



Efficient Race Detection with Futures

Robert Utterback
Monmouth College
rutterback@
monmouthcollege.edu

Kunal Agrawal
Washington University in
St. Louis
kunal@wustl.edu

Jeremy Fineman
Georgetown University
jfineman@
cs.georgetown.edu

I-Ting Angelina Lee
Washington University in
St. Louis
angelee@wustl.edu

Abstract

This paper addresses the problem of provably efficient and practically good on-the-fly determinacy race detection in task parallel programs that use futures. Prior works on determinacy race detection have mostly focused on either task parallel programs that follow a series-parallel dependence structure or ones with unrestricted use of futures that generate arbitrary dependences. In this work, we consider a restricted use of futures and show that we can detect races more efficiently than with general use of futures.

Specifically, we present two algorithms: MultiBags and MultiBags+. MultiBags targets programs that use futures in a restricted fashion and runs in time $O(T_1 \alpha(m, n))$, where T_1 is the sequential running time of the program, α is the inverse Ackermann's function, m is the total number of memory accesses, n is the dynamic count of places at which parallelism is created. Since α is a very slowly growing function (upper bounded by 4 for all practical purposes), it can be treated as a close-to-constant overhead. MultiBags+ is an extension of MultiBags that target programs with general use of futures. It runs in time $O((T_1 + k^2) \alpha(m, n))$ where T_1 , α , m and n are defined as before, and k is the number of future operations in the computation. We implemented both algorithms and empirically demonstrate their efficiency.

*CCS Concepts • **Theory of computation** → *Dynamic graph algorithms*; • **Software and its engineering** → *Software testing and debugging*; Parallel programming languages; • **Computing methodologies** → *Parallel computing methodologies*;

*Keywords dynamic program analysis; determinacy race; race detection; series-parallel maintenance

1 Introduction

Races constitute a major source of errors in parallel programs. Since they lead to nondeterministic program behaviors, they are extremely challenging to detect and debug. In this work,

we focus on the problem of race detection for task-parallel programs, where the programmer denotes the *logical* parallelism of the computation using high-level parallel control constructs provided by the platform, and lets the underlying runtime system perform the necessary scheduling and synchronization. Examples of task parallel platforms include OpenMP [42], Intel's TBB [29, 46], IBM's X10 [14], various Cilk dialects [17, 25, 30, 34, 36], and Habanero dialects [5, 12].

In the context of task parallel programs, the focus is typically on detecting **determinacy races** [19] (also called **general races** [40]), which occur when two or more logically parallel instructions access the same memory location and at least one access is a write. In the absence of a determinacy race, a task parallel program for a given input behaves *deterministically*.

Over the years, researchers have proposed several determinacy race algorithms [6, 19–21, 39, 44, 45, 53, 57, 60] for task parallel code. These algorithms perform race detection **on the fly** as the program executes, and consist of two main components: (1) an **access history** that keeps track of previous readers and writers for each memory location; and (2) a **reachability data structure** for maintaining and querying whether two instructions are logically in parallel. On each memory access, the detector checks whether the current access is logically parallel with the previous accessors (stored in the access history) to determine whether a race exists.

Most prior work focuses on a restricted set of computations, namely computations that can be represented as **series-parallel dags (SP dags)** [58] with nice structural properties, such as ones generated using fork-join parallelism (i.e., `spawn/sync` or `async/finish`). Prior works show that one can race detect computations that are SP dags efficiently by exploiting the structural properties. In particular, the reachability data structure can be maintained and queried with *no asymptotic overhead* for both serial [6, 21] and parallel executions [57]. Moreover, the access history needs to store only a constant number of accessors per memory location to correctly race detect for such computations [19, 20, 39].

The use of **futures** has become a popular way to extend fork-join parallelism. Since their proposal [4, 24] in the late 70s, futures have been incorporated into various parallel platforms [3, 12–14, 23, 26, 33, 38]. Researchers have studied scheduling bounds [2, 9] and cache efficiency [27, 51] for using futures with fork-join computations. Kogan and Herlihy [32] study linearizability of concurrent 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.

PPoPP '19, February 16–20, 2019, Washington, DC, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6225-2/19/02...\$15.00

<https://doi.org/10.1145/3293883.3295732>

accessed using futures. Surendran and Sarkar [52] proposed using futures to automatically parallelize programs.

The use of futures can form arbitrary dependencies, and thus computations generated by a parallel program that uses futures are no longer series-parallel. However, not much work has been done on race detecting programs with more general dependence structures.

Two prior works exist on race detection for programs that use futures and both are sequential (no known parallel algorithms exist). An algorithm proposed by Surendran and Sarkar [54] has high overheads — the running time is $O(T_1(f + 1)(k + 1))$ where T_1 is the **work**, or sequential running time of the program without race detection, f is number of future objects and k is the number of future operations. That is, the running time of the race detection algorithm increases quadratically with the total number of futures used in the program. More recently, Agrawal et al. [1] present a sequential algorithm to perform race detection on SP dags with k added non-series-parallel edges in $O(T_1 + k^2)$ time, which is the best known running time. The algorithm is difficult to implement however, since it requires storing all the nodes in the computation graph and traversing the graph during execution to update labels. Thus, no actual implementation of the algorithm exists to date.

Contributions

While prior work on race detection has focused on either structured SP dags or unrestricted use of futures that generates arbitrary dependences, we consider a restricted use of futures. Researchers have observed in other contexts [27] that using futures in a specific restricted manner can reduce scheduling and cache overheads. We call such a use case as **structured** futures and observe that the restricted use allows race detection to be performed much more efficiently than general use of futures. This class of futures is quite natural and can be checked with program analysis. We provide the precise definition in Section 2; informally, it requires that the instruction that creates the future is sequentially before the instruction that uses the handle.

We present two practical algorithms for race detecting programs with futures: MultiBags and MultiBags+. The main contribution for both algorithms is a novel reachability data structure. Both algorithms run the program sequentially for a given input and report a race if and only if one exists, following the same correctness criteria as prior work. MultiBags focuses on structured use of futures and incurs very little overhead — a multiplicative overhead in the inverse Ackermann’s function, which is upper bounded by 4 for all practical purposes [16]. MultiBags+ is an extension of MultiBags, which handles general use of futures and has overhead comparable to the state-of-the-art theoretical algorithm [1] (i.e., multiplicative overhead of the inverse Ackermann’s function) and can be implemented efficiently. We have implemented both algorithms and empirically evaluated them.

The empirical results show that both algorithms can maintain reachability efficiently for their designated use cases.

Specifically, we make the following contributions:

- **MultiBags:** We propose MultiBags, an algorithm to race detect programs that use structured futures (Section 4). We prove its correctness and show that it race detects in $O(T_1\alpha(m, n))$ time where T_1 is the work, α is the inverse Ackermann’s function, m is the number of memory accesses, and n is the dynamic count of places at which parallelism is created. Since α is a very slow-growing function, this bound is essentially $O(T_1)$ for all intents and purposes.
- **MultiBags+:** We propose MultiBags+, an algorithm to race detect programs that use general futures (Section 5). We prove its correctness and show that it race detects in $O(T_1 + k^2)\alpha(m, n)$ time where T_1 , α , m , and n are defined as above, and k is the number of future operations in the computation (Section 5). Compared to the state-of-the-art algorithm proposed by Agrawal et al. [1], MultiBags+’s running time has a multiplicative overhead of the inverse Ackermann’s function. Unlike the state-of-the-art, however, MultiBags+’s relative simplicity allows it to be implemented efficiently in practice. We provide a more detailed comparison between our MultiBags+ algorithm and the state-of-the-art [1] in Section 5.
- **FutureRD:** We have built a prototype race detector called FutureRD based on MultiBags and MultiBags+. Empirical evaluation with FutureRD shows that our algorithms allow reachability to be maintained efficiently, incurring almost no overhead (geometric means of 1.06× and 1.40× overhead for MultiBags and MultiBags+, respectively). The overall race detection incurs geometric means of 20.48× and 25.98× overhead, respectively.

2 Preliminaries and Definitions

Parallel control constructs: Our algorithms are described assuming parallelism in programs is generated using four primitives: `spawn`, `sync`, `fut-create` and `get`. The algorithms themselves are general and can be applied to platforms that use other constructs that generate similar types of dags. We assume that `spawn` and `sync` are used to generate fork/join or series/parallel structures. In particular, for the purposes of this paper, function F can **spawn** off a child function G , invoking G without suspending the continuation of F , thereby creating parallelism; similarly, F can invoke **sync**, joining together all previously spawned children within the functional scope.¹ We assume **fut-create** and **get** primitives are used to create and join futures, respectively. Like `spawn`, one can precede a function call to G in F with `fut-create`, which allows G to execute without suspending F . Unlike `spawn`, however, a function call invoked with `fut-create`

¹Some constructs, such as `async/finish` primitives have slightly different restrictions, they still generate SP dags and our algorithms can be modified to apply to these programs.

can escape the scope of a sync — a subsequent sync joins together previously spawned functions but does not wait for function calls preceded by `fut-create` to return. Instead, `fut-create` of a function instance G returns a **future handle** h , and the program must explicitly invoke `get` on h to join with G . A call of `get` on h **block** until G completes.

Modeling parallel computations: One can model the execution of a parallel program for a given input as a **dag** (**directed acyclic graph**) G_{full} , whose nodes are **strands** — sequence of instructions containing no parallel control — and edges are control dependencies among strands. The dag unfolds dynamically as the program executes. A strand u is sequentially before another strand v (denoted by $u < v$) if there is a path from u to v in the dag; two nodes u and v are **logically parallel** if there is no path from one to the other. The performance of a computation can be measured in two terms: the **work** T_1 of the computation is the execution time of the computation on a single processor; the **span** (also called **depth** or **critical-path length**) T_∞ of the computation is its execution time on an infinite number of processors (or, longest sequential path through the dag).

Series-parallel dags: Computations which use only `spawn` and `sync` can be modeled as **series-parallel dags** (**SP-dag**) [58] that have a single **source** node with no incoming edges and a single **sink** node with no outgoing edges. Upon the execution of a `spawn`, a **fork node** is created with two outgoing edges: one leads to the first strand in the spawned child function and one leads to the continuation of the parent. Upon the execution of a `sync`, a **join node** is created, that has two or more incoming edges, joining the previously spawned subcomputations.²

SP dags can be constructed recursively as follows.

- **Base Case:** the dag consists of a single node that is both the source and the sink.
- **Series Composition:** let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be SP-dags on distinct nodes. Then a series composition G is formed by adding an edge from $sink(G_1)$ to $source(G_2)$ with $source(G) = source(G_1)$ and $sink(G) = sink(G_2)$.
- **Parallel Composition:** let $G_L = (V_L, E_L)$ and $G_R = (V_R, E_R)$ be SP-dags on distinct nodes. Then the parallel composition G is formed as follows: add a fork node f with edges from f to both sources, and a join node j with edges from both sinks to j . $source(G) = f$ and $sink(G) = j$. We refer to G_L and G_R as the **left subdag** and **right subdag**, respectively, of both the fork f and join j .

Adding futures: We model computations that employ futures in addition to `spawn` and `sync` as a set of independent SP dags connected to each other via **non-SP edges** due to `fut-create` and `get` calls. If a function F spawns a function G , then the strands of F and G are part of the same SP dag. However, if function G calls H via `fut-create`, then the

first strand, say v , of H is the source of a different SP dag. The last strand of H will be the sink node of this SP dag. Therefore, if a program calls `fut-create` f times, then it has $f + 1$ SP dags (f future functions plus the main function) connected to each other via non-SP edges.

These non-SP edges are incident on strands that end with `fut-create` and ones that immediately follow strands that ended with `get`. A strand u in function G that ends with $h = \text{fut-create}(F)$ has two outgoing edges — one non-SP edge to the first strand in F , and one SP edge to the continuation in G . We say that u is the **creator** of F denoted by $creator(F_j)$. The first strand of F is the source of a new SP-dag which contains all strands of F and the functions it calls (recursively) using `spawn` and the last strand of F is the sink of this SP-dag. Similarly, strand u in H immediately follows a `get(h)` call where h is the future handle for future F has two incoming edges — one SP edge from the strand that ended with the `get` call in the current function H and one non-SP edge from the last strand of the F . We say u is the **getter** of F_j denoted as $getter(F)$. We say that $u <_{SP} v$ if there is a path from u to v using only SP edges.

This model is quite general and subsumes computations that can arise from futures [3, 12–14, 23, 26, 33, 38] or other future-like (such as “put” and “get” [11, 56]) parallel constructs proposed in the literature. Therefore, our algorithm would work on all of these primitives.

Structured futures: We place the same restrictions on structured futures as in prior work [27]: (1) **Single-touch:** Every future handle is called with `get` at most once. (2) **No race on future handles:** There is a sequential dependence in the program from the point where a future is created (via `fut-create` which initializes a future handle) to the point where it is read (via `get`). More precisely, if strand u terminates with a $f = \text{fut-create}(F)$ call and strand v terminates with a `get(f)` call, then $u < v$ in the computation.

Eager execution: Both our algorithms execute the computation sequentially and execute the program in **depth-first eager** order. When the execution reaches `fut-create(F)` (after executing `creator(F)`) call or a `spawn(F)` call, it always executes the function F . When F returns, then the next node of the parent function (the one after the continuation edge) is executed. This execution order automatically has the property that all functions that must join at a sync point have already returned when the execution reaches the sync; therefore, the execution never blocks at a sync. Similarly, for structured futures, this execution has the property that the execution will never block at a `get`. For general futures, we restrict our attention to computations where the use of futures is **forward-pointing**: for every future F , $creator(F)$ executes before $getter(F)$ in the depth-first eager execution. Without this restriction, sequential execution of the original program could deadlock, in which case our algorithm race detects up to the point where it deadlocks.

²For simplicity, our analyses assume that each join has exactly two incoming edges. One can modify the algorithms to the more general case easily.

3 Managing Access History

As mentioned in Section 1, there are two important components in a race detector: access history and reachability data structure. MultiBags and MultiBags+ differ in how they maintain reachability but manage access history similarly. This section discusses how they manage access history — for each memory location ℓ , the access history maintains enough information about the previous accesses to ℓ so that future accesses to ℓ can detect races.

When race detecting a series-parallel program, it is sufficient to store a constant number of previous reader strands and a single previous writer strand in the access history [19, 39]. When a strand s accesses a memory location ℓ , it checks if some subset (based on whether s is reading or writing) of ℓ 's previous accessors are in parallel with s . Therefore, each memory access leads to at most a constant number of queries into the reachability data structure.

This property no longer holds for programs with futures, however. The access history for a memory location ℓ still holds only one writer strand, namely the most recent writer strand, $\text{last-writer}(\ell)$. However, it must now store an arbitrarily number of readers. Race detection proceeds as follows. Whenever a strand s reads from a memory location ℓ , the detector checks the reachability data structure to determine whether s is logically parallel with $\text{last-writer}(\ell)$; if so, a race is reported. Otherwise, s is added to $\text{reader-list}(\ell)$. When a strand s writes to a memory location ℓ , the race detector must check s against all readers in $\text{reader-list}(\ell)$ and with $\text{last-writer}(\ell)$. If s is in parallel with any of them, then it declares a race. Otherwise, the $\text{reader-list}(\ell)$ is set to \emptyset and s is stored as $\text{last-writer}(\ell)$. We can empty the $\text{reader-list}(\ell)$ without missing any races because anything that executes later that would be in parallel with these readers must also be in parallel with s (the new $\text{last-writer}(\ell)$), and a race will be reported with s .

A key thing to notice here is the following: the total number of queries into the reachability data structure (i.e., checking one access against another for race) is bounded by the total number of memory accesses in the computation. Since we can empty all the $\text{reader-list}(\ell)$ whenever we encounter a $\text{last-writer}(\ell)$, a reader is checked against some writer at most twice: when it is inserted into the reader list, and right before it is removed. Thus, the total number of queries to the reachability data structure is bounded by $O(T_1)$, where T_1 is the work of the computation. We shall relate this observation formally to the performance bound of MultiBags and MultiBags+ later in Sections 4 and 5.

4 MultiBags for Structured Futures

We now describe MultiBags, which can race-detect programs with structured futures in time $O(T_1 \alpha(m, n))$ where T_1 is the work of the program, α is the inverse Ackermann's function, m is the number of memory accesses in the program

and n is the number of `spawn` and `fut-create` calls. Since the inverse Ackermann's function is a very slowly growing function, the bound is close to optimal.

Notation: Note that programmatically, `spawn` and `sync` are subsumed by `fut-create` and `get` since we can convert a `spawn` to `fut-create` and `sync` to a series of `get` calls, one on each function spawned in the current function scope. In the case of general use of futures discussed in Section 5, we distinguish between SP edges generated by `spawn` and `sync` and non-SP edges generated by `fut-create` and `get` since the bound depends on k , the number of `get` calls, and converting all `sync` calls to `get` calls will increase this number. For structured futures, however, the bound does not depend on k ; therefore, for simplicity in this section, we assume that we only have `fut-create` and `get` constructs to create parallelism.

The computation dag consists of three kinds of nodes — regular strands with one incoming and one outgoing edge, **creator** strands which end with a `fut-create` call with two outgoing edges, and **getter** strands (the continuation after a `get` call) with two incoming edges. It also consists of three kinds of edges: **spawn** edges are edges from creator nodes to the first strand of the future; **join** edges are edges from last strand of a future to getter nodes; all other edges (between strands of the same function instance) are **continue** edges.

4.1 Algorithm

This algorithm is similar to the SP-Bags algorithm for detecting races for series-parallel programs [19]. As with SP-Bags, we use a fast disjoint-set data structure [55] to maintain a dynamic collection D of disjoint sets with three operations:

- $A = \text{MAKE-SET}(D, x)$: Creates a new set $A = \{x\}$ and adds it to the disjoint-set data structure D .
- $A = \text{UNION}(D, A, B)$: Unions the set B into A and destroys B . We will sometimes overload notation and say $\text{UNION}(D, x, y)$ where x and y are elements in the set instead of sets. This means that we union the set containing y into the set containing x .
- $\text{FIND}(D, x)$ returns the set that contains the element x .

In this section, we only have one disjoint-set data structure and thus D is implicit.

Like SP-Bags, MultiBags depends on the depth-first eager execution of the computation. MultiBags maintains a bag (a set in the data structure) for each function instance F which has been created and for which `get` has not yet been called (these bags can be stored with the future handle). This bag is labeled either as an S -bag, represented by S_F or a P -bag, represented by P_F . The algorithm maintains these bags as shown in Figure 1. The strands of a particular function F are always added to S_F before they execute.

The algorithm looks similar to SP-Bags [19]. The main difference is that when the function G returns, its S -bag S_G is renamed as P_G bag; in SP-bags, S_G would be unioned with

```

F calls  $f = \text{fut-create}(G)$  where  $u$  is the first strand of  $G$ :
1  $S_G = \text{MAKE-SET}(u)$ 
G returns to F:
2  $P_G = S_G$ .
F calls  $y = \text{get}(f)$  where  $f$  is  $G$ 's handle:
3  $S_F = \text{UNION}(S_F, P_G)$ 
// Called when strand  $v$  accesses memory location  $\ell$ 
// previously accessed by  $u$  in a conflicting way:
QUERY( $u, v$ ) // return TRUE iff  $u < v$ 
4 if FIND( $u$ ) is an S bag, return TRUE
5 else return FALSE

```

Figure 1. Pseudocode for MultiBags. The top part shows how MultiBags maintains the S and P bags when it encounters future constructs. The bottom part shows the operation of checking for races upon a memory access.




Figure 2. An example execution of MultiBags on a program with structured use of futures. In this program, there is always a sequential dependence between each future's creator and its corresponding getter's immediate predecessor in the same function (e.g., B 's creator and getter are 1 and 16 respectively, and the immediate predecessor of the getter in the same function is 15). This dag is not a series-parallel dag, as the spawning and joining of function instances are not well nested. The table shows the state of the disjoint-set data structure for maintaining reachability immediately before the execution of each strand in the order of the execution.

P_F , the parent function of G . However, this small difference is crucial for handling programs with structured futures rather than series-parallel programs. In addition, the proof for correctness of this algorithm for structured futures is significantly different than the proof of correctness of SP-bags for series-parallel programs.

Figure 2 shows the operation of this algorithm on an example program which uses structured futures. Each rectangle is a function instance and nodes are strands. The straight dashed lines going towards the left represent `fut-create` edges while the curved dashed lines represent `get` edges. Consider step 12 when the first node of function F is executing. All nodes except node 4 are sequentially before this strand and are correspondingly in some S -bag. Node 4 is in parallel with this strand and is in a P -bag.

4.2 Proof of Performance and Correctness

The performance bound of MultiBags and its proof is similar to that of SP-Bags, but the correctness analysis is quite different, which we will focus most of our attention on.

Proof of Performance of MultiBags

Theorem 4.1. *The running time of MultiBags when detecting races for a program with work T_1 is $T_1\alpha(m, n)$ where m is the number of memory accesses and n is the number of `fut-create` calls.*

Proof. The fast disjoint-set data structure provides the bound of amortized time $O(\alpha(m, n))$ per operation where m is the number of operations and n is the number of sets. For our program, m is at most the number of memory accesses and n is the number of `fut-create` calls. Note here, again, that unlike series parallel computations, each write may generate multiple queries; however, for the reason as explained in Section 3, the total number of queries is bounded by the two times the total number of memory accesses since each writer removes the entire reader-list. Therefore, the total running time is $O(T_1\alpha(m, n))$. \square

Intuition for the Correctness Proof

In order to argue that MultiBags is correct, we must prove the following theorem.

Theorem 4.2. *If the currently executing strand is v , then a previously executed strand u is currently in an S bag iff $u < v$.*

In order to prove this theorem, we define two more terms. A node u is a **spawn predecessor** of a node v if there is a path from u to v which consists of only spawn and continue edges. A node u is a **join predecessor** of v if there is a path from u to v that consists of only join and continue edges. Spawn and join successors are defined in the symmetric way. We will overload notation and say that a strand u is a spawn predecessor of a function F if there is a path from u to the first strand of F that consists of only spawn and continue edges and similarly a strand v is a join successor of F if there is a path from the last strand of F to v . Each node is its own spawn and join predecessor and successor.

The algorithm works due to the following observations.³ We say that a function is **active** if it has started executing, but has not completed (returned). While a function F is active, S_F exists and P_F doesn't because P_F is only created upon F 's return. All strands of an active function F are in S_F . After a function has returned, S_F is destroyed; P_F exists if `get` has not been called on F 's future handle and if P_F exists, then all strands of F are in P_F . After a function has joined (`get` has been called) then neither S_F nor P_F exist.

Property 1 is a property of eager executions.

Property 1. *When a strand v is currently executing, all spawn predecessors of v are part of some active function. The converse is also true; all strands w that are part of active functions are spawn predecessors of v .*

We need two other properties. The first one is a static property of paths in a program with structured futures as follows. If there is a path from a strand u to strand v ,⁴ then there must be a path where the first (possibly empty) part of the path consists of only join and continue edges, while the second (possibly empty) part of the path contains only spawn and continue edges. In other words, there is a path where no join edges follow spawn edges. More formally, for any two nodes u and v , if $u < v$, then we can find a node w where u is a join-predecessor of w and w is a spawn predecessor of v . Combining with Property 1, we get:

Observation 1. *Consider a completed strand u and a currently-executing strand v where $u < v$. The furthest join successor of u , say w , must be part of an active function.*

This observation allows us to concern ourselves only with paths that go through nodes of active functions. In particular, to detect races, it is sufficient to try to check if the furthest join successor w of any previously executed node u is part of an active function. If so, the observation implies that $u < w$ and we already know from Property 1 that $w < v$; therefore, $u < v$. If not, then $u \not< v$.

The second property is a dynamic property of MultiBags which allows to precisely check this. In particular, it states the following (which combined with the previous observation gives us the theorem):

Observation 2. *Consider an already completed node u . Say, at time t , u 's furthest join successor w is part of a function G . If G is active, then u is in G 's S bag, otherwise u is in G 's P bag.*

An informal argument about why this property is true follows: In Line 3, MultiBags unions P_G into S_F when G joins with an active function F . This suggests the following: Consider a particular strand u in function G . Say at time t , w is the furthest join successor strand of u which has executed

³We implicitly assume that when we refer to any strand (or function), it is either currently executing or has already executed (we have no knowledge of strands or functions that are still to execute).

⁴In general, there can be many paths from u to v .

(or is executing). (This strand is well defined since there is no branching in a path that contains only join and continue edges.) Say w is part of function F . Then at time t , u is in F 's bag. F 's bag is an S bag if F is active. Therefore, to check if u 's furthest join successor w is active, it suffices to check if u is in an S bag, which is precisely what MultiBags does.

If we combine Observations 1 and 2, we get the theorem. Say v is executing at time t and consider a previously executed node u and say w is the furthest join successor of u . If $u < v$, then by the first observation, w is part of an active function and therefore, by the second observation, u is in an S bag. On the other hand, if $u \parallel v$, then w can not be part of an active function — if it was, then by Property 1, $w < v$ and therefore, $u < v$ (a contradiction). Therefore, u is in a P bag.

When the computation does a memory access, it simply checks if the previous accesses are in an S . In order to do queries, the algorithm operates as follows:

1. Currently executing strand s reads location ℓ : check if `last-writer(ℓ)` is in a P bag; if so, declare a race. Otherwise, append r to `reader-list(ℓ)`
2. Currently executing strand s writes to location ℓ : check if any reader $r \in \text{reader-list}(\ell)$ is in a P bag; if so, declare a race. Otherwise, empty the reader list and set `last-writer(ℓ) = w` .

Proof of Correctness of MultiBags

First, we prove the static property stated in Observation 1. In order to prove it, we first define a canonical order on the futures of the program and show that we can always order the functions in this canonical order.

Lemma 4.3. *We can always find a canonical order.*

Proof. The order starts with the main function; we then progressively add futures to the computation. A future F can only be added if its creator strand and getter strand have already both been added.

We will induct on adding new futures. We can always start since the main function is added first. At some point, we have already added some futures, say a set S . Next, we will add the future F such that there is no future G such that `creator(F) < creator(G)` where both creator nodes are in S . There must be some such future (if there are more than one, we can choose arbitrarily). Since no future was created sequentially after F was created, and the strand before the `getter(F)` is sequentially after `creator(F)`, `getter(F)` must also be in S ; therefore, it is legal for us to pick F as the next future in our canonical order. \square

We now use this canonical order to induct on the futures. In particular, we can show that the static property stated in Observation 1 (stated more formally and completely in the following lemma) by inducting on the futures in the canonical order.

Lemma 4.4. *If $u < v$, then there exists a path from u to v that contains two sections: the first path (possibly empty) contains only join and continue edges and the second part (possibly empty) contains only spawn and continue edges. In other words there is never a spawn edge followed by a join edge on this path. In addition, this path is unique. Therefore, If $u < v$, then there is some node w (possibly u or v) which is a join successor of u and a spawn predecessor of v .*

Proof. Induct on futures in the canonical order (which we can always find according to Lemma 4.3) and show that this is true as we add futures one by one.

Base case: We first have only the main strand, so this is true trivially.

Inductive case: Assume that after we have added a set S of futures, the statement is true. We now add a new future F' .

Consider any nodes u and v in this new dag where $u < v$. If neither u nor v are in F' , then the addition of F' does not add any new paths between u and v (since the only new path added is between $\text{creator}(F')$ and $\text{getter}(F')$ and there was already a path between them before we added F'). In addition, any new path added does have a spawn followed by a join — therefore, the uniqueness is preserved. Therefore, we only need consider pairs where either u or v are in F' . If u is in F' and v is not, then the path from u to v must go from the last strand of F' to $\text{getter}(F')$ and then to v . By inductive hypothesis, the path from $\text{getter}(F')$ already follows the desired property and the path from u to $\text{getter}(F')$ only contains join and continue edges. Therefore, the property still holds. A symmetric argument applies when v is in F' . \square

We then prove the dynamic property stated in Observation 2 by looking at the execution as it unfolds. For each function F , we define its **operating function** G as the function containing the “furthest join descendant” of the last executed strand of F . An active function is its own operating function. If a function is not active (it has returned), it may be **confluent** if its operating function is active; otherwise it is **non-confluent**. By definition, the operating function of a non-confluent function is always non-confluent. A confluent function can never be its own operating function, but a non-confluent function F may be its own operating function if its $\text{getter}(F)$ has not yet executed.

The following lemma is proved by induction on the program as it executes.

Lemma 4.5. (a) *When a function F is active, all its strands are in its S bag.* (b) *If a function F is confluent, then all its strands are in its operating function G 's S bag.* (c) *If a function F is non-confluent, then all its strands are in its operating function G 's P bag.*

Proof. When a function is first called, it has an S bag and its strands are placed in the S bag. They remain in this S bag while it is active. Once the function returns, all its items

move to a P bag. For the other two statements, we induct on time after F returns.

Base Case: When F returns, $\text{getter}(F)$ has not yet been called. Therefore, it is its own operating function; it is non-confluent; and all its strands are in its own P bag.

Inductive Case: We will do this by two cases:

Case 1: F is non-confluent and G is its (non-confluent) operating function; by inductive hypothesis, all strands of F are in G 's P bag. The only thing that can change the location of its strands is if $\text{getter}(G)$ executes, say by function H . At this point, H (currently active) becomes the operating function for both G and F — therefore, F is now confluent. All strands of F (and incidentally G) move to H 's S bag.

Case 2: F is confluent and G is its (active) operating function; by inductive hypothesis, all strands of F are in G 's S bag. The only thing that changes the location of F 's strands is if G returns. At this point G becomes non-confluent (since it is no longer active); therefore F also becomes non-confluent. All of F 's strands move to G 's P bag. \square

The combination of static and dynamic properties leads to the proof of correctness. The intuition is that if a function F is confluent, then there is some strand w in its (active) operating procedure which is a join successor of all strands of F and a spawn predecessor of currently executing strand.

Theorem 4.2. *If the currently executing strand is v , then a previously executed strand u is currently in an S bag iff $u < v$.*

Proof. By Lemma 4.4, we know that if $u < v$, then we can find a node w such that u is a join predecessor of w and w is a spawn predecessor of v . By Property 1, since v is executing, the function containing w , say G , is still active. Therefore, by definition, the function containing u is confluent. Therefore, by Lemma 4.5, u is an S bag.

If u does not precede v , then there is no path from u to v , and u cannot have a path to any strand w in any active function (otherwise by Property 1, since w has path to v , u will also have a path to v). Thus, by definition, u is non confluent. By Lemma 4.5, u is in a P bag. \square

5 MultiBags+ for General Futures

We now consider general use of futures for programs that use both spawn/sync constructs and also futures. In particular, we consider programs where most of the parallelism is created using spawn and sync, but there are also k future get operations. For these programs, we provide a race detection algorithm that runs in total time $O(T_1\alpha(m, n) + k^2)$, where T_1 is the work of the program, α is the inverse Ackermann's function, m is the number of memory accesses in the program and n is the number of spawn and fut-create calls. To put this bound in context, a series-parallel program has $k = 0$ — in this case (and in fact, for any program where $k = O(\sqrt{T_1})$), the MultiBags+ runs in time $O(T_1\alpha(m, n))$. Since the inverse

Ackermann’s function grows slowly (upper bounded by 4), this bound is close to asymptotically optimal.

As mentioned in Section 2, MultiBags+ depends on eager execution of the computation and we assume that our futures are forward-pointing. Therefore, the depth-first execution never blocks on a get call since the corresponding future has already finished executing.

Notation: Unlike in Section 4, we must distinguish between *spawn* and *fut-create* (similarly, between *sync* and *get*) for MultiBags+. The computation dag consists of five kinds of nodes: (1) regular strands with one incoming and one outgoing edge; (2) *spawn* strands which end with a *spawn* instruction and have with two outgoing edges; (3) *creator* strands which end with a *fut-create* instruction and have with two outgoing edges; (4) *sync* strands which begin immediately after a instruction and have with two incoming edges; and (5) *getter* strands which begin immediately after *get* instruction and have with two incoming edges. Some strands can have two incoming and two outgoing edges (if they start immediately after a *get* or *sync* instruction and end with a *spawn* or *fut-create*); these strands are correspondingly in both categories.

The computation dag also consists of five kinds of edges: *spawn* edges are from *spawn* nodes to the first strand of the corresponding spawned function; *join* edges are from last strand of a spawned function to the corresponding *sync* node; *create* edges from the creator strand to the first stand of the future function; and *get* edges from the last strand of a future function to the corresponding *getter* node. Each future can have multiple *get* edges if it is a multi-touch future.⁵

Reachability data structures: Recall that we can model computations that employ futures as a set of series-parallel dags (SP dags) plus some non-SP edges (Section 2). When we need to check if $u < v$, if they are already in the same SP dag (i.e., $\text{SP-Dag}(u) = \text{SP-Dag}(v)$), as defined in Section 2), the disjoint-set data structure maintained by MultiBags can readily answer the reachability query correctly. We only run into trouble due to use of general futures when $\text{SP-Dag}(u) \neq \text{SP-Dag}(v)$ but a path exists between them via (possibly more than one) non-SP edges.

Thus, MultiBags+ maintains two data structures, and both utilize fast disjoint-set data structure described in Section 4. The first disjoint-set data structure, called D_{SP} , is virtually identical to the data structure used in Section 4: we maintain S and P bags for each function, and the operations performed on D_{SP} are identical to those described in Section 4 in the case of *fut-create*, and *spawn* is treated in the same way as *fut-create*. In addition, *sync* is treated like *get*. The only difference is that we do not perform any operation on this data structure on *get* (since we allow multi-touch futures). We can use D_{SP} to correctly answer reachability

query between two strands if they are in the same SP dag. Intuitively, for reasons similar to the structured case, all strands that are currently stored in an S bag are sequentially before the currently executing strand. Note that all strands that are stored in a P bag are not necessarily in parallel with the current strand, due to non-SP edges — we will use the second data structure to answer that query.

The second data structure handles the additional complication in reachability query when two strands are connected via non-SP edges. This data structure has two components: (1) a disjoint-set data structure called D_{NSP} that maintains a collection of disjoint-set, and each strand is added to D_{NSP} when encountered; and (2) a separate dag called \mathcal{R} that contains some of the sets from D_{NSP} . The high-level idea is that these sets are made of connected series-parallel subdags of the original dag G_{full} . For any two nodes u and v in different SP dags, MultiBags+ ensures that $u < v$ in G_{full} iff $\text{FIND}(D_{NSP}, u) < \text{FIND}(D_{NSP}, v)$ in \mathcal{R} (the sets they are in are connected in \mathcal{R}).

We call sets also in \mathcal{R} as the **attached sets**, which store nodes that are subdags which start and/or end with *creator* or *getter* strands.⁶ \mathcal{R} explicitly maintains reachability relationship that arises due to non-SP edges between nodes in the attached sets. \mathcal{R} is simply a dag (with each node being an attached set), but it is not series-parallel. Thus, to answer reachability queries quickly between nodes in \mathcal{R} , MultiBags+ maintains a full transitive closure of all sets in \mathcal{R} — whenever a set is added to \mathcal{R} , its reachability from all sets already added to \mathcal{R} is explicitly computed and stored. Therefore, one can check if $A < B$ in \mathcal{R} in constant time.

If every set could be in \mathcal{R} we would be done. We must keep \mathcal{R} small, however, since every time we add a set to \mathcal{R} we compute a full transitive closure, which is expensive. It turns out that it is difficult to simultaneously put all strands in attached sets and keep \mathcal{R} small. In order to cope with this, some strands are in **unattached sets**, which are only stored in D_{NSP} . Intuitively, an unattached set contains nodes of a complete series-parallel subdag which have no incident non-SP edges. Each unattached set U has two additional fields, **attached predecessor** and **attached successor**, which point to attached sets that act as U ’s proxies when querying \mathcal{R} . U ’s attached predecessor, denoted as $U.attPred$, is set when U is created; therefore, it always points to some attached set. U ’s attached successor, denoted as $U.attSucc$, is set at some later point; it either points to some attached set or may be null. An attached set is always its own attached predecessor and successor. We will overload notation and say that node u ’s attached predecessor is $\text{FIND}(D_{NSP}, u)$ ’s attached predecessor (and similarly for attached successor).

Answering queries: Figure 3 shows how the reachability data structures are queried to find out if a path exists between

⁵For context, in Section 4, both *spawn* and *create* edges were called *spawn* edges and both *join* and *get* edges were called *join* edges.

⁶This is not quite accurate; for technical reasons, some attached sets start/end with regular, *spawn*, and *join* nodes as well.

QUERY(u, v) // return TRUE iff $u < v$ in G_{full} 1 if FIND(D_{SP}, u) is an S -bag, // Query D_{SP} first 2 return TRUE 3 $S_v = \text{FIND}(D_{NSP}, v)$ 4 if S_v is unattached 5 $S_v = S_v.attPred$	6 $S_u = \text{FIND}(D_{NSP}, u)$ 7 if S_u is unattached 8 $S_u = S_u.attSucc$ 9 if $S_u = \text{NULL}$ return FALSE 10 $ans = \text{query } \mathcal{R} \text{ to determine if } S_u < S_v$ 11 return ans
--	---

Figure 3. Code for querying reachability.

u is the first strand of the computation: 1 add u to an attached set with no predecessor. <hr/> Function F calls spawn (G): // u is the strand in F immediately before the spawn // v is the strand in F right after spawn // w is the first strand of G 2 $S_G = \text{MAKE-SET}(D_{SP}, w)$ 3 $U_v = \text{MAKE-SET}(D_{NSP}, v)$ and make U_v unattached. 4 $U_v.attPred = \text{FIND}(D_{NSP}, u).attPred$ 5 $U_w = \text{MAKE-SET}(D_{NSP}, w)$ and make U_w unattached. 6 $U_w.attPred = \text{FIND}(D_{NSP}, u).attPred$ <hr/> Function F calls fut-create (G): // u is the strand in F immediately before the fut-create // v is the strand in F immediately after the fut-create // w is the first strand of G 7 $S_G = \text{MAKE-SET}(D_{SP}, w)$ 8 ATTACHIFY (u) 9 $A_v = \text{MAKE-SET}(D_{NSP}, v)$ and make A_v attached. 10 Add an arc from $\text{FIND}(D_{NSP}, u)$ to A_v in \mathcal{R} . 11 $A_w = \text{MAKE-SET}(D_{NSP}, w)$ and make A_w attached. 12 Add an arc from $\text{FIND}(D_{NSP}, u)$ to A_w in \mathcal{R} <hr/> Function G returns: 13 $P_G = S_G$; deallocate S_G <hr/> Function F calls get (G): // u is the strand in F that ended with the get // v is the strand immediately after u in F // w is the last strand of G . 14 ATTACHIFY (u) 15 $A_v = \text{MAKE-SET}(D_{NSP}, v)$ and make A_v attached. 16 Add an arc from $\text{FIND}(D_{NSP}, u)$ to A_v in \mathcal{R} 17 Add an arc from $\text{FIND}(D_{NSP}, w)$ to A_v in \mathcal{R} ; // $\text{FIND}(D_{NSP}, w)$ is guaranteed to be attached.	ATTACHIFY (u) // make the set containing u attached if not already. 18 $U_u = \text{FIND}(D_{NSP}, u)$ 19 if U_u is unattached 20 mark U_u as attached 21 add U_u to \mathcal{R} 22 add the arc ($U_u.attPred, U_u$) to \mathcal{R} <hr/> Function F calls sync with child function G : 23 $S_F = \text{UNION}(D_{SP}, S_F, P_G)$; deallocate P_G 24 look at the corresponding fork 25 let f be the strand immediately preceding the fork 26 let s_1 and s_2 be f 's two immediate successors of f // i.e., the first strand of G and the first strand of the continuation 27 let j be the strand immediately after the sync 28 let t_1 and t_2 be j 's immediate predecessors of the sync // i.e., the last strand of the G and the continuation 29 if neither $\text{FIND}(D_{NSP}, t_1)$ nor $\text{FIND}(D_{NSP}, t_2)$ is attached // No non-SP edges 30 $\text{UNION}(D_{NSP}, f, t_1)$ 31 $\text{UNION}(D_{NSP}, f, t_2)$ 32 $\text{UNION}(D_{NSP}, f, \text{MAKE-SET}(j))$ 33 elseif both $\text{FIND}(D_{NSP}, t_1)$ and $\text{FIND}(D_{NSP}, t_2)$ are attached 34 ATTACHIFY (f) 35 add arc ($\text{FIND}(D_{NSP}, f), \text{FIND}(D_{NSP}, s_1)$) to \mathcal{R} 36 add arc ($\text{FIND}(D_{NSP}, f), \text{FIND}(D_{NSP}, s_2)$) to \mathcal{R} 37 $A_j = \text{MAKE-SET}(D_{NSP}, j)$ and make A_j attached 38 add a node A_j to \mathcal{R} 39 add arc ($\text{FIND}(D_{NSP}, t_1), A_j$) to \mathcal{R} 40 add arc ($\text{FIND}(D_{NSP}, t_2), A_j$) to \mathcal{R} 41 else let t_a be the attached one and t_u be the unattached one // correspondingly s_a is attached and s_u is unattached 42 if $\text{FIND}(D_{NSP}, f)$ is not attached 43 $\text{UNION}(D_{NSP}, s_a, f)$ 44 $\text{UNION}(D_{NSP}, t_a, \text{MAKE-SET}(j))$ 45 $\text{FIND}(D_{NSP}, t_u).attSucc = \text{FIND}(D_{NSP}, j)$ 46
---	---

Figure 4. The actions taken by the algorithm to maintain D_{SP} , D_{NSP} and \mathcal{R}

some previously executed node u and the currently executing node v . In the first part of the query (lines 1–2), we query D_{SP} and if u is in the S bag, then we can conclude that $u <_{full} v$ and return. If u is in the P bag, then we check if attached successor of u precedes the attached predecessor of v in \mathcal{R} ; if so, we say that $u <_{full} v$. Otherwise u is in parallel with v .

Maintaining D_{SP} , D_{NSP} , and \mathcal{R} : Figure 4 shows the code for maintaining reachability relationships between nodes in the computation. The first thing we do during a **spawn**, **fut-create**, **return** and **sync** is to manipulate D_{SP} (lines 2, 7, 13, and 23) in a manner identical to Section 4.⁷

Now let's consider the manipulations of D_{NSP} and \mathcal{R} . It uses an auxiliary function **ATTACHIFY**(u), which simply

checks if $U_u = \text{FIND}(D_{NSP}, u)$ is an unattached set, and if so, converts it into an attached set by adding it to \mathcal{R} and adding an edge from $U_u.attPred$ to U_u in \mathcal{R} .

The attached and unattached sets change as the execution continues. MultiBags+ unions sets in D_{NSP} growing both attached and unattached sets. Two attached sets are never unioned together. Whenever we union an attached set and an unattached set, we always union the unattached set into the attached set; therefore, the resulting set is attached and remains in \mathcal{R} . On the other hand, an unattached set contains nodes of a complete series-parallel subdag which have no incident non-SP edges. In particular, consider a parallel composition of two series-parallel subdags G_1 and G_2 . Say G_1 has no incident non-SP edges. Then all nodes of G_1 constitute

⁷Since we are assuming binary forking, we sync with one function at a time.




Figure 5. An example of general futures. The left figure shows the full dag. The dashed edges are fut-create and get edges. The numbering shows the order in which the nodes execute in depth-first eager execution. The orange ovals (solid outline) are attached sets while the green ovals (dashed outlines) are unattached. The right figure shows \mathcal{R} at the end of the computation. The thick blue lines indicate an attached predecessor relationship while a thick red line indicates an attached successor relationship.

an unattached set if either the join node that joins G_1 and G_2 has not executed yet, or if G_2 has an incident non-SP edge. At a high-level, this design allows us to have the property that each non-SP edge only leads to a constant number of attached sets. Since only attached sets are in \mathcal{R} , this idea allows us to keep \mathcal{R} small, allowing us to get good performance.

Figure 5 shows a computation dag with futures and its corresponding \mathcal{R} at the end of the computation. As can be seen in this example, sets X , Y and Z are all unattached and each have the above property.

Due to space limitations, all the proofs are omitted to the arXiv version of the paper⁸. Here we simply state the performance and correctness theorems.

Theorem 5.1. *MultiBags+ detects races in time $O(T_1\alpha(m,n) + k^2)$ for programs with T_1 work, k get calls, m memory accesses, and n number of strands.*

Theorem 5.2. *If the program is executed in a depth first eager execution order, $\text{QUERY}(u, v)$ returns true iff $u < v$ in G_{full} .*

Comparison to Algorithm by Agrawal et al. [1]

As we have fully described the MultiBags+ algorithm, we discuss the differences between MultiBags+ and the the-of-the-art algorithm by Agrawal et al. [1] and provide an analytical analysis as to why the algorithm by Agrawal et al. is much more challenging to implement in practice.

The algorithm by Agrawal et al. utilizes the following data structures to answer reachability queries: 1) an order-maintenance data structure for answering series-parallel queries; 2) the full computation DAG to update and maintain “anchor-predecessors” and “proxies” used to infer “anchor-successors;” and 3) a reachability matrix \mathcal{R} which contains

anchor nodes to answer reachability queries involving non-SP edges. The functionalities served by these data structures are similar to that of D_{SP} , D_{NSP} , and \mathcal{R} in MultiBags+; in particular, their algorithm utilize anchor-predecessors and anchor-successors to allow for correct reachability queries involving non-SP edges, sharing similar roles as the attached predecessors and attached successors in MultiBags+. The main difference is in the second data structure and how the anchor-predecessors and anchor-successors are maintained.

In the algorithm by Agrawal et al., the mechanism for maintaining anchor-predecessors and proxies (which are used to infer anchor-successors) are more complex. In particular, to maintain anchor-predecessors, the algorithm maintains the full computation dag, and each strand (a node in the dag) explicitly stores its anchor-predecessor. However, anchor-predecessors can sometimes change as the program executes. When that occurs, the algorithm must explicitly traverse subpart of the dag and update some of the predecessors explicitly. The asymptotic complexity of such updates is still ok because the paper argues that a strand’s anchor predecessor can only change a constant number of times.

Similarly, the algorithm maintains a proxy per strand, used to infer a strand’s anchor-successor. A proxy for a strand is stored instead of its anchor-successor is because, while an anchor-predecessor of a node can change a constant number of times, its anchor-successor can change many times. Thus instead, the algorithm maintains a proxy, which indirectly allows the algorithm to deduce its anchor-successor. Like the anchor-predecessor, a proxy of a node can only change a constant number of times, and when that occurs, the algorithm again explicitly traverses the relevant subdag and updates the proxies explicitly.

We argue that this algorithm is harder to implement and likely has higher overheads due to the following reasons. First, explicitly maintaining the entire program dag and also storing each strand’s anchor-predecessor and proxy would be more memory intensive than keeping these strands in union-find data structures which are tagged appropriately. Second, explicit dag traversals in order to update proxies and anchor-predecessors of nodes would be expensive (even though the asymptotic complexity is manageable). This prior work establishes the state-of-the art time bound for race detecting programs that use general futures, but no implementation exists.

6 Experimental Evaluation

This section empirically evaluates FutureRD that implements MultiBags and MultiBags+ described earlier. We first evaluate the practical efficiency of these algorithms and then the performance difference between them, focusing on the impact of the additional k^2 overhead that MultiBags+ incurs, where k is the number of get operations.

⁸ Available at <https://arxiv.org/abs/1901.00622>

Implementation of FutureRD

FutureRD works by instrumenting parallel program executions: upon the execution of a parallel construct (i.e., `spawn`, `sync`, `fut-create`, and `get`), it invokes the necessary operations to update the reachability data structures; likewise, upon the execution of a memory access, it invokes the necessary operations to update the access history data structure and query both data structures.

We use Intel Cilk Plus [28] as our language front end, which is a C/C++ based task parallel platform that readily supports fork-join parallelism. Cilk Plus does not currently support the use of futures, however, so we have implemented our own future library. Since our race detector executes the program sequentially with eager evaluation of futures, the future library never actually interacts with Cilk Plus runtime during race detection.

Both MultiBags and MultiBags+ utilize disjoint-set data structures to maintain reachability as described in Section 4). MultiBags+ additionally needs to maintain \mathcal{R} as part of its reachability data structure (defined in Section 5). Conceptually, \mathcal{R} is simply a boolean reachability matrix where each cell (i, j) indicates whether there is a path from attached set i to attached set j . FutureRD maintains \mathcal{R} as a vector of bit vectors, representing the reachability between any two sets using a single bit. Whenever an edge is added to \mathcal{R} , reachability is transitively propagated via parallel bit operations.

FutureRD maintains the access history like a two-level direct-mapped cache, and keeps track of the reader list and last writer at four-byte granularity (all our benchmarks perform four-byte or larger accesses). That is, to query or update readers/writers for an address a , the more significant bits of a are used to index into the top-level table and the rest of the bits are used to index into the second-level table.

Experimental Setup

We evaluate FutureRD using six benchmarks: longest-common subsequence (`lcs`), Smith-Waterman (`sw`), matrix multiplication without temporary matrices (`mm`), binary tree merge (`bst`) as described by Blelloch and Reid-Miller [10], Heart Wall Tracking (`heartwall`), and Dedup (`dedup`). Heart Wall Tracking and Dedup both contain parallel patterns that cannot be easily implemented using fork-join constructs alone. The Heart Wall Tracking algorithm is adapted from the Rodinia benchmark suite [15] that tracks the movement of a mouse heart over a sequence of ultrasound images. Dedup is a compression program that exhibits pipeline parallelism [8], taken from the Parsec benchmark suite [7]. All but dedup have two implementations: structured and general futures; dedup does not utilize the flexibility of general futures. Figure 6 shows input and base case sizes of the benchmarks used for the experiments. The input and base case sizes shown are the default used for Figures 7 and 8. In these experiments, we use base case $B = \sqrt{N}$ for `lcs`, `mm`, and `sw`

to keep the work the same for the baseline, MultiBags, and MultiBags+ (since MultiBags+ has k^2 additional overhead). We vary the base case size for experiments shown in Figure 9.

We ran our experiments on an Intel Xeon E5-4620 with 32 2.20-GHz cores on four sockets. Each core has a 32-KByte L1 data cache, 32-KByte L1 instruction cache, a 256-KByte L2 cache. There is a total of 500 GB of memory, and each socket shares a 16-MByte L3-cache. All benchmarks are compiled with LLVM/Clang 3.4.1 with `-O3 -fno` running on Linux kernel version 3.10. Each data point is the average of 5 runs with standard deviation less than 5% with the exception of running dedup with full race detection, which sees a standard deviation under 9%.

Practical Efficiency of FutureRD

First, we evaluate the overhead of FutureRD and show that the algorithms can be implemented efficiently. To get the sense of where the overhead comes from, we ran the application benchmarks with four configurations:

- **baseline**: running time without race detection;
- **reachability**: running time with only the reachability components, including the instrumentation overhead to capture parallel control constructs;
- **instrumentation**: running time with memory-access instrumentation overhead on top of the reachability configuration, but does not maintain or query the access history;
- **full**: running time with the full race detection overhead.

Figure 7 shows the list of programs that employ structured futures running with different configurations, where FutureRD maintains reachability using the MultiBags algorithm. First, observe that the reachability configuration incurs almost no overhead, except for `bst`, which has very little work per parallel construct. Since the operations on the disjoint-set data structure are very efficient, as long as there is sufficient work per parallel construct, the overhead of maintaining reachability in MultiBags should be low. These program contains large number of memory accesses, however, and thus adding instrumentation for memory accesses alone incurs additional 2–4.5× overhead.

Going from the instrumentation configuration to the full race detection incurs another 6–10× overhead, with the exception of dedup. We expect the additional overhead incurred to be about 8–10× because the full configuration transforms every memory access into updates to access history and queries to both access history and reachability data structures. Thus, each memory access is translated into a few function calls and several pointer chases to multiple data structures. The benchmark `heartwall` only incurs additional 6×, because it spends non-negligible amount of time performing I/O (reading in image files). Finally, dedup is an outlier because dedup calls into a dynamic library to perform compression, which we could not recompile to include instrumentation. Thus, any memory accesses performed within

bench	N	base case	reads	writes	structured			general		
					fork-join ops	future ops	# strands	fork-join ops	future ops	# strands
lcs	16384	128	1.27e09	2.69e08	0	32768	49152	13	48897	65302
sw	2048	64	8.59e09	4.20e06	0	2048	3072	13	3009	4054
mm	2048	64	1.72e10	1.80e08	858	65536	99690	42142	65536	177902
heartwall	10 (images)	n/a	1.73e10	1.64e08	85	9180	13935	4261	10557	22242
dedup	"large"	n/a	8.20e08	1.89e07	0	375856	563784	0	375856	563784
bst	8e6 & 4e6	depth limit= 5	4.51e08	2.55e08	62	0	93	62	434	744

Figure 6. Input sizes and characteristics of benchmarks used.

bench	baseline	reachability	instr	full
lcs	2.19	2.23 (1.02×)	6.65 (3.04×)	54.27 (24.77×)
sw	14.78	14.25 (0.96×)	28.79 (1.95×)	325.10 (22.00×)
mm	13.94	13.82 (0.99×)	58.84 (4.22×)	468.75 (33.61×)
heartwall	13.86	13.77 (0.99×)	63.39 (4.58×)	340.04 (24.54×)
dedup	12.38	12.15 (0.98×)	13.79 (1.11×)	26.43 (2.14×)
bst	1.37	1.92 (1.41×)	2.65 (1.94×)	10.94 (8.02×)

Figure 7. The execution times for the benchmarks using structured futures, shown in seconds, with MultiBags used for race detection. Numbers in the parentheses show the overhead compared to the baseline.

bench	baseline	reachability	instr	full
lcs	2.03	2.30 (1.14×)	6.47 (3.19×)	54.95 (27.13×)
sw	14.73	14.65 (0.99×)	27.87 (1.89×)	380.19 (25.82×)
mm	13.13	15.07 (1.15×)	64.04 (4.88×)	498.65 (37.99×)
heartwall	13.82	13.89 (1.00×)	56.58 (4.09×)	487.95 (35.31×)
dedup	12.11	27.73 (2.29×)	29.60 (2.44×)	52.39 (4.33×)
bst	1.44	6.01 (4.16×)	6.79 (4.70×)	18.18 (12.60×)

Figure 8. The execution times for the benchmarks using general futures, shown in seconds, with MultiBags+ used for race detection. Numbers in the parentheses show the overhead compared to the baseline.

the library do not incur additional overhead. Since the compression takes up a substantial amount of execution time, the additional overhead is small.

Figure 8 shows the runtime of programs that employ general futures where FutureRD maintains reachability using the MultiBags+ algorithm. The additional overhead incurred going from one configuration to the next is similar to Figure 7 except the higher overhead from MultiBags+ is evident in the reachability configuration.

Over five benchmarks (excluding dedup, since we could not instrument its compression library), we see a geometric mean overhead of 1.06× and 1.40× to maintain reachability using MultiBags and MultiBags+, respectively. Full race detection exhibits 20.48× and 25.98× overhead, respectively.

Comparison between MultiBags and MultiBags+

Next, we compare the performance difference between MultiBags and MultiBags+. To evaluate the overhead difference between them, we run the same programs (i.e., with structured futures) with both algorithms. Although MultiBags+ is designed for general futures, it also works with programs

that use structured futures, albeit with an additional k^2 overhead, where k is the number of get calls.

For lcs, sw, and mm, k is dictated by how much the base case is coarsened — the smaller the base case, the more get calls, and the higher k is (which leads to higher overhead). Runtimes shown before used base case of $B = \sqrt{N}$ to keep the work asymptotically the same across baseline, MultiBags, and MultiBags+. Now we decrease the base case size below (i.e., increase k) to see how the overhead of MultiBags+ changes compared with the overhead of MultiBags.

bench	base case	baseline	reachability	
			MultiBags	MultiBags+
lcs	128	2.20	2.25 (1.02×)	2.39 (1.09×)
lcs	64	2.14	2.20 (1.03×)	4.68 (2.19×)
lcs	32	2.14	2.09 (0.98×)	39.82 (18.63×)
sw	64	14.57	14.13 (0.97×)	14.74 (1.01×)
sw	32	14.68	14.56 (0.99×)	14.25 (0.97×)
mm	64	13.83	13.97 (1.01×)	14.52 (1.05×)
mm	32	13.08	13.12 (1.00×)	49.11 (3.75×)

Figure 9. The execution times under the baseline and reachability configurations (both MultiBags and MultiBags+) for a subset of benchmarks implemented with structured futures. Numbers in the parentheses show the overhead compared to the baseline.

Figure 9 shows the measurements for running programs with structured futures using MultiBags and MultiBags+ in the reachability configuration with different base cases. The overhead difference between MultiBags and MultiBags+ can readily be observed in Figures 7 and 8 — compared to MultiBags, MultiBags+ incurs $2 + \times$ more overhead running dedup and $3 + \times$ more running bst for maintaining reachability. Here we show additional numbers for benchmarks where varying base case sizes changes k .

The measurements with lcs and mm bear out the extra overhead of MultiBags+. The lcs benchmark has $\Theta(n^2)$ work versus $(n/B)^2$ futures, while mm has more work ($\Theta(n^3)$), but also requires $(n/B)^3$ futures. With a higher ratio of futures to total work, the overhead is more apparent. Moreover, the memory required for the reachability matrix \mathcal{R} becomes substantial for small base cases, adding more overhead. The sw benchmark, however, has $\Theta(n^3)$ work compared to $(n/B)^2$ futures, so the effect of smaller base cases is small.

7 Related Work

Besides works discussed in Section 1, researchers have considered race detection for other structured computations. Dimitrov et al. [18] propose a sequential near-optimal race detection algorithm for two-dimensional dags which also exhibit nice structural properties. Subsequently, Xu et al. [60] propose a race detector for two-dimensional dags with asymptotically optimal parallel running time. Lee and Schardl [35] propose a sequential race detector for fork-join computations with reductions, where the computation dag is almost series-parallel except when reductions are performed.

Beyond task parallel code, there is a rich literature on race detection for programming models that generate nondeterministic computations, such as ones that employ persistent threads and locks. For such models, since the dag necessarily depends on the schedule, the best correctness guarantee that a race detector can provide is for a given program, for a given input, and for a given *schedule*. Early work [48, 59] employs lock-set algorithm, which provides wide coverage but can lead to many false positives, because it cannot precisely capture happens-before (HB) relations formed between threads.

A vector-clock (VC) based algorithm such as one proposed by Flanagan and Freund [22] can capture HB precisely for a given schedule. Such algorithm can be used on computation with arbitrary dependences, but naively applying it to task parallel code would be impractical, since it requires storing a VC of length n with each each memory location querying against it per access, incurring a multiplicative factor of n overhead on top of the work, where n is the number of strands, which can be on the order of millions.

In the context of race detecting nondeterministic code, researchers have investigated hybrid approaches incorporating VC and lock-set [41, 43, 49, 61] to trade-off precisions and coverage. More recently, researchers have proposed predictive analysis to explore alternative feasible schedules among close by instructions to increase the coverage (e.g. [31, 37, 47, 50]) while keeping the precision.

8 Conclusion

In this paper, we have shown that race detection can be performed more efficiently when the program employs only structured futures. As such, an interesting question to ask is how much benefit does the general use of futures provides. The flexibility of general futures lends itself to express parallelism that are not strictly fork-join and can potentially allow for slightly higher parallelism (such as in the case of pipelined binary-tree merge `bst` due to [10]). Moreover, we find that, writing code using structured futures can require more effort from the programmer, due to the fact that it needs to be single touch, such as in the case of dynamic programming examples `lcs` and `sw`. However, for most benchmarks,

these benefits are not as evident. As we have shown experimentally, the additional $O(k^2)$ overhead (k being the number of future operations) can impact performance in practice.

Currently, both of our algorithms execute the computation serially. An interesting avenue of future work is how to parallelize race detection for programs that use futures. Both of our algorithms depend on the depth-first execution order and extension to parallel execution appears to be non-trivial.

Acknowledgements

This research was supported in part by National Science Foundation under grants CCF-1150036, CCF-1527692, CCF-1733873, and XPS-1439062. We thank the referees and our shepherd for their excellent comments.

A Artifact Appendix

A.1 Abstract

Our implementation is open source and available at <https://github.com/wustl-pctg/futurerd.git>. The version used for artifact evaluation has DOI 10.5281/zenodo.2510564. The library is provided under an MIT license, though other code packaged with it is licensed separately. A brief summary of the artifact is provided below; more details can be found in the repository `README.md` file. Please send feedback or file bug reports by opening issues in the Github repository.

A.2 Software Dependencies

The headers installed by `gcc 6+` are likely to cause problems when compiling the compiler; we recommend using the Docker container provided (instructions in `README.md`). To fully reproduce the results, link-time optimization should be used (`-flto`) with the GNU gold linker installed as `ld`. The dedup benchmark requires several dependencies which are documented in the repository.

The benchmark script requires GNU `datamash`, which can be installed using `apt-get` in Ubuntu 14+ or can be obtained from <https://www.gnu.org/software/datamash>. Bash 4+ should be used to run the scripts.

A.3 Evaluation

After compiling all components and running the benchmark script, full results can be found in the files `times.ss.csv` (benchmarks used MultiBags race detection algorithm with structured futures), `times.ns.csv` (benchmarks used MultiBags+ algorithm with structured futures), and `times.nn.csv` (benchmarks used MultiBags+ algorithm with general futures). Although absolute times will differ on your machine, you should see similar relative overhead for the benchmarks. Compare the results to figures 7 and 8 in the paper. Comparison to figure 11 can be done by modifying the benchmark scripts and the base cases of the `lcs`, `sw`, and `mm` (called `matmul_z` in the repository).

References

- [1] Kunal Agrawal, Joseph Devietti, Jeremy T. Fineman, I-Ting Angelina Lee, Robert Utterback, and Changming Xu. 2018. Race Detection and Reachability in Nearly Series-parallel DAGs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '18)*. Society for Industrial and Applied Mathematics, New Orleans, Louisiana, 156–171. <http://dl.acm.org/citation.cfm?id=3174304.3175277>
- [2] Nimar S. Arora, Robert D. Blumofe, and C. Greg Plaxton. 1998. Thread Scheduling for Multiprogrammed Multiprocessors. In *10th Annual ACM Symposium on Parallel Algorithms and Architectures*. 119–129.
- [3] Arvind, R.S. Nikhil, and K.K. Pingali. 1986. I-structures: Data Structures for Parallel Computing. In *Proceedings of the Graph Reduction Workshop*.
- [4] Henry C. Baker, Jr. and Carl Hewitt. 1977. The incremental garbage collection of processes. *SIGPLAN Notices* 12, 8 (1977), 55–59.
- [5] Rajkishore Barik, Zoran Budimlić, Vincent Cavé, Sanjay Chatterjee, Yi Guo, David Peixotto, Raghavan Raman, Jun Shirako, Sağnak Taşlılar, Yonghong Yan, Yisheng Zhao, and Vivek Sarkar. 2009. The Habanero Multicore Software Research Project. In *Proceedings of the 24th ACM SIGPLAN Conference Companion on Object Oriented Programming Systems Languages and Applications (OOPSLA '09)*. ACM, Orlando, Florida, USA, 735–736.
- [6] 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 *16th Annual ACM Symposium on Parallel Algorithms and Architectures*. 133–144.
- [7] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In *PACT*. ACM, 72–81.
- [8] Christian Bienia and Kai Li. 2010. Characteristics of Workloads Using the Pipeline Programming Model. In *ISCA*. Springer-Verlag, 161–171.
- [9] Guy E. Blelloch, Phillip B. Gibbons, Yossi Matias, and Girija J. Narlikar. 1997. Space-Efficient Scheduling of Parallelism with Synchronization Variables. In *9th Annual ACM Symposium on Parallel Algorithms and Architectures*. 12–23.
- [10] Guy E. Blelloch and Margaret Reid-Miller. 1997. Pipelining with futures. In *SPAA*. ACM, 249–259.
- [11] Zoran Budimlić, Michael Burke, Vincent Cavé, Kathleen Knobe, Geoff Lowney, Ryan Newton, Jens Palsberg, David Peixotto, Vivek Sarkar, Frank Schlimbach, and Sağnak Taşlılar. 2010. Concurrent Collections. *Journal of Scientific Programming* 18, 3-4 (Aug. 2010), 203–217.
- [12] Vincent Cavé, Jisheng Zhao, Jun Shirako, and Vivek Sarkar. 2011. Habanero-Java: the new adventures of old X10. In *Proceedings of the 9th International Conference on Principles and Practice of Programming in Java (PPPJ '11)*. 51–61.
- [13] Rohit Chandra, Anoop Gupta, and John L. Hennessy. 1994. COOL: An Object-Based Language for Parallel Programming. *IEEE Computer* 27, 8 (Aug. 1994), 13–26.
- [14] 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 *20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*. 519–538.
- [15] Shuai Che, Michael Boyer, Jiayuan Meng, David Tarjan, Jeremy W. Sheaffer, Sang-Ha Lee, and Kevin Skadron. 2009. Rodinia: A benchmark suite for heterogeneous computing. In *2009 IEEE International Symposium on Workload Characterization (IISWC)*. 44–54.
- [16] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms* (third ed.). The MIT Press.
- [17] John S. Danaher, I-Ting Angelina Lee, and Charles E. Leiserson. 2008. Programming with exceptions in JCilk. *Science of Computer Programming* 63, 2 (Dec. 2008), 147–171.
- [18] Dimitar Dimitrov, Martin Vechev, and Vivek Sarkar. 2015. Race Detection in Two Dimensions. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '15)*. ACM, Portland, Oregon, USA, 101–110. <https://doi.org/10.1145/2755573.2755601>
- [19] Mingdong Feng and Charles E. Leiserson. 1997. Efficient Detection of Determinacy Races in Cilk Programs. In *Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*. 1–11.
- [20] Mingdong Feng and Charles E. Leiserson. 1999. Efficient Detection of Determinacy Races in Cilk Programs. *Theory of Computing Systems* 32, 3 (1999), 301–326.
- [21] Jeremy T. Fineman. 2005. *Provably Good Race Detection That Runs in Parallel*. Master's thesis. Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA.
- [22] Cormac Flanagan and Stephen N. Freund. 2009. FastTrack: efficient and precise dynamic race detection. *SIGPLAN Not.* 44, 6 (June 2009), 121–133.
- [23] Matthew Fluet, Mike Rainey, John Reppy, and Adam Shaw. 2010. Implicitly Threaded Parallelism in Manticore. *Journal of Functional Programming* 20, 5-6 (Nov. 2010), 537–576. <https://doi.org/10.1017/S0956796810000201>
- [24] D.P. Friedman and D.S. Wise. 1978. Aspects of Applicative Programming for Parallel Processing. *IEEE Trans. Comput.* C-27, 4 (1978), 289–296.
- [25] Matteo Frigo, Charles E. Leiserson, and Keith H. Randall. 1998. The Implementation of the Cilk-5 Multithreaded Language. In *PLDI*. ACM, 212–223.
- [26] Robert H. Halstead, Jr. 1985. Multilisp: A Language for Concurrent Symbolic Computation. *ACM TOPLAS* 7, 4 (Oct. 1985), 501–538.
- [27] Maurice Herlihy and Zhiyu Liu. 2014. Well-structured Futures and Cache Locality. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '14)*. ACM, Orlando, Florida, USA, 155–166. <https://doi.org/10.1145/2555243.2555257>
- [28] Intel 2013. Intel® Cilk™ Plus. <https://www.cilkplus.org/>. (2013).
- [29] Intel Corporation 2012. *Intel(R) Threading Building Blocks*. Intel Corporation. Available from http://software.intel.com/sites/products/documentation/doclib/tbb_sa/help/index.htm.
- [30] Intel Corporation 2013. *Intel® Cilk™ Plus Language Extension Specification, Version 1.1*. Intel Corporation. Document 324396-002US. Available from http://cilkplus.org/sites/default/files/open_specifications/Intel_Cilk_plus_lang_spec_2.htm.
- [31] Dileep Kini, Umang Mathur, and Mahesh Viswanathan. 2017. Dynamic Race Prediction in Linear Time. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2017)*. ACM, New York, NY, USA, 157–170. <https://doi.org/10.1145/3062341.3062374>
- [32] Alex Kogan and Maurice Herlihy. 2014. The Future(s) of Shared Data Structures. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing (PODC '14)*. ACM, Paris, France, 30–39. <http://doi.acm.org/10.1145/2611462.2611496>
- [33] David A. Kranz, Robert H. Halstead, Jr., and Eric Mohr. 1989. Mul-T: A High-Performance Parallel Lisp. In *PLDI*. ACM, 81–90.
- [34] I-Ting Angelina Lee, Silas Boyd-Wickizer, Zhiyi Huang, and Charles E. Leiserson. 2010. Using Memory Mapping to Support Cactus Stacks in Work-Stealing Runtime Systems. In *PACT*. ACM, 411–420.
- [35] I-Ting Angelina Lee and Tao B. Schardl. 2015. Efficiently Detecting Races in Cilk Programs That Use Reducer Hyperobjects. In *SPAA '15: Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '15)*. ACM, Portland, Oregon, USA, 111–122. <http://doi.acm.org/10.1145/2755573.2755599>
- [36] Charles E. Leiserson. 2010. The Cilk++ Concurrency Platform. *J. Supercomputing* 51, 3 (2010), 244–257.
- [37] Peng Liu, Omer Tripp, and Xiangyu Zhang. 2016. IPA: Improving Predictive Analysis with Pointer Analysis. In *Proceedings of the 25th*

- International Symposium on Software Testing and Analysis (ISSTA 2016)*. ACM, New York, NY, USA, 59–69. <https://doi.org/10.1145/2931037.2931046>
- [38] Li Lu, Weixing Ji, and Michael L. Scott. 2014. Dynamic Enforcement of Determinism in a Parallel Scripting Language. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '14)*. ACM, Edinburgh, United Kingdom, 519–529. <https://doi.org/10.1145/2594291.2594300>
- [39] John Mellor-Crummey. 1991. On-the-fly Detection of Data Races for Programs with Nested Fork-Join Parallelism. In *Proceedings of Supercomputing '91*. 24–33.
- [40] Robert H. B. Netzer and Barton P. Miller. 1992. What are Race Conditions? *ACM Letters on Programming Languages and Systems* 1, 1 (March 1992), 74–88.
- [41] Robert O'Callahan and Jong-Deok Choi. 2003. Hybrid Dynamic Data Race Detection. In *Proceedings of the Ninth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '03)*. ACM, New York, NY, USA, 167–178.
- [42] OpenMP 4.0 2013. *OpenMP Application Program Interface, Version 4.0*.
- [43] Eli Pozniansky and Assaf Schuster. 2003. Efficient On-the-fly Data Race Detection in Multithreaded C++ Programs. (2003), 179–190.
- [44] Raghavan Raman, Jisheng Zhao, Vivek Sarkar, Martin Vechev, and Eran Yahav. 2010. Efficient Data Race Detection for Async-Finish Parallelism. In *Runtime Verification*, Howard Barringer, Ylies Falcone, Bernd Finkbeiner, Klaus Havelund, Insup Lee, Gordon Pace, Grigore Rosu, Oleg Sokolsky, and Nikolai Tillmann (Eds.). Lecture Notes in Computer Science, Vol. 6418. Springer Berlin / Heidelberg, 368–383.
- [45] Raghavan Raman, Jisheng Zhao, Vivek Sarkar, Martin Vechev, and Eran Yahav. 2012. Scalable and Precise Dynamic Datarace Detection for Structured Parallelism. In *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '12)*. 531–542.
- [46] James Reinders. 2007. *Intel Threading Building Blocks: Outfitting C++ for Multi-core Processor Parallelism*. O'Reilly Media, Inc.
- [47] Mahmoud Said, Chao Wang, Zijiang Yang, and Karem Sakallah. 2011. Generating Data Race Witnesses by an SMT-based Analysis. In *Proceedings of the Third International Conference on NASA Formal Methods (NFM'11)*. Springer-Verlag, Berlin, Heidelberg, 313–327. <http://dl.acm.org/citation.cfm?id=1986308.1986334>
- [48] Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro, and Thomas Anderson. 1997. Eraser: A Dynamic Race Detector for Multi-Threaded Programs. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles (SOSP)*.
- [49] Konstantin Serebryany and Timur Iskhodzhanov. 2009. ThreadSanitizer: Data Race Detection in Practice. In *Proceedings of the Workshop on Binary Instrumentation and Applications (WBLA '09)*. ACM, New York, New York, 62–71.
- [50] Yannis Smaragdakis, Jacob Evans, Caitlin Sadowski, Jaeheon Yi, and Cormac Flanagan. 2012. Sound Predictive Race Detection in Polynomial Time. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '12)*. ACM, New York, NY, USA, 387–400. <https://doi.org/10.1145/2103656.2103702>
- [51] Daniel Spoonhower, Guy E. Blelloch, Phillip B. Gibbons, and Robert Harper. 2009. Beyond Nested Parallelism: Tight Bounds on Work-stealing Overheads for Parallel Futures. In *Proceedings of the Twenty-first Annual Symposium on Parallelism in Algorithms and Architectures (SPAA '09)*. ACM, Calgary, AB, Canada, 91–100. <https://doi.org/10.1145/1583991.1584019>
- [52] Rishi Surendran and Vivek Sarkar. 2016. Automatic Parallelization of Pure Method Calls via Conditional Future Synthesis. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 2016)*. ACM, New York, NY, USA, 20–38. <https://doi.org/10.1145/2983990.2984035>
- [53] Rishi Surendran and Vivek Sarkar. 2016. Brief Announcement: Dynamic Determinacy Race Detection for Task Parallelism with Futures. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '16)*. ACM, Asilomar State Beach, CA, USA, 95–97.
- [54] Rishi Surendran and Vivek Sarkar. 2016. *Dynamic Determinacy Race Detection for Task Parallelism with Futures*. Springer International Publishing, Cham, 368–385. https://doi.org/10.1007/978-3-319-46982-9_23
- [55] Robert Endre Tarjan. 1975. Efficiency of a Good But Not Linear Set Union Algorithm. *J. ACM* 22, 2 (April 1975), 215–225.
- [56] Sağnak Taşlılar and Vivek Sarkar. 2011. Data-Driven Tasks and Their Implementation. In *Proceedings of the 2011 International Conference on Parallel Processing (ICPP '11)*. IEEE Computer Society, Taipei City, Taiwan, 652–661.
- [57] Robert Utterback, Kunal Agrawal, Jeremy Fineman, and I-Ting Angelina Lee. 2016. Provably Good and Practically Efficient Parallel Race Detection for Fork-Join Programs. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '16)*. ACM, Asilomar State Beach, CA, USA, 83–94.
- [58] Jacobo Valdes. 1978. *Parsing Flowcharts and Series-Parallel Graphs*. Ph.D. Dissertation. Stanford University. STAN-CS-78-682.
- [59] Christoph von Praun and Thomas R. Gross. 2001. Object Race Detection. In *Proceedings of the 16th ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications (OOPSLA '01)*. ACM, Tampa Bay, FL, USA, 70–82.
- [60] Yifan Xu, I-Ting Angelina Lee, and Kunal Agrawal. 2018. Efficient Parallel Determinacy Race Detection for Two-dimensional Dags. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '18)*. ACM, Vienna, Austria, 368–380. <http://doi.acm.org/10.1145/3178487.3178515>
- [61] Yuan Yu, Tom Rodeheffer, and Wei Chen. 2005. RaceTrack: Efficient Detection of Data Race Conditions via Adaptive Tracking. In *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles (SOSP '05)*. ACM, New York, NY, USA, 221–234.