

An Ownership Policy and Deadlock Detector for Promises

Caleb Voss
cvoss@gatech.edu
Georgia Institute of Technology

Vivek Sarkar
vsarkar@gatech.edu
Georgia Institute of Technology

Abstract

Task-parallel programs often enjoy deadlock freedom under certain restrictions, such as the use of structured join operations, as in Cilk and X10, or the use of asynchronous task futures together with deadlock-avoiding policies such as Known Joins or Transitive Joins. However, the *promise*, a popular synchronization primitive for parallel tasks, does not enjoy deadlock-freedom guarantees. Promises can exhibit deadlock-like bugs; however, the concept of a deadlock is not currently well-defined for promises.

To address these challenges, we propose an ownership semantics in which each promise is associated to the task which currently intends to fulfill it. Ownership immediately enables the identification of bugs in which a task fails to fulfill a promise for which it is responsible. Ownership further enables the discussion of deadlock cycles among tasks and promises and allows us to introduce a robust definition of deadlock-like bugs for promises.

Cycle detection in this context is non-trivial because it is concurrent with changes in promise ownership. We provide a lock-free algorithm for precise runtime deadlock detection. We show how to obtain the memory consistency criteria required for the correctness of our algorithm under TSO and the Java and C++ memory models. An evaluation compares the execution time and memory usage overheads of our detection algorithm on benchmark programs relative to an unverified baseline. Our detector exhibits a 12% (1.12×) geometric mean time overhead and a 6% (1.06×) geometric mean memory overhead, which are smaller overheads than in past approaches to deadlock cycle detection.

CCS Concepts: • Computing methodologies → Parallel programming languages; Concurrent algorithms.

Keywords: promises, deadlocks, parallel programming

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 '21, February 27–March 3, 2021, Virtual Event, Republic of Korea

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8294-6/21/02...\$15.00

<https://doi.org/10.1145/3437801.3441616>

Listing 1. A deadlock?

```
1 Promise p,q;
2 t1 = async { ... };
3 t2 = async {
4   p.get(); // stuck
5   q.set();
6 };
7 q.get(); // stuck
8 p.set();
```

1 Introduction

The task-parallel programming model is based on the principle that structured parallelism (using high-level abstractions such as spawn-sync [20, 47], async-finish [12, 26, 33], futures [24, 36], barriers [47], and phasers [10, 50]) is a superior style to unstructured parallelism (using explicit low-level constructs like threads and locks). Structured programming communicates programmer intent in an upfront and visible way, providing an accessible framework for reasoning about complex code by isolating and modularizing concerns. However, the *promise* construct, found in mainstream languages including C++ and Java, introduces an undesirable lack of structure into task-parallel programming. A promise generalizes a future in that it need not be bound to the return value of a specific task. Instead, any task may elect to supply the value, and the code may not clearly communicate which task is intended to do so.

Promises provide point-to-point synchronization wherein one or more tasks can await the arrival of a payload, to be produced by another task. Although the promise provides a safe abstraction for sharing data across tasks, there is no safety in the kinds of inter-task blocking dependencies that can be created using promises. The inherent lack of structure in promises not only leads to deadlock-like bugs in which tasks block indefinitely due to a cyclic dependence, but such bugs are not well-defined and are undetectable in the general case due to the lack of information about which task is supposed to fulfill which promise.

A deadlock-like cycle may only be detected once all tasks have terminated or blocked. For example, the Go language runtime reports a deadlock if no task is eligible to run [25]. However, if even one task remains active, this technique cannot raise an alarm. An example of such a program is in Listing 1; the root task and t_2 are in a deadlock that may be hidden if t_1 is a long-running task, such as a web server. An alternative detection approach is to impose timeouts

on waits, which is only a heuristic solution that may raise an alarm when there is no cycle. In both of these existing approaches, the detection mechanism may find the deadlock some time *after* the cycle has been created. It is instead more desirable to detect a cycle immediately when it forms.

1.1 Promise Terminology

There is inconsistency across programming languages about what to call a promise and sometimes about what functionality “promise” refers to. The synchronization primitive we intend to discuss is called by many names, including promise [36], handled future [46], completable future [48], and one-shot channel [16]. For us, a promise is a wrapper for a data payload that is initially absent; each *get* of the payload blocks until the first and only *set* of the payload is performed. Setting the payload may also be referred to as completing, fulfilling, or resolving the promise.

Some languages, such as C++, divide the promise construct into a pair of objects; in this case, “promise” refers only to the half with a setter method, while “future” refers to the half with a getter method. In Java, the `CompletableFuture` class is a promise, as it implements the `Future` interface and additionally provides a setter method.

Habanero-Java introduced the data-driven future [51], which is a promise with limitations on when gets may occur. When a new task is spawned, the task must declare up front which promises it intends to consume. The task does not become eligible to run until all such promises are fulfilled.

In JavaScript, the code responsible for resolving a promise must be specified during construction of the promise [44]. This is a limitation that makes deadlock cycles impossible, although the responsible code may omit to resolve the promise altogether, leading to unexecuted callbacks.

Promises may provide a synchronous or an asynchronous API. The Java concurrency library provides both, for example [48]. The synchronous API consists of the *get* and *set* methods. The asynchronous API associates each of the synchronous operations to a new task. A call to *supplyAsync* binds the eventual return value of a new task to the promise. The *then* operation schedules a new task to operate on the promise’s value once it becomes available. The asynchronous API can be implemented using the synchronous API. Conversely, the synchronous API can be implemented using continuations and an asynchronous event-driven scheduler [32]. We focus on the synchronous API in this work.

1.2 Two Bug Classes

We identify two kinds of synchronization bug in which the improper use of promises causes one or more tasks to block indefinitely:

1. the *deadlock cycle*, in which tasks are mutually blocked on promises that would be set only after these tasks unblock, and

Listing 2. An omitted set?

```

1 Promise r, s;
2 t3 = async { // should set r, s
3   t4 = async { // should set s
4     // (forgot to set s)
5   }
6   r.set();
7 };
8 r.get();
9 s.get(); // stuck

```

2. the *omitted set*, in which a task is blocked on a promise that no task intends to set.

However, neither of these bugs manifests in an automatically recognizable way at runtime unless every task in the program is blocked. In fact, the definitions of these bugs describe conditions which cannot generally be detected. What does it mean for no task to *intend* to set a promise? What does it mean that a task *would* set a promise once the task unblocks? In a traditional deadlock, say one involving actual locks, the cycle is explicit: Task 1 holds lock *A* and blocks while acquiring lock *B*, because task 2 is holding lock *B* and concurrently blocked during its acquisition of lock *A*. Intention to release a lock (thereby unblocking any waiters) is detectable by the fact that a task holds the lock. But we currently have no concept of a task “holding” a promise and no way to tell that a task intends to set it.

1.3 Need for Ownership Semantics

Consider the small deadlock in Listing 1. Two promises, *p*, *q*, are created. Task *t*₂ waits for *p* prior to setting *q*, whereas the root task waits for *q* prior to setting *p*. Clearly a deadlock cycle arises? Not so fast. To accurately call this pattern a deadlock cycle requires knowing that task *t*₁ will not ever set *p* or *q*. Such a fact about what *will* not happen is generally not determinable from the present state without an offline program analysis. For this reason, a deadlock cycle among promises evades runtime detection unless the cycle involves every currently executing task.

Now consider the bug in Listing 2. Two promises, *r*, *s*, are created. According to the comments, task *t*₃ is responsible for setting both, and it subsequently delegates the responsibility for *s* to *t*₄. However, *t*₄ fails to perform its intended behavior, terminating without setting *s*. The root task then blocks on *s* forever. If a bug has occurred, we would like to raise an alarm at runtime when and where it occurs. Where is this bug? Should the root task not have blocked on *s*? Should *t*₄ have set *s*? Should *t*₃ have set *s*? The blame cannot be attributed, and the bug may, in fact, be in any one of the tasks involved. Furthermore, *when* does this bug occur? The *symptom* of the bug manifests in the indefinite blocking of the root task, potentially *after* *t*₄ terminates successfully. If some other task may yet set *s*, then this bug is not yet confirmed to have

Listing 3. An omitted set in Amazon AWS SDK (v2) [31]. Code abbreviated and inlined for clarity.

```

1 private CompletableFuture<Void> cf;
2
3 public void onComplete () {
4     ...
5     if (streamChecksumInt != computedChecksumInt)
6     {
7         onError (...); // Assumed to fulfill promise
8         return; // Don't fulfill promise again
9     }
10    ...
11    cf.complete(null); // Fulfills promise
12 }
13
14 public void onError (Throwable t) {
15     // Originally a no-op. Fixed to:
16     cf.completeExceptionally (t);
17 }

```

occurred. Omitted sets evade runtime detection and, even once discovered, evade proper blame assignment.

We propose to augment the task creation syntax (**async** in our examples) to carry information about promise ownership and responsibility within the code itself, not in the comments. In doing so, omitted sets become detectable at runtime with blame appropriately assigned. Moreover, programmer intent is necessarily communicated in the code. Finally, in knowing which task is expected to set each promise, it becomes possible to properly discuss deadlock cycles among promises.

1.4 Omitted Set in the Wild

An example of an omitted set bug was exhibited by the Amazon Web Services SDK for Java (v2) when a certain checksum validation failed [31]. An abbreviated version of the code is given in Listing 3; line 16 was absent prior to the bug fix. The control flow ensures that either exception handling code or non-exceptional code was executed, not both (line 8) [43]. However, only the non-exceptional code would set the value of a `CompletableFuture` (Java’s promise) to indicate the work was completed (line 11), whereas the `onError` method would take no action. If checksum validation failed after a file download, any consumer tasks waiting for the download to complete would block indefinitely. A month later, the omitted set bug was identified and corrected by adding line 16 [2].

When this bug arises at runtime, the symptom (the blocked consumer) is far from the cause (the omitted set), and the bug is not readily diagnosable. If the runtime could track which tasks are responsible for which promises, then this bug could be detected and reported as an exception as soon as the responsible task terminates. Using our approach, the bug would be detected when the task running the `onComplete` callback finishes, and the alarm would name the offending task and the unfulfilled promise.

1.5 Contributions

In this work, we propose the addition of *ownership semantics* for promises which enables a task’s intention to set a promise to be reflected in the runtime state. In so doing,

1. we enable a precise definition of a deadlocked cycle of promises in terms of runtime state;
2. we define a second kind of blocking bug, the *omitted set*, which does not involve a cycle;
3. we require important programmer intent to be encoded explicitly and to respect a runtime-verifiable policy, thereby enabling structured programming for promises.

In addition to these theoretical contributions,

1. we introduce a new lock-free algorithm for detecting our now-identifiable deadlock-cycle and omitted-set bugs *when they occur*;
2. we identify properties critical for establishing the correctness of the algorithm under weak memory consistency and show how to ensure these properties hold under the TSO, Java, and C++ memory models;
3. we prove that our algorithm precisely detects every deadlock without false alarms;
4. we experimentally show that a Java implementation has low execution time and memory usage overheads on nine benchmarks relative to the original, unverified baseline (geometric mean overheads of 1.12× and 1.06×, respectively).

2 Ownership Policy

In promise-based synchronization, a task does not directly await another task; it awaits a promise, thereby *indirectly* waiting on whichever task fulfills that promise. It is a runtime error to fulfill a promise twice, so there ought to be one and only one fulfilling task. However, the relationship between a promise and the task which *will* fulfill it is not explicit and inhibits the identification of deadlocks. To make this relationship explicit and meaningful, we say that each promise is *owned* by exactly one task at any given time. The owner is responsible for fulfilling the promise eventually, or else handing ownership off to another task. Ownership hand-offs may only occur at the time of spawning a new task. We augment the **async** keyword, used to spawn tasks, with a list of promises currently owned by the parent task that should be transferred to the new child.

2.1 Language Extension

We define an abstract language, showing only its synchronization instructions and leaving its sequential control flow and other instructions unspecified. For simplicity, we have abstracted away the payload values of promises and refer to individual promises by globally unique identifiers.

Definition 2.1. The \mathcal{L}_p language consists of task-parallel programs, P , whose synchronization instructions have the syntax

new p | **set** p | **get** p | **async** (p_1, \dots, p_n) $\{P\}$

where n may be 0.

The instruction **new** p represents the point of allocation for the promise p , and we assume well-formed programs do not allocate a given p twice or operate on p prior to its allocation. Each invocation of **get** p blocks the current task until after **set** p has been invoked for the first (and only) time.

The **async** block creates a new task to execute a sub-program P ; the block is annotated with a list of promises, which should be moved from the parent task to the new task. In many task-parallel languages, **async** automatically creates a future which can be used to retrieve the new task's return value. We can readily reproduce this behavior using promises in the pattern **new** p ; **async** (p, \dots) $\{\dots$; **set** $p\}$.

Definition 2.2. The *ownership policy*, \mathcal{P}_o , maintains state during the execution of an \mathcal{L}_p program in the form of a map $\text{owner} : \text{Promise} \rightarrow \text{Task} \cup \{\text{null}\}$ according to these rules:

1. When task t executes **new** p , set $\text{owner}(p) := t$.
2. When task t spawns task t' as **async** (p_1, \dots, p_n) $\{P\}$, prior to t' becoming eligible to run, ensure $\text{owner}(p_i) = t$ and update $\text{owner}(p_i) := t'$ for each p_i .
3. When task t terminates, ensure the set of promises $\text{owner}^{-1}(t)$ is empty.
4. When task t executes **set** p , ensure that $\text{owner}(p) = t$ and set $\text{owner}(p) := \text{null}$.

These four rules together ensure that there is at least one **set** for each promise, with omitted sets being detected by rule 3. Rule 4 guarantees there is at most one **set**.

Our proposed modification to the program given in Listing 1 is to annotate the **async** in line 3 as **async** (q), indicating that t_2 takes on the responsibility to set q . It is now possible to trace the cycle when it occurs: the root task awaits q , owned by t_2 , awaiting p , owned by the root task. It is clear that t_1 , whose **async** is not given any parameters, is not involved as it can set neither p nor q (rule 4).

The proposed modification to the program given in Listing 2 is to write **async** (r, s) in line 2 and **async** (s) in line 3. That is, the information already present in the comments is incorporated into the code itself. The moment t_4 terminates, the runtime can observe that t_4 still holds an outstanding obligation to set s . We treat this as an error immediately (rule 3), irrespective of whether any task is awaiting s .

2.2 Algorithm for Ownership Tracking

Algorithm 1 implements the \mathcal{P}_o policy by providing code to be run during **new**, **async**, and **set** operations. Each promise has an owner field to store the task that is currently its owner,

Algorithm 1 Promise Ownership Management

```

1: procedure NEW()
2:    $t \leftarrow \text{currentTask}()$ 
3:    $p \leftarrow \{\text{owner} : t\}$             $\triangleright$  C: atomic, Java: volatile
4:   append  $p$  to  $t.\text{owned}$ 
5:   return  $p$ 

6: procedure ASYNC( $P, f$ )
7:    $t \leftarrow \text{currentTask}()$ 
8:   assert  $p.\text{owner} = t$  forall  $p \in P$ 
9:    $t' \leftarrow \{\text{owned} : P,$ 
10:    waitingOn :  $\text{null}\}$             $\triangleright$  C: atomic, Java: volatile
11:   remove all of  $P$  from  $t.\text{owned}$ 
12:    $p.\text{owner} \leftarrow t'$  forall  $p \in P$ 
13:   do asynchronously
14:      $\text{setCurrentTask}(t')$ 
15:      $f()$ 
16:     assert  $t'.\text{owned}$  is empty
17:   return  $t'$ 

18: procedure INIT( $\text{main}$ )
19:    $\text{setCurrentTask}(\text{null})$ 
20:   ASYNC( $[], \text{main}$ )

21: procedure SET( $p, v$ )
22:    $t \leftarrow \text{currentTask}()$ 
23:   assert  $p.\text{owner} = t$ 
24:    $p.\text{owner} \leftarrow \text{null}$ 
25:   remove  $p$  from  $t.\text{owned}$ 
26:    $\text{set\_impl}(p, v)$ 

```

and each task has an associated owned list that maintains the inverse map, owner^{-1} . The functions *currentTask* and *getCurrentTask* interact with thread-local storage.

In compliance with \mathcal{P}_o rule 1, the New procedure creates a promise owned by the currently running task (line 3) and adds this promise to that task's owned list (line 4).

ASYNC(P, f) schedules f to be called asynchronously as a new task and moves the promises listed in P into this task. These promises are first confirmed to belong to the parent task (line 9), then moved into the child task (lines 9–12), in accordance with rule 2. (Line 10 is in preparation for Algorithm 2, presented in section 3.) Once the child task terminates, rule 3 requires that the task not own any remaining promises (line 16). The INIT procedure shows how to set up a root task to execute the main function.

Finally, SET(p, v) achieves rule 4, checking that the current task owns p and marking p as fulfilled by assigning it no owner (lines 23–25). The procedure then invokes the underlying mechanism for actually setting the promise value to v (line 26).

As an example of how Algorithm 1 enforces compliance with \mathcal{P}_o , refer again to Listing 2. When promise s is first

created, it belongs to the root task (Algorithm 1 line 4). If the **async** that creates t_4 is annotated with s , then Algorithm 1 line 12 changes the owner of s to t_4 . Since t_4 does not set s , upon termination of t_4 , an assertion fails in Algorithm 1 line 16. The offending task, t_4 , and the outstanding promise, s , are directly identifiable and can be reported in the alarm.

3 Deadlock Detection Algorithm

Now that we have established the relationship between promises and tasks, it is possible to describe what a deadlock is. A deadlock is a cycle of n tasks, t_i , and n promises, p_i , such that t_i awaits p_i while p_i is owned by $t_{i+1} \pmod n$. The information required to identify such a deadlock is, for the first time, made available explicitly at runtime through the use of the \mathcal{P}_o policy. We can now develop a runtime detection mechanism to identify deadlocks based on this information and raise an alarm as soon as one is created.

3.1 Approach

Even assuming sequential consistency, the algorithm for finding such a cycle is non-trivial. Conceptually, whenever a **get** p is executed by t , t must alternately traverse owned-by and waits-for edges to see if the path of dependences returns to t . If another task, t' , is encountered which is not currently awaiting a promise, this proves that progress is still being made and there is no deadlock (yet). In this case, t passes verification and commits to blocking on p . Should this path of dependences grow due to a subsequent **get** p' by t' , then the same algorithm runs again in task t' to verify that the new waits-for edge does not create a deadlock.

Crucially, during verification t must establish a waits-for edge to mark that it is awaiting p prior to traversing the dependence path. That is, a waits-for edge is created before it is determined that t will be allowed to await p . A two-task cycle shows what would go wrong if this procedure is not followed. If t begins to verify its wait of p (say, owned by t') without marking that t is awaiting p , and concurrently t' begins to verify its wait of p' (owned by t) without marking that t' is awaiting p' , then each task may find that the other is apparently not awaiting any promises at this time, and both commit to blocking, creating an undetected deadlock. However, by ensuring that each task marks itself as awaiting a promise prior to verifying whether that wait is safe, we guarantee the last task to arrive in the formation of a deadlock cycle will be able to detect this cycle.

A second consideration is how this approach handles concurrent transfer of promise ownership or concurrent fulfillment of promises. Suppose that while the cycle detection algorithm is traversing a dependence path, an earlier promise in the path is transferred to a new owner or is fulfilled, thereby invalidating the remainder of the traversed path. Failure to handle this correctly could result in an alarm when there is no deadlock. The first observation we make

Algorithm 2 Deadlock Cycle Detection

```

1: procedure GET( $p_0$ )
2:    $t_0 \leftarrow \text{currentTask}()$ 
3:    $t_0.\text{waitingOn} \leftarrow p_0$  ▷ C: seq_cst
4:   ▷ TSO: memory fence
5:    $i \leftarrow 0$ 
6:    $t_{i+1} \leftarrow p_i.\text{owner}$ 
7:   while  $t_{i+1} \neq t_0$  do
8:     if  $t_{i+1} = \text{null}$  then break
9:      $p_{i+1} \leftarrow t_{i+1}.\text{waitingOn}$  ▷ C: acquire
10:    if  $p_{i+1} = \text{null}$  then break
11:    if  $t_{i+1} \neq p_i.\text{owner}$  then break
12:     $i \leftarrow i + 1$ 
13:     $t_{i+1} \leftarrow p_i.\text{owner}$ 
14:  try
15:    assert  $t_{i+1} \neq t_0$ 
16:    return  $\text{get\_impl}(p_0)$ 
17:  finally
18:     $t_0.\text{waitingOn} \leftarrow \text{null}$  ▷ C: release

```

is that this scenario cannot arise for any but the most recent promise encountered on the path. If p_0 is owned by t_1 , awaiting p_1 , owned by t_2 , then it is impossible for p_0 to move into a new task or to become fulfilled, since its current owner, t_1 , is blocked (or about to block, pending successful verification). The concern is only that t_2 has not yet blocked and may transfer or fulfill p_1 . The natural solution is that when traversing the dependence path, upon reaching each promise in the path we must go back and double-check that the *preceding* promise still belongs to the task it belonged to in the previous iteration and is still unfulfilled. If this check fails, then the present verification passes because progress is still being made.

3.2 Detection Algorithm

The deadlock detector occupies the implementation of the **get** instruction, given in Algorithm 2. This detector can thereby raise an alarm in a task as soon as the task attempts a deadlock-forming await of a promise. At the time of raising an alarm, the available diagnostic information that can be reported includes the task, the awaited promise, as well as every other task and promise in the cycle, if desired.

For a preliminary understanding of the procedure's logic, we assume sequential consistency in this section. Upon entering GET, the currently executing task records the promise that it will be waiting on (line 3). This waitingOn field was initialized to *null* in Algorithm 1 line 10, and is always reset to *null* upon exiting GET (Algorithm 2 line 18), either normally (line 16) or abnormally (line 15). Doing so makes the algorithm robust to programs with more than one deadlock.

The loop in the detection algorithm traverses the chain of alternating owner and waitingOn fields. If task t is waiting on promise p , which is owned by a task t' , then t is effectively

waiting on whatever t' awaits. In traversing this chain, if t finds that it is transitively waiting on itself, then we have identified a deadlock (lines 7, 15). If the algorithm reaches the end of this chain without finding t again, as indicated by finding a *null* value in line 8 (p_i is already fulfilled) or in line 10 (t_{i+1} is not awaiting a promise), then it is safe to commit to a blocking wait on the desired promise (line 16). Recall that $p_i.owner$ is *null* after p_i has been fulfilled, and $t_{i+1}.waitingOn$ is *null* when t_{i+1} is not currently executing GET.

In order to guarantee that an apparent cycle always corresponds to a real deadlock, even under concurrent updates to promises, we rely on line 11 to establish that task t_{i+1} was waiting on promise p_{i+1} while t_{i+1} was still the owner of promise p_i . This is achieved by reading the owner field both before (line 6, 13) and after (line 11) reading the `waitingOn` field (line 9). If the task observes the owner of p_i to have changed, it turns out that it is safe to abandon the deadlock check and commit to the blocking wait.

In sections 4–5, we will move to a weaker memory model. There are two crucial points to remember. We must preserve the ability to reason temporally over the edges in the dependence path, and we must guarantee that at least one task entering a deadlock can observe the existence of the whole deadlock cycle.

4 Weakly Consistent Definition of Deadlock

With a few tweaks, we can obtain a correctness guarantee for our deadlock detector under a weak memory model, which implies the same guarantee under any stronger model, including sequential consistency. First, we must define this weak memory model and give a definition of deadlock that is compatible with it.

In practice, we do not want to assume that maps such as the owner field have a single, globally consistent state that is observed by all tasks. Machines and languages often have weaker consistency guarantees, and there are performance costs for requesting stronger consistency due to the synchronization required. Instead, we will assume a weak memory model and use unsynchronized accesses whenever possible.

We now define this weak memory model, which we will use to establish the correctness of our deadlock detection algorithm under models at least as strong as this one.

Definition 4.1. The *happens-before* (h.b.) order is a partial order over the instructions in a program execution that subsumes the intra-task program order and, upon spawning each new task, the ordering of Algorithm 1 line 14 (the start of the new task) after Algorithm 1 line 12 (the last action of the parent task before spawning). The reverse of happens-before is *happens-after*.

Definition 4.2. With respect to a given memory location, a read may only *observe* a (not necessarily unique) last write

which happens-before it or any write with which the read is not h.b. ordered. Two writes or a write and read of the same location which are not h.b. ordered are *racing*.

A typical language has a more refined happens-before ordering and definition of observable writes, especially relating to reads-from edges on promises; however, we will not need to appeal to such edges in our formalism.

Definition 4.3. A program in \mathcal{L}_p is *well-formed* if, in every execution, for each promise, p , there is at most one **new** p instruction, and each **set**, **get**, or **async** instruction referring to p happens-after such a **new** p .

We note that although the owners of different promises may be updated concurrently, it is not possible in Algorithm 1 for a write-write race to occur on the same owner field.

Lemma 4.4. Consider an execution of a well-formed program. If w_1, w_2 are two writes to $p.owner$ in Algorithm 1, then w_1 and w_2 are not racing. Further, if r is a read of $p.owner$ by task t , and r observes the value to be t , then r does not race with the write it observes.

Proof. The two claims can be shown together. Line 3 represents the initialization of the owner field and so happens-before every other write to it. The writes in lines 12 and 24 each happen-after a read of the same field observes the value to be the currently executing task (lines 8, 23). Take this together with the fact that there are only two ways to set $p.owner$ to t : line 3, executed by t itself, or line 12, executed by the parent of t prior to spawning t . In either case, writing t to $p.owner$ happens-before any read of $p.owner$ by t itself. \square

Since we do not assume a globally consistent state, we have to be careful in the definition of deadlock cycle. Two tasks need not agree on the value of $owner(p)$ for a given promise, p . Instead of freely referring to $owner$ as a map $Promise \rightarrow Task \cup \{null\}$, we must additionally state which task's perspective is being used to observe the owner map.

Definition 4.5. A non-empty set of tasks, T , is in a *deadlock cycle* if for every task $t \in T$,

1. t is executing **get** p_t for some promise, p_t ,
2. there exists a task, o_{p_t} , also in T which observes that $owner(p_t) = o_{p_t}$,

and T is minimal with respect to these constraints. The set of promises associated to the deadlock is $\{p_t \mid t \in T\}$.

The subtle point in this definition is that task o_{p_t} necessarily has the most up-to-date information about the owner of p_t , since o_{p_t} is itself the owner. Per lemma 4.4, we know that all the writes to $p_t.owner$ are ordered and that o_{p_t} is observing the last such write, since only o_{p_t} is capable of performing the next write to follow the observed one.

5 Correctness under Weak Consistency

Algorithm 2 correctly and precisely detects all deadlocks under our weak memory consistency model with some additional specific consistency requirements on certain accesses. We define these requirements, show how to meet them in each of the TSO, Java, and C++ memory models, and then prove the algorithm raises an alarm exactly when there is a deadlock.

5.1 Requirements

In order to prove correctness, we require the following additional memory consistency.

1. There is a total order, $<$, over all instances of the write in Algorithm 2 line 3, across all memory locations. Let $w_1 < w_2$. Any write preceding and including w_1 in h.b. order is visible to any read following w_2 in h.b. order.
2. The consistency of any owner field is expected to follow from release-acquire semantics for any waitingOn field. Specifically, let w_1 be an Algorithm 1 line 3 or line 12 write to an owner field, let w_2 be an Algorithm 2 line 3 write to a waitingOn field, let r_2 be an Algorithm 2 line 9 read, and let r_1 be an Algorithm 2 line 11 read. Suppose w_1, r_1 refer to the same location, as do w_2, r_2 . If w_1 happens-before w_2 , if w_2 is visible to r_2 , and if r_2 happens-before r_1 , then w_1 is visible to r_1 .
3. The write in Algorithm 2 line 18 must not become visible until the fulfillment of p_0 is visible (Algorithm 1 line 24) or it is determined that an exception should be raised (Algorithm 2 line 15).

These three requirements are readily attained in TSO, Java, and C++ as follows.

- Under TSO, a memory fence is needed in Algorithm 2 line 4 to achieve requirement 1 by ordering line 9 after line 3 and sequentializing all instances of line 4 with each other. TSO naturally achieves requirement 2 by respecting the local store order, as well as requirement 3 by not allowing the line 18 write to become visible early. Note that the loop contains no fences.
- Under the Java memory model, it suffices to mark the two fields, owner and waitingOn, as volatile to satisfy all three requirements. This eliminates all write-read data races. Remember that there are no write-write races (see lemma 4.4). In the absence of any races on these two fields, the Java memory model guarantees sequential consistency with respect to these fields.
- In C++ both of the fields must be `std::atomic` to eliminate data races, but this alone is insufficient. Algorithm 2 line 3 must be tagged as a `std::memory_order_seq_cst` access to achieve requirement 1, establishing a total order over these writes and subsuming release

consistency. Line 9 must then be tagged `std::memory_order_acquire` to achieve requirement 2. And finally, line 18 must be `std::memory_order_release` to satisfy 3.

5.2 Correctness

Under the preceding consistency requirements, we can now prove important theoretical guarantees of correctness for our deadlock detector. Throughout, we consider an execution of a well-formed program (definition 4.3).

We first show that Algorithm 2 raises no false alarms.

Theorem 5.1. *If task t fails the assertion in line 15 during $GET(p)$, then a deadlock cycle exists, involving t and p .*

Proof. We have $t_0 = t$ and $p_0 = p$. If the execution had broken out of the while loop in line 8, 10, or 11, then the assertion would have succeeded. Therefore, it is the loop condition that fails. Upon reaching line 12 in each iteration, we have found $p_i.owner$ to be t_{i+1} both before and after we found $t_{i+1}.waitingOn$ to be p_{i+1} . Therefore, we know 1) that at one time t_{i+1} was the owner of p_i , and 2) that while t_{i+1} still observed itself to own p_i , t_{i+1} had invoked $GET(p_{i+1})$. This follows from memory consistency requirement 2. At this point in the reasoning, we do not yet know if t_{i+1} still the owner of p_i or if t_{i+1} is still awaiting p_{i+1} .

When the loop (lines 7–13) terminates with $t_{i+1} = t_0$, since t_0 is the current task, we deduce that the final t_{i+1} , set by line 6 or 13, is the current owner of p_i . For all k modulo $i + 1$, t_k at one time concurrently observed itself to be the owner of p_{k-1} and was in a call to $GET(p_k)$. This meets our definition of deadlock. \square

The following series of lemmas builds to the theorem that Algorithm 2 detects every deadlock.

Definition 5.2. In a deadlock cycle comprising tasks T , a t^* task is a task in T to which the line 3 write by every task in T is visible.

Lemma 5.3. *Every deadlock cycle has a t^* task.*

Proof. Corollary to memory consistency requirement 1. \square

A t^* task, which need not be unique, should be thought of as the (or a) last task to enter the deadlock.

Lemma 5.4. *If a program execution exhibits a deadlock cycle comprising tasks T and promises P , when a t^* task calls GET it constructs a sequence $\{t_i\}_i$ that is a subset of T and a sequence $\{p_i\}_i$ that is a subset of P .*

Proof. We have $t_0 = t^* \in T$ and, by definition, $p_0 \in P$. If the loop immediately terminates, then $t_1 = t_0 \in T$, and we are done. Otherwise, the values of t_{i+1} and p_{i+1} inductively depend on t_i and p_i . By definition of deadlock, one of the tasks in T , call it o_{p_i} , observes itself to be the owner of p_i . The most recent write to $p_i.owner$ (recall all the writes are ordered by lemma 4.4) occurred in program order before o_{p_i} 's

line 3 write. Therefore, memory consistency requirement 1 establishes that t^* must read $t_{i+1} = o_{p_i} \in T$ in line 11. By definition of t^* and by memory consistency requirement 3, we see that line 9 observes t_{i+1} 's line 3 write, not its line 18 write. Thus, $p_{i+1} \in P$ by definition of deadlock. \square

Lemma 5.5. *If a program execution exhibits a deadlock cycle comprising tasks T , no t^* task executes a diverging loop (lines 7–13) in its call to GET.*

Proof. Suppose, during the call to GET by t^* , the loop does not terminate. Thus $t_i \neq t_0$ for any $i > 0$. But by lemma 5.4, the infinite sequence $\{t_i\}_i$ is a subset of T . Therefore, T , in fact, exhibits a smaller cycle not involving t_0 , violating the minimality condition in the definition of deadlock cycle. \square

Theorem 5.6. *If a program execution exhibits a deadlock cycle comprising tasks T and promises P , at least one task in T fails the assertion in Algorithm 2 line 15.*

Proof. Suppose for the sake of contradiction that a deadlock cycle arises and yet no assertion fails. So every task $t \in T$ enters the GET procedure and either blocks at line 16 on a promise in P or diverges in an infinite loop.

No task exits the loop by failing the loop condition, $t_{i+1} \neq t_0$, since this would directly fail the assertion in line 15.

For each invocation of GET by a t^* task, the loop cannot break in line 8 or line 10 because lemma 5.4 implies no tasks or promises in the sequence are *null*. If the loop breaks in line 11, then t^* has observed the owner of p_i to change from one read to the next. This is impossible: both reads observe the current owner, o_{p_i} , by the same reasoning as in the proof of lemma 5.4. Finally, the loop cannot diverge for t^* , by lemma 5.5. Since there exists at least one t^* task, by lemma 5.3, we have a contradiction. \square

Corollary 5.7 (to theorems 5.1, 5.6). *Algorithm 2 is precise and correct, guaranteeing the existence of a deadlock when an alarm is raised and raising an alarm upon every deadlock.*

6 Implementation and Evaluation

We have implemented ownership semantics with omitted set and deadlock detection in Java. We give a brief discussion of some of the practical considerations in the design of this implementation. We then present the results of a performance evaluation on a set of benchmark programs.

6.1 Objected-Oriented Promise Movement

Introducing an explicit conception of ownership is minimally disruptive. It is already the case that every promise is fulfilled by at most one task, since two sets cause a runtime error. We only ask that the programmer identify this task by leveraging the existing structure of **async** directives. However, for large, complex synchronization patterns that rely on many promises, it can become tedious for a programmer to specify all the relevant promises, one by one.

Listing 4. Object-oriented approach to promise movement.

```

1 class Channel<T> implements PromiseCollection {
2     class Payload {
3         T value;
4         Promise<Payload> next;
5     }
6
7     Promise<Payload> producer = new Promise<>();
8     Promise<Payload> consumer = producer;
9
10    @Override // from PromiseCollection
11    Iterable<Promise<?>> getPromises () {
12        // Return the set of all promises that
13        // should be moved when this object moves
14        return Collections.singleton(producer);
15    }
16
17    void send (T value) {
18        // Fulfills one promise; allocates another
19        Promise<Payload> next = new Promise<>();
20        producer.set({value, next});
21        producer = next;
22    }
23
24    void stop () {
25        // Fulfills a promise
26        producer.set(null);
27    }
28
29    T recv () {
30        Payload p = consumer.get();
31        consumer = p.next;
32        return p.value;
33    }
34 }
35
36 void main () {
37     Channel<Integer> ch = new Channel<>();
38     ch.send(1);
39     async (ch) { // Move entire channel
40         ch.send(2);
41         ch.stop();
42         // No remaining promises
43     }
44     // No remaining promises
45     ch.recv(); // 1
46     ch.recv(); // 2
47 }

```

In our Java implementation, an object-oriented approach can reduce the burden of identifying which promises should be moved to new tasks. In our Java implementation of these language features, classes containing many promises may implement a PromiseCollection interface so that moving a composite object to a new task is equivalent to moving each of its constituent promises. A channel class is shown in Listing 4, illustrating that complex and versatile primitives can

be built on top of promises with the aid of PromiseCollection. This class behaves like a promise that can be used repeatedly, where the n th `recv` operation obtains the value from the n th `send` operation. This behavior depends on dynamically allocated promises, and the responsibility for the sending end of the channel is associated not to the ownership of a single promise, but to the ownership of different promises at different times. It is abstraction-breaking to ask the channel user to manually specify which promise to move to a new task in order to effectively move the sending end of the channel. Instead, we give the impression that the channel object itself is movable like a promise (line 39), since it is a PromiseCollection, and the implementation of `async` relies on the `getPromises` method (line 11) to determine which promises should be moved.

6.2 Exception Handling

In an implementation of Algorithm 1, some care must go into an exception handling mechanism. What code is capable of and responsible for recovering from the failed assertion in line 16? And what happens if a task terminates early, with unfulfilled promises, because of an exception?

Observe that line 16 occurs within an asynchronous task after the user-supplied code for that task has completed. One solution is to add a parameter to `ASYNC` so that the user can supply a post-termination exception handler, which accepts the list of unfulfilled promises, t' .owned, as input. Indeed, the fix for the AWS omitted set bug included such a mechanism (not shown in Listing 3) [2]. Alternatively, the runtime could automatically fulfill every unfulfilled promise upon an assertion failure in line 16. Some APIs, including in C++ and Java, provide an exceptional variant of the completion mechanism for promises [36, 48]. In our implementation, we use this mechanism to propagate an exception through the promises that were left unfulfilled.

Finally, observe that the correctness of Algorithm 1 only depends on knowing when a task's owned list is empty. Therefore, the owned list could be correctly replaced with a counter, which would at least reduce the memory footprint of ownership tracking, if not also the execution time of maintaining a list. However, doing so would mean that an assertion failure in line 16 could not indicate *which* promises went unfulfilled. Therefore, the implementation we evaluate uses an actual list.

6.3 Benchmarks

We evaluate the execution time and memory usage overheads introduced by our promise deadlock detector on nine task-parallel programs. The overheads are measured relative to the original, unverified baseline versions.

1. Conway [56] parallelizes a 2D cellular automaton by dividing the grid into chunks. We adapted the code from C to Java, using our Channel class (Listing 4)

in place of MPI primitives used by worker tasks to exchange chunk borders with their neighbors.

2. Heat [9] simulates diffusion on a one-dimensional surface, with 50 tasks operating on chunks of 40,000 cells for 5000 iterations. Neighboring tasks again use Channel in place of MPI primitives.
3. QSort sorts 1M integers using a parallelized divide-and-conquer recursion; the partition phase is not parallelized. This is a standard technique for parallelizing Quicksort [19] and has been previously implemented using the Habanero-Java Library [33]. We implemented the finish construct, which awaits task termination using promises.
4. Randomized distributes 5000 promises over 2535 tasks spawned in a tree with branching factor of 3. Each task awaits a random promise with probability 0.8 before performing some work, fulfilling its own promises, and awaiting all its child tasks. We chose a random seed that does not construct a deadlock.
5. Sieve counts the primes below 100,000 with a pipeline of tasks, each filtering out the multiples of an earlier prime. A similar program is found in prior work [45].
6. SmithWaterman (adapted from HClib [26]; also used in prior work [15, 55]) aligns DNA sequences having 18,000–20,000 bases. Each task operates on a 25×25 tile.
7. Strassen (such a program is found in the Cilk, BOTS, and KASTORS suites [17, 20, 53]) multiplies sparse 128×128 matrices containing around 8000 values. Divide-and-conquer recursion issues asynchronous addition and multiplication tasks, up to depth 5.
8. StreamCluster (from PARSEC [5]) computes a streaming k -means clustering of 102,400 points in 128 dimensions, using 8 worker tasks at a time. We replaced the OpenMP barriers with promises in an all-to-all dependence pattern.
9. StreamCluster2 reduces synchronization in StreamCluster by replacing some of the all-to-all patterns with all-to-one when it is correct to do so. We also correct a data race in the original implementation.

All benchmarks were run on a Linux machine with a 16-core AMD Opteron processor under the OpenJDK 11 VM with a 1 GB memory limit. A thread pool schedules asynchronous tasks by spawning a new thread for a new task when all existing threads are in use. This execution strategy is necessary in general for promises because there is no *a priori* bound on the number of tasks that can block simultaneously. We measured both execution time and, in a separate run, average memory usage by sampling every 10 ms. Each measurement is averaged over thirty runs within the same VM instance, after five discarded warm-up runs; this is a standard technique to mitigate the variability of JVM overheads, including JIT compilation [22].

Table 1. Mean execution time and memory overheads for verification.

Benchmark	Time		Memory		Tasks	Gets/ms	Sets/ms
	Baseline (s)	Overhead	Baseline (MB)	Overhead			
Conway	4.43	1.01×	314.06	0.98×	101	361.74	361.58
Heat	5.06	1.00×	51.28	1.00×	51	98.92	98.89
QSort	3.14	0.98×	115.92	1.08×	786035	250.13	250.12
Randomized	2.99	0.98×	6.90	1.01×	2535	2.52	1.51
Sieve	1.24	2.07×	140.39	1.18×	9594	37 285.39	74 547.63
SmithWaterman	4.26	1.10×	444.44	1.40×	569857	536.08	401.53
Strassen	0.58	1.04×	116.69	1.00×	58998	102.20	544.11
StreamCluster	14.48	1.19×	91.02	0.95×	33	39.27	274.89
StreamCluster2	16.81	0.99×	89.96	0.99×	33	17.92	125.93
Geometric Mean Overhead		1.12×	1.06×				

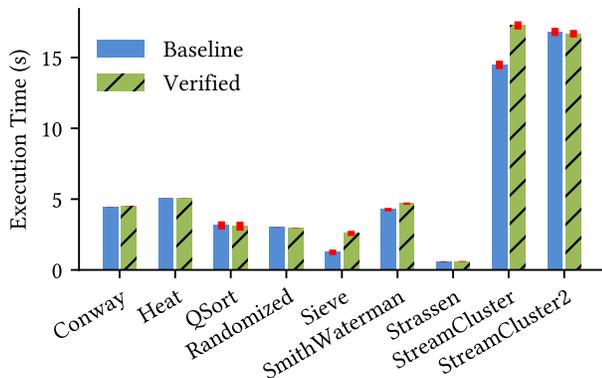
**Figure 1.** Execution times for each benchmark showing the mean with a 95% confidence interval (red).

Table 1 gives the unverified baseline measurements for each program and the overhead factors introduced by the verifiers. The table also gives the geometric mean of overheads across all benchmarks. There is an overall factor of 1.12× in execution time and 1.06× in memory usage. The total number of tasks in the program and the average rates of promise get and set actions per millisecond (with respect to the baseline execution time) are also reported. Figure 1 represents the execution times of each benchmark, showing the 95% confidence interval. The low overheads indicate that our deadlock detection algorithm does not introduce serialization bottlenecks.

The overall execution time overheads are within 1.1× for each of Conway, Heat, QSort, Randomized, SmithWaterman, Strassen, and StreamCluster2. The same is true of the memory overheads for this subset of benchmarks, excepting SmithWaterman. In many cases, the verified run narrowly out-performs the baseline, which can be attributed to perturbations in scheduling and garbage collection.

It is worth noting that the execution overhead for Sieve is in excess of 2×. Sieve has the single highest rate of get operations by an order of magnitude (over 37,000, compared to SmithWaterman’s 536). The Sieve program requires almost 9594 tasks to be live simultaneously, each waiting on the next, with the potential to form very long dependence chains for Algorithm 2 to traverse.

We can also remark on the 1.4× memory overhead in SmithWaterman. Unlike Conway, Heat, Sieve, and both of the StreamCluster benchmarks, in which most promises are allocated by the same task that fulfills them, SmithWaterman (and Randomized) allocates all promises in the root task and moves them later. In maintaining the owned lists in Algorithm 1, one can make trade-offs between speed and space. Our implementation favors speed, so instead of literally removing a promise p from t .owned in lines 11 and 25, we simply rely on the fact that p .owner $\neq t$ anymore to detect that p should no longer be counted in line 16.

For comparison with deadlock verification in other settings, the Armus tool [14] can identify barrier deadlocks as soon as they occur, with execution overheads of up to 1.5× on Java benchmarks. Our benchmark results represent an acceptable performance overhead when one desires runtime-identifiable deadlocks and omitted sets with attributable blame.

7 Related Work

Task-parallel programming is prevalent in a variety of languages and libraries. Multilisp [28] is one of the earliest languages with futures, a mechanism for parallel execution of functional code. Fork-join parallelism is employed in Cilk [20], and the more general async-finish with futures model was introduced in X10 [12]. Habanero-Java [10] modernized X10 as an extension to Java and, later, as a Java library, HJlib [33]; this language incorporates additional synchronization primitives, such as the phaser [50] and the data-driven future [51], which is a promise-like mechanism. Many other

languages, libraries, and extensions include spawning and synchronizing facilities, whether for threads or lightweight tasks, including Chapel [11], Fortress [3], OpenMP [47], Intel Threading Building Blocks [34], Java [24], C++17 [36], and Scala [27].

The promise, as we define it, can be traced back to the I-structures of the Id language [4], which are also susceptible to deadlock. Cells of data in an I-structure are uninitialized when allocated, may be written to at most once, and support a read operation that blocks until the data is available.

The classic definition of a deadlock is found in Isloor and Marsland [35], which is primarily concerned with concurrent allocation of limited resources. Solutions in this domain fall into the three categories of Coffman: static prevention, runtime detection, and run-time avoidance [13].

We consider logical deadlocks, which are distinct from resource deadlocks in that there is an unresolvable cyclic dependence among computational results. Solutions in the logical deadlock domain include techniques that dynamically detect cycles [29, 30, 37, 38, 40, 54], that raise alarms upon the formation or possible formation of cycles [1, 6, 14, 15, 23, 55], that statically check for cycles through analysis [42, 45, 57] or through type systems [7, 52], or that preclude cycles by carefully limiting the blocking synchronization semantics available to the programmer, either statically or dynamically [10, 12, 15, 50, 55]. The present work includes a dynamic, precise cycle detection algorithm, enabled only by the introduction of a structured ownership semantics on the otherwise unrestricted promise primitive.

Futures are a special case of promises where each one is bound to a task whose return value is automatically put into the promise. Transitive Joins [55] and its predecessor, Known Joins [15], are policies with runtime algorithms for deadlock detection on futures. They are, in general, not applicable to promises. These two techniques impose additional structure on the synchronization pattern by limiting the set of futures that a given task may await at any given time.

Recent work identifies the superior flexibility of promises over futures with the problematic loss of a guarantee that they will be fulfilled and develops a *forward* construct as a middle-ground [18]. Forwarding can be viewed in terms of delegating promise ownership, but it is restricted in that 1) it moves only a single promise into a new task, and 2) in particular, it moves only the implicit promise that is used to retrieve a task's return value. In terms of futures, forwarding amounts to re-binding a future to new task.

Other synchronization constructs benefit from similar annotations to the one we have proposed for promises. This includes event-driven programming models where events have similar semantics to that of promises. JavaScript, though a single-threaded language, still uses an asynchronous task model to schedule callbacks on an event loop [39], and could benefit from our approach. Likewise, our approach is directly applicable to multithreaded execution models, such

as Concurrent Collections [8] and the Open Community Runtime [41], that use event-driven execution as a fundamental primitive. As another example, the MPI blocking receive primitive must name the sending task; from this information a waits-for graph for deadlock detection can be directly constructed [29]. In addition, nonblocking communications in MPI use `MPI_Request` objects in a manner similar to promises, and the `MPI_Wait` operation akin to the get operation on promises.

Languages with barriers and phasers sometimes require the participating tasks to *register* with the construct [50]. Notably, this kind of registration is absent from the Java API, which is problematic for the Armus deadlock tool [14]. In that work, registration annotations had to be added to the Java benchmarks in order to apply the Armus methodology.

In this work, we considered programs which only use promises for blocking synchronization, and we constrained ownership transfer to occur only when a task is spawned. Since a promise can have multiple readers or no readers at all, it is not possible in principle to use one promise to synchronize the ownership hand-off of a second promise between two existing tasks. We cannot guarantee that the receiving task exists and is unique. In future work, one could consider a slightly higher abstraction in the form of a pair of promises acting like a rendezvous, which is a primitive in languages like Ada and Concurrent C [21]. Such a synchronization pattern could be leveraged to hand off promise ownership since there would be a guaranteed single receiving task.

The Rust language incorporates affine types in its move semantics to ensure that certain objects have at most one extant reference at all times [49]. The movement of promise ownership from one task to another and the obligation to fulfill each promise exactly once may be expressible at compile time through the use of a linear type system, which restricts references to exactly one instance.

8 Conclusion

We have introduced an ownership mechanism for promises, whereby each task is responsible for ensuring that all of its owned promises are fulfilled. This mechanism makes it possible to identify a bug, called the omitted set, at runtime when the bug actually occurs and to report which task is to blame for the error. The ownership mechanism also makes it meaningful, for the first time, to formally define, discuss, and detect deadlock cycles among tasks synchronizing with promises. Such a bug is now detectable as soon as the cycle forms.

In our approach, any code that spawns a new asynchronous task must name the promises which are to be transferred to the new task. The programmer must already be aware of this critical information in order to even informally reason about omitted set and deadlock bugs. We now ask that it be explicitly notated in the code.

We provided an algorithm to check for compliance with the ownership policy at runtime, thereby detecting omitted sets, as well as an algorithm for detecting deadlock cycles using ownership information. Both types of bug are detected when they occur, not after-the-fact. Our deadlock detector is provably precise and correct under a weak memory model and we described how to obtain this correct behavior under the TSO, Java, and C++ memory models. Every alarm corresponds to a true deadlock and every deadlock results in an alarm. Experimental evaluation demonstrates that our lock-free approach to deadlock detection exhibits low execution time and memory overheads relative to an uninstrumented baseline.

Acknowledgments

This work is supported by the National Science Foundation under Collaborative Grant No. 1822919 and Graduate Research Fellowship Grant No. 1650044.

References

- [1] Rahul Agarwal and Scott D. Stoller. 2006. Run-Time Detection of Potential Deadlocks for Programs with Locks, Semaphores, and Condition Variables. In *Proc. 2006 Worksh. on Parallel and Distributed Systems: Testing and Debugging (PADTAD '06)*. ACM, New York, NY, 51–60.
- [2] Dongie Agnir. 2019. *Call exceptionOccurred in case of stream error*. Amazon Web Services. Retrieved 30 July 2020 from <https://github.com/aws/aws-sdk-java-v2/commit/bfdd0d2063>
- [3] Eric Allen, David Chase, Joe Hallett, Victor Luchangco, Jan-Willem Maessen, Sukyoung Ryu, Guy L. Steele Jr., and Sam Tobin-Hochstadt. 2008. *The Fortress Language Specification*. Sun Microsystems, Inc.
- [4] Arvind, Rishiyur S. Nikhil, and Keshav K. Pingali. 1989. I-Structures: Data Structures for Parallel Computing. *ACM Trans. Program. Lang. Syst.* 11, 4 (1989), 598–632.
- [5] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In *Proc. 17th Int'l. Conf. on Parallel Architectures and Compilation Techniques (PACT '08)*. ACM, New York, NY, 72–81.
- [6] Gérard Boudol. 2009. A Deadlock-Free Semantics for Shared Memory Concurrency. In *Proc. 6th Int'l. Coll. on Theoretical Aspects of Computing (ICTAC '09)*. Springer, Berlin, Germany, 140–154.
- [7] Chandrasekhar Boyapati, Robert Lee, and Martin Rinard. 2002. Ownership Types for Safe Programming: Preventing Data Races and Deadlocks. In *Proc. 17th ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '02)*. ACM, New York, NY, 211–230.
- [8] Zoran Budimlić, Michael Burke, Vincent Cavé, Kathleen Knobe, Geoff Lowney, Ryan Newton, Jens Palsberg, David Peixotto, Vivek Sarkar, Frank Schlimbach, and Saĝnak Taşlılar. 2010. Concurrent Collections. *Sci. Program.* 18, 3–4 (Aug. 2010), 203–217. <https://doi.org/10.1155/2010/521797>
- [9] John Burkardt. 2016. *HEAT_MPI: Solve the 1D Time Dependent Heat Equation using MPI*. Florida State University. Retrieved 13 August 2020 from https://people.sc.fsu.edu/~jburkardt/cpp_src/heat_mpi/heat_mpi.html
- [10] Vincent Cavé, Jisheng Zhao, Jun Shirako, and Vivek Sarkar. 2011. Habanero-Java: The New Adventures of Old X10. In *Proc. 9th Int'l. Conf. on Principles and Practice of Programming in Java (PPPJ '11)*. ACM, New York, NY, 51–61.
- [11] Bradford L. Chamberlain, David Callahan, and Hans P. Zima. 2007. Parallel Programmability and the Chapel Language. *Int'l. Journal of High Performance Computing Applications* 21, 3 (2007), 291–312.
- [12] 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 *Proc. 20th ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '05)*. ACM, New York, NY, 519–538.
- [13] E. G. Coffman, M. Elphick, and A. Shoshani. 1971. System Deadlocks. *ACM Comput. Surv.* 3, 2 (1971), 67–78.
- [14] Tiago Cogumbreiro, Raymond Hu, Francisco Martins, and Nobuko Yoshida. 2018. Dynamic Deadlock Verification for General Barrier Synchronisation. *ACM Trans. Program. Lang. Syst.* 41, 1, Article 1 (2018), 38 pages.
- [15] Tiago Cogumbreiro, Rishi Surendran, Francisco Martins, Vivek Sarkar, Vasco T. Vasconcelos, and Max Grossman. 2017. Deadlock Avoidance in Parallel Programs with Futures: Why Parallel Tasks Should Not Wait for Strangers. *Proc. ACM Program. Lang.* OOPSLA, Article 103 (2017), 26 pages.
- [16] Alex Crichton. 2020. *futures::channel*. Retrieved 30 July 2020 from <https://docs.rs/futures/0.3.5/futures/channel>
- [17] Alejandro Duran, Xavier Teruel, Roger Ferrer, Xavier Martorell, and Eduard Ayguade. 2009. Barcelona OpenMP Tasks Suite: A Set of Benchmarks Targeting the Exploitation of Task Parallelism in OpenMP. In *Proc. 2009 Int'l. Conf. on Parallel Processing (ICPP '09)*. IEEE Computer Society, Washington, DC, 124–131.
- [18] Kiko Fernandez-Reyes, Dave Clarke, Elias Castegren, and Huu-Phuc Vo. 2018. Forward to a Promising Future. In *Int'l. Conf. on Coordination Models and Languages*. Springer, Cham, Switzerland, 162–180.
- [19] Rhys S. Francis and Linda J.H. Pannan. 1992. A parallel partition for enhanced parallel QuickSort. *Parallel Comput.* 18, 5 (1992), 543–550.
- [20] Matteo Frigo, Charles E. Leiserson, and Keith H. Randall. 1998. The Implementation of the Cilk-5 Multithreaded Language. In *Proc. ACM SIGPLAN 1998 Conf. on Programming Language Design and Implementation (PLDI '98)*. ACM, New York, NY, 212–223.
- [21] Narain H. Gehani and William D. Roome. 1988. Rendezvous Facilities: Concurrent C and the Ada Language. *IEEE Transactions on Software Engineering* 14, 11 (1988), 1546–1553.
- [22] Andy Georges, Dries Buytaert, and Lieven Eeckhout. 2007. Statistically Rigorous Java Performance Evaluation. In *Proc. 22nd ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '07)*. ACM, New York, NY, 57–76.
- [23] Prodromos Gerakios, Nikolaos Papaspyrou, Konstantinos Sagonas, and Panagiotis Vekris. 2011. Dynamic Deadlock Avoidance in Systems Code Using Statically Inferred Effects. In *Proc. 6th Worksh. on Programming Languages and Operating Systems (PLOS '11)*. ACM, New York, NY, Article 5, 5 pages.
- [24] B. Goetz, T. Peierls, J. Bloch, J. Bowbeer, D. Lea, and D. Holmes. 2006. *Java Concurrency in Practice*. Pearson Education, London, England.
- [25] Google. 2014. *Go 1.3 Release Notes*. Retrieved 30 July 2020 from <https://golang.org/doc/go1.3>
- [26] Habanero Extreme Scale Software Research Lab. 2020. *Habanero-C Library*. Retrieved 15 February 2020 from <https://github.com/habanero-ric/hclib>
- [27] Philipp Haller, Aleksandar Prokopec, Heather Miller, Viktor Klang, Roland Kuhn, and Vojin Jovanovic. 2012. *Futures and Promises*. École polytechnique fédérale de Lausanne. Retrieved 24 August 2020 from <https://docs.scala-lang.org/overviews/core/futures.html>
- [28] Robert H. Halstead, Jr. 1985. Multilisp: A Language for Concurrent Symbolic Computation. *ACM Trans. Program. Lang. Syst.* 7, 4 (1985), 501–538.
- [29] Tobias Hilbrich, Bronis R. de Supinski, Martin Schulz, and Matthias S. Müller. 2009. A Graph Based Approach for MPI Deadlock Detection. In *Proc. 23rd Int'l. Conf. on Supercomputing (ICS '09)*. ACM, New York, NY, 296–305.

- [30] Tobias Hilbrich, Joachim Protze, Martin Schulz, Bronis R. de Supinski, and Matthias S. Müller. 2012. MPI Runtime Error Detection with MUST: Advances in Deadlock Detection. In *Proc. Int'l. Conf. on High Performance Computing, Networking, Storage and Analysis (SC '12)*. IEEE Computer Society, Los Alamitos, CA, Article 30, 11 pages.
- [31] Oliver Hsu. 2019. *S3: FileAsyncResponseTransformer future does not complete when checksum error occurs*. Amazon Web Services. Retrieved 30 July 2020 from <https://github.com/aws/aws-sdk-java-v2/issues/1279>
- [32] Shams Imam and Vivek Sarkar. 2014. Cooperative Scheduling of Parallel Tasks with General Synchronization Patterns. In *European Conference on Object-Oriented Programming (ECOOP '14)*. Springer, Berlin, Germany, 618–643.
- [33] Shams Imam and Vivek Sarkar. 2014. Habanero-Java Library: A Java 8 Framework for Multicore Programming. In *Proc. 2014 Int'l. Conf. on Principles and Practices of Programming on the Java Platform: Virtual Machines, Languages, and Tools (PPPJ '14)*. ACM, New York, NY, 75–86.
- [34] Intel. 2020. *Intel Threading Building Blocks Developer Guide*. Intel.
- [35] S. Sreekanth Isloor and T. Anthony Marsland. 1980. The Deadlock Problem: An Overview. *Computer* 13, 9 (1980), 58–78.
- [36] ISO. 2017. *ISO/IEC 14882:2017: Programming Languages — C++*. International Organization for Standardization, Geneva, Switzerland.
- [37] Bettina Krammer, Tobias Hilbrich, Valentin Himmler, Blasius Czink, Kiril Dichev, and Matthias S. Müller. 2008. MPI Correctness Checking with Marmot. In *Tools for High Performance Computing*. Springer, Berlin, Germany, 61–78.
- [38] Bettina Krammer, Matthias S. Müller, and Michael M. Resch. 2004. MPI Application Development Using the Analysis Tool MARMOT. In *Proc. Int'l. Conf. on Computational Science (ICCS '04)*. Springer, Berlin, Germany, 464–471.
- [39] Matthew C. Loring, Mark Marron, and Daan Leijen. 2017. Semantics of Asynchronous JavaScript. In *Proc. 13th ACM SIGPLAN Int'l. Symp. on Dynamic Languages (DLS '17)*. ACM, New York, NY, 51–62.
- [40] Glenn Luecke, Hua Chen, James Coyle, Jim Hoekstra, Marina Kraeva, and Yan Zou. 2003. MPI-CHECK: A Tool for Checking Fortran 90 MPI Programs. *Concurrency and Computation: Practice and Experience* 15, 2 (2003), 93–100.
- [41] T. G. Mattson, R. Cledat, V. Cavé, V. Sarkar, Z. Budimčić, S. Chatterjee, J. Fryman, I. Ganev, R. Knauerhase, Min Lee, B. Meister, B. Nickerson, N. Pepperling, B. Seshasayee, S. Tasirlar, J. Teller, and N. Vrvilo. 2016. The Open Community Runtime: A runtime system for extreme scale computing. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*. 1–7. <https://doi.org/10.1109/HPEC.2016.7761580>
- [42] Mayur Naik, Chang-Seo Park, Koushik Sen, and David Gay. 2009. Effective Static Deadlock Detection. In *Proc. 31st Int'l. Conf. on Software Engineering (ICSE '09)*. IEEE Computer Society, Washington, DC, 386–396.
- [43] Varun Nandi. 2019. *Don't call onComplete after onError in ChecksumValidatingSubscriber#onComplete method which results in NPE*. Amazon Web Services. Retrieved 30 July 2020 from <https://github.com/aws/aws-sdk-java-v2/commit/eaecf99a02>
- [44] Mozilla Developer Network. 2020. *Promise – JavaScript | MDN*. Retrieved 5 August 2020 from https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/Promise
- [45] Nicholas Ng and Nobuko Yoshida. 2016. Static Deadlock Detection for Concurrent Go by Global Session Graph Synthesis. In *Proc. 25th Int'l. Conf. on Compiler Construction (CC '16)*. ACM, New York, NY, 174–184.
- [46] Joachim Niehren, Jan Schwinghammer, and Gert Smolka. 2005. A Concurrent Lambda Calculus with Futures. In *Int'l. Worksh. on Frontiers of Combining Systems (FroCoS '05)*. Springer, Berlin, Germany, 338–356.
- [47] OpenMP Architecture Review Board 2018. *OpenMP Application Programming Interface*. OpenMP Architecture Review Board.
- [48] Oracle. 2020. *CompletableFuture (Java SE 14 & JDK 14)*. Retrieved 30 July 2020 from <https://docs.oracle.com/en/java/javase/14/docs/api/java.base/java/util/concurrent/CompletableFuture.html>
- [49] Rust Lang. 2020. *The Rust Programming Language*. Retrieved 12 August 2020 from <https://doc.rust-lang.org/1.8.0/book/index.html>
- [50] Jun Shirako, David M. Peixotto, Vivek Sarkar, and William N. Scherer, III. 2008. Phasers: A Unified Deadlock-Free Construct for Collective and Point-to-Point Synchronization. In *Proc. 22nd Ann. Int'l. Conf. on Supercomputing (ICS '08)*. ACM, New York, NY, 277–288.
- [51] Saġnak Taşirlar and Vivek Sarkar. 2011. Data-Driven Tasks and Their Implementation. In *2011 Int'l. Conf. on Parallel Processing (ICPP '11)*. ACM, New York, NY, 652–661.
- [52] Vasco T. Vasconcelos, Francisco Martins, and Tiago Cogumbreiro. 2010. Type Inference for Deadlock Detection in a Multithreaded Polymorphic Typed Assembly Language. In *Proc. 2nd Int'l. Worksh. on Programming Language Approaches to Concurrency and Communication-eentric Software (Electronic Proceedings in Theoretical Computer Science, Vol. 17)*, Alastair R. Beresford and Simon Gay (Eds.). Open Publishing Association, 95–109.
- [53] Philippe Virouleau, Pierrick Brunet, François Broquedis, Nathalie Fumento, Samuel Thibault, Olivier Aumage, and Thierry Gautier. 2014. Evaluation of OpenMP Dependent Tasks with the KASTORS Benchmark Suite. In *Int'l. Worksh. on OpenMP (IWOMP '14)*. Springer, Cham, Switzerland, 16–29.
- [54] Anh Vo, Ganesh Gopalakrishnan, Robert M. Kirby, Bronis R. de Supinski, Martin Schulz, and Greg Bronevetsky. 2011. Large Scale Verification of MPI Programs Using Lamport Clocks with Lazy Update. In *Proc. 2011 Int'l. Conf. on Parallel Architectures and Compilation Techniques (PACT '11)*. IEEE Computer Society, Washington, DC, 330–339.
- [55] Caleb Voss, Tiago Cogumbreiro, and Vivek Sarkar. 2019. Transitive Joins: A Sound and Efficient Online Deadlock-Avoidance Policy. In *Proc. 24th ACM SIGPLAN Symp. on Principles and Practice of Parallel Programming (PPoPP '19)*. ACM, New York, NY, 378–390.
- [56] Aaron Weeden. 2012. *Parallelization: Conway's Game of Life*. The Shodor Education Foundation. Retrieved 13 August 2020 from <http://www.shodor.org/petascale/materials/UPModules/GameOfLife/>
- [57] Amy Williams, William Thies, and Michael D. Ernst. 2005. Static Deadlock Detection for Java Libraries. In *Proc. 19th European Conf. on Object-Oriented Programming (ECOOP '05)*. Springer, Berlin, Germany, 602–629.

A Artifact Description

The evaluated artifact which accompanied this work consists of the Java implementation of our algorithms, the benchmark programs, and the means of conducting the experiment described in section 6.

A.1 Implementation

We created a simple runtime environment for executing task-parallel Java programs in which promise ownership is managed by an implementation of Algorithm 1 and every promise *get* operation invokes Algorithm 2 to determine if it will create a deadlock. The implementation is mindful of the Java memory model, following the requirements laid out in section 5.1.

Section 6.1 and section 6.2 discuss some practical considerations in our implementation regarding how we can encapsulate promise movement specifications in an object-oriented way and how we approach exception handling in a parallel program.

A.2 Programs

Included in the artifact are twelve short test programs to verify that the implementation raises an alarm for deadlocks, omitted sets, and other mismanagement of promise ownership, while not raising alarms on well behaved programs. These programs serve to sanity-check the correctness of the implementation.

Also included are the nine benchmark programs, as described in section 6.3. When run under conditions comparable to our experimental setup (under OpenJDK 11 on a 16-core AMD Opteron CPU), it is expected that the benchmarks produce performance results similar in character to table 1 and fig. 1.