mellor-Crummey. 2018. Dynamic Data Race Detection for OpenMP Programs. In *SC*. Article 61, 12 pages.
- [15] Dileep Kini, Umang Mathur, and Mahesh Viswanathan. 2017. Dynamic Race Prediction in Linear Time. In *PLDI*. 157–170.

- [16] Shivam Kumar, Anupam Agrawal, and Swarnendu Biswas. 2021. Efficient Data Race Detection of Async-Finish Programs Using Vector Clocks. *CoRR* abs/2112.04352 (2021).
- [17] Doug Lea. 2000. A Java Fork/Join Framework. In ACM Conference on Java Grande. New York, NY, USA, 36–43.
- [18] N. G. Leveson and C. S. Turner. 1993. An Investigation of the Therac-25 Accidents. *IEEE Computer* 26, 7 (July 1993), 18–41.
- [19] Li Lu, Weixing Ji, and Michael L. Scott. 2014. Dynamic Enforcement of Determinism in a Parallel Scripting Language. In *PLDI*. 519–529.
- [20] Daniel Marino, Madanlal Musuvathi, and Satish Narayanasamy. 2009. LiteRace: Effective Sampling for Lightweight Data-Race Detection. In *PLDI*. 134–143.
- [21] Michael McCool, James Reinders, and Arch Robison. 2012. Structured Parallel Programming: Patterns for Efficient Computation (first ed.). Morgan Kaufmann Publishers Inc.
- [22] John Mellor-Crummey. 1991. On-the-Fly Detection of Data Races for Programs with Nested Fork-Join Parallelism. In SC. 24–33.
- [23] Mayur Naik and Alex Aiken. 2007. Conditional Must Not Aliasing for Static Race Detection. In POPL. 327–338.
- [24] Nicholas Nethercote and Julian Seward. 2007. Valgrind: A Framework for Heavyweight Dynamic Binary Instrumentation. In *PLDI*. 89–100.
- [25] Robert H. B. Netzer and Barton P. Miller. 1992. What Are Race Conditions? Some Issues and Formalizations. ACM Letters on Programming Language and Systems 1, 1 (March 1992), 74–88.
- [26] Robert O'Callahan and Jong-Deok Choi. 2003. Hybrid Dynamic Data Race Detection. In PPoPP. 167–178.
- [27] PCWorld. 2012. Nasdaq's Facebook Glitch Came From Race Conditions. Online. http://www.pcworld.com/article/255911/nasdaqs_ facebook_glitch_came_from_race_conditions.html
- [28] Yuanfeng Peng, Benjamin P. Wood, and Joseph Devietti. 2017. PARSNIP: Performant Architecture for Race Safety with No Impact on Precision. In *MICRO*. 490–502.
- [29] K. Poulsen. 2004. SecurityFocus News: Tracking the blackout bug. http://www.securityfocus.com/news/8412
- [30] Raghavan Raman, Jisheng Zhao, Vivek Sarkar, Martin Vechev, and Eran Yahav. 2010. Efficient Data Race Detection for Async-Finish Parallelism. In RV. 368–383.
- [31] Raghavan Raman, Jisheng Zhao, Vivek Sarkar, Martin Vechev, and Eran Yahav. 2012. Scalable and Precise Dynamic Datarace Detection for Structured Parallelism. In *PLDI*. 531–542.
- [32] Jake Roemer, Kaan Genç, and Michael D. Bond. 2018. High-coverage, Unbounded Sound Predictive Race Detection. In *PLDI*. 374–389.
- [33] Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro, and Thomas Anderson. 1997. Eraser: A Dynamic Data Race Detector for Multi-threaded Programs. In SOSP. 27–37.
- [34] Koushik Sen. 2008. Race Directed Random Testing of Concurrent Programs. In *PLDI*. 11–21.
- [35] Konstantin Serebryany, Alexander Potapenko, Timur Iskhodzhanov, and Dmitriy Vyukov. 2012. Dynamic Race Detection with LLVM Compiler. In *RV*. 110–114.
- [36] Julian Shun, Guy E. Blelloch, Jeremy T. Fineman, Phillip B. Gibbons, Aapo Kyrola, Harsha Vardhan Simhadri, and Kanat Tangwongsan. 2012. Brief Announcement: The Problem Based Benchmark Suite. In SPAA. 68–70.
- [37] Robert Utterback, Kunal Agrawal, Jeremy T. Fineman, and I-Ting Angelina Lee. 2016. Provably Good and Practically Efficient Parallel Race Detection for Fork-Join Programs. In SPAA. 83–94.
- [38] Benjamin P. Wood, Luis Ceze, and Dan Grossman. 2014. Low-Level Detection of Language-Level Data Races with LARD. In ASPLOS. 671–686.
- [39] Adarsh Yoga, Santosh Nagarakatte, and Aarti Gupta. 2016. Parallel Data Race Detection for Task Parallel Programs with Locks. In FSE. 833–845.